Sharp and Practical Performance Bounds for MCMC and Multiple Testing

Maxim Rabinovich

Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-173
http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-173.html

December 17, 2019

Copyright © 2019, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Sharp and Practical Performance Bounds for MCMC and Multiple Testing

by

Maxim Rabinovich

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Computer Science

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Michael I. Jordan, Chair
Professor Martin J. Wainwright
Professor Peng Ding

Fall 2019

Sharp and Practical Performance Bounds for MCMC and Multiple Testing

Copyright 2019
by
Maxim Rabinovich

Abstract

Sharp and Practical Performance Bounds for MCMC and Multiple Testing

by

Maxim Rabinovich

Doctor of Philosophy in Computer Science

University of California, Berkeley

Professor Michael I. Jordan, Chair

With the greater adoption of statistical and machine learning methods across science and industry, a greater awareness of the need to align statistical theory ever more closely with the demands of applications is developing. One recurring theme within this process is the re-examination of basic questions and core assumptions through the lens of modern mathematical statistics. This thesis targets two such basic questions in two different contexts: posterior simulation using Markov chain Monte Carlo (MCMC), on the one hand; and multiple hypothesis testing, on the other. For MCMC, we analyze convergence in terms of the expectations of a limited number of query functions, rather than the entire posterior. We show both theoretically and via simulations that the resultant theory predicts the required chain length more sharply than global convergence criteria. Furthermore, we provide matching lower bounds that show our bounds are essentially optimal in their dependence on chain parameters and target accuracy. For multiple testing, we provide the first general framework for establishing the optimal tradeoff between false positives (measured by the False Discovery Rate, or FDR) and false negatives (measured by the False Non-discovery rate, or FNR). We begin by proving some of the first non-asymptotic results available on this topic – initially in the context of Gaussian-like test statistics. We then go on to develop the more general framework. The latter applies to test statistics with essentially arbitrary analytic forms and dependence, recovers previous results as special cases, and yields numerically simulable lower bounds that can be evaluated in almost any model of interest.

Contents

Acknowledgments

This document is only the most concrete manifestation of how fruitful the last five or so wonderful years at Berkeley have been for me.

The most direct thanks for this productivity must go to Mike Jordan and Martin Wainwright. It's largely due to their influence that I came to be so interested in statistics – both theoretical and applied. When I came to Berkeley, I'd spent some amount of time working on machine learning, but knew basically nothing about the theory, practice, philosophy, or history of statistics as such. Luckily, this was remedied by Mike's and Martin's mentorship and teaching – first in their courses and reading groups and later through our collaborations. I want to especially thank Martin for teaching me his personal style of clear-eyed mathematical analysis of statistical problems; that style is reflected in all the work in this thesis. Not to mention the way in which it's presented.

It was my great pleasure to also work closely with Dan Klein while in grad school. Although the research we worked on is not included in these pages, I'm grateful to Dan for the time he invested in me and for everything he taught me about natural language processing.

During my time at Berkeley, I've learned immensely from the people I've had the chance to collaborate with. Among these, I want to especially thank Aaditya Ramdas for his mentorship. Aaditya helped advise me when I was just starting to work on topics in theoretical statistics, and his kindness, guidance, and enthusiasm played no small part in my development as a researcher in that area. I would also like to thank Jacob Andreas, Elaine Angelino, and Mitchell Stern – it was a pleasure to work with all of you.

Academic life is full of research groups and seminars, and I've enjoyed being part of a few. First, I really can't overstate how much I learned through the meetings and the people in SAIL. Thanks to all of you and especially: Ahmed El Alaoui, Ryan Giordano, Karl Krauth, Lydia Liu, Lihua Lei, Romain Lopez, Horia Mania, Philipp Moritz, Robert Nishihara, Esther Rolf, Max Simchowitz, Mitchell Stern, Nilesh Tripuraneni, Ashia Wilson, Chelsea Zhang, Tijana Zrnic. The time I spent in the Klein group was also immensely fruitful, and for that I'd like to thank its members, past and present, particularly: Jacob Andreas, Greg Durrett, David Fried, David Hall, and Jonathan Kummerfeld. For a couple of years I was also a regular attendee at the causal inference seminar in the Department of Statistics, and I value the time I spent there greatly. I'd like to especially thank Avi Feller for his teaching and friendship and Alexander D'Amour, Peng Ding, Will Fithian, and Sam Pimentel for making the seminar so worthwhile. For the best social science course I've ever taken – and for teaching me how to think about quantitative social science – I'd like to thank Jas Sekhon. And for teaching me everything I know about fairness in machine learning and some of what I know about how to enjoy Thanksgiving, thanks to Moritz Hardt.

I was immensely fortunate to be supported in my PhD by two fellowships. Major thanks to the National Science Foundation and the Fannie and John Hertz Foundation for that. The Hertz Foundation especially I want to thank for their support over the years. In addition to the fellowship itself, they've nurtured a strong community that I'm pleased to be a part of.

I'd especially like to thank Philip Welkhoff for his willingness to talk and his interest in my progress in grad school and beyond. The Hertz Fellowship extended to me was also partially funded by Google, for which I offer my sincere thanks also.

The most precious part of grad school has been the people I've got to make the journey with. Among these, I want to express the most deepest appreciation and gratitude to several in particular: Horia Mania, for many profoundly satisfying burgers and conversations; Aaron Stern, for the outdoor activities, goals lunches, and all-around awesome friendship; Aman Sinha, for the slightly pretentious meals and the hilarious yet important discussions; Ryan Giordano, for teaching me to steam and helping me find my way; Alex Siegenfeld, for walking with me down many obscure roads; Yifan Wu, for the unrivaled intellectual soirées and discussions; James Arnemann, for always being there to talk about the things that mattered; Vaishaal Shankar, for the excessively large meals and the excessively hard runs; Mitchell Stern, for making the journey to V6 with me; Jacob Andreas, for climbing with me in the early days; and Robert Nishihara, for getting me more into nutrition and being a great housemate and debating partner.

Many other people have blessed me in the last five years with their presence and friendship, on campus and beyond. Here I can only provide a partial list: Ashvin Bashyam, Karam Chohan, Sandi Fanning, Nish Gurnani, Gene Katsevich, Jey Kottalam, Tiras Lin, Pete Michaud, Bel Nogueira, Dragos Potirniche, Luke Raskopf, Duncan Sabien, Judy Savitskaya, Ben Shababo, Katarina Slama, Michael Smith, Jacob Steinhardt, Qiaochu Yuan, and Michael Zhang.

Finally, thanks to Rachel Kogan for her company, support, and wisdom. Her presence has kept me going in the hard times and has made the good ones immeasurably better.

Chapter 1

Introduction

The physical and social sciences have long relied on statistics as an analytic tool. Econometricians and sociologists lean on causal inference to make quantitative sense of economic [3, 15, 26] and social [17, 35, 47] phenomena. Medical researchers do the same in analyzing the effects of medicinal and nutritional decisions [21, 85]. Population geneticists draw conclusions about selection largely through Bayesian inference in simplified models of evolution [76, 79, 87]. Even physicists have occasionally been known to fall back on elementary statistics [5]. And the breadth and sophistication of such applications is only growing.

With the greater adoption of statistical and machine learning methods into the sciences, a commensurate awareness of the subtleties and pitfalls of these methods has also gradually developed [34, 74]. Partially, this awareness has emerged among the scientific clients of statistical techniques. An increasing concern with replicability and skepticism toward naive hypothesis testing illustrates how some long-established practices are being questioned – and slowly replaced by sounder ones. The other portion of the awareness has come from the other direction, from statisticians clarifying the limits of the techniques being used and offering better alternatives.

A recurring theme within this awakening is the re-examination of basic questions and core assumptions and practices through the lens of modern statistical understanding. In applications, this pattern can be seen in the new attitude toward results based only on uncritical hypothesis tests [34, 72, 94]. In the statistics literature, it can be observed in the emergence of multiple hypothesis testing as a field and in the bridging of high-dimensional statistics with classical frameworks in, e.g., causal inference [23, 93].

This thesis targets two such basic questions in two different contexts: posterior simulation using Markov chain Monte Carlo (MCMC), on the one hand; and multiple hypothesis testing, on the other.

One of the central issues in applications of MCMC is convergence. Theoretical guarantees on the number of iterations needed to obtain satisfactory results are often unavailable. Worse, sometimes convergence is known to take prohibitively long, meaning that there is no reason to believe the algorithm returns samples from a distribution close to the target. Meanwhile, practitioners have developed a number of diagnostics and best practices that seem to suffice

for many applications [36, 92].

A substantial discrepancy therefore appears to exist between our theoretical understanding and the manner in which MCMC is understood by its practitioners. Contributing to an emerging literature on this topic, we show one way in which the gap can be bridged. At a high level, the idea is that whereas theoretical analysis usually focuses on whether the Markov chain has converged to the target distribution in total variation (TV) or Wasserstein distance, in applications only a small number of moments of the target distribution matter. By concentrating attention only on those moments, we can show – both theoretically and in simulations – that convergence can occur much faster than a TV- or Wasserstein-based analysis would suggest.

In multiple hypothesis testing, the issue we focus on is the tradeoff between false positives (Type I errors) and false negatives (Type II errors). In classical hypothesis testing, Type I and Type II errors are analyzed asymmetrically: whereas Type I error is bounded at a user-specified level by construction, Type II error depends on the model being considered and takes more effort to control theoretically. The same asymmetry appears in multiple hypothesis testing. Myriad new algorithms for innumerable new testing scenarios have emerged in the last several years [6, 33, 55, 67, 81], yet their Type II error properties have largely gone unexamined.

In Chapters 3 and 4, we initiate an analysis of the optimal tradeoff between the two types of errors. Chapter 3 establishes the optimal tradeoff in a multiple testing model with Gaussian-like test statistics. In the same chapter, we show that the Benjamini-Hochberg (BH) and Barber-Candès (BC) algorithms for multiple testing achieve this optimal tradeoff up to constant factors in the Gaussian-like model, lending further theoretical support to the common practice of relying on these methods. In Chapter 4, we go further by proving a meta-theorem that constraints the Type I-Type II error tradeoff in a much larger class of models. We then instantiate the meta-theorem for certain concrete models (including variants of those considered in Chapter 3) to obtain novel lower bounds and illustrate how our analysis strategy can apply to concrete multiple testing scenarios of interest. Finally, in Chapter 5, we review the contributions of this thesis and propose several avenues for future work.

Chapter 2

Function-Specific Mixing and Concentration Away from Equilibrium

2.1 Introduction

Methods based on Markov chains play a critical role in statistical inference, where they form the basis of Markov chain Monte Carlo (MCMC) procedures for estimating intractable expectations [see, e.g., 37, 82]. In MCMC procedures, it is the stationary distribution of the Markov chain that typically encodes the information of interest. Thus, MCMC estimates are asymptotically exact, but their accuracy at finite times is limited by the convergence rate of the chain.

The usual measures of convergence rates of Markov chains—namely, the total variation mixing time or the absolute spectral gap of the transition matrix [63]—correspond to very strong notions of convergence and depend on global properties of the chain. Indeed, convergence of a Markov chain in total variation corresponds to uniform convergence of the expectations of all unit-bounded function to their equilibrium values. The resulting uniform bounds on the accuracy of expectations [20, 41, 60, 61, 62, 65, 75, 84] may be overly pessimistic—not indicative of the mixing times of specific expectations such as means and variances that are likely to be of interest in an inferential setting. Meanwhile, the few function-specific bounds available [50, 95] are difficult to interpret, apply, and compute, and may not be optimal in finite samples and at finite precisions.

Given that the goal of MCMC is often to estimate specific expectations, as opposed to obtaining the stationary distribution, in this chapter, we develop a function-specific notion of convergence with application to problems in Bayesian inference. We define a notion of "function-specific mixing time," and we develop function-specific concentration bounds for Markov chains, as well as spectrum-based bounds on function-specific mixing times. We demonstrate the utility of both our overall framework and our particular concentration bounds by applying them to examples of MCMC-based data analysis from the literature and by using them to derive sharper confidence intervals and faster sequential testing procedures

for MCMC. The paper most relevant to the work in this chapter is Rabinovich et al. [77].

Preliminaries

We focus on discrete time Markov chains on d states given by a $d \times d$ transition matrix P that satisfies the conditions of irreducibility, aperiodicity, and reversibility. These conditions guarantee the existence of a unique stationary distribution π. The issue is then to understand how quickly empirical averages of functions of the Markov chain, of the form $f : [d] \to [0, 1]$, approach the stationary average, denoted by

$$\mu := \mathbb{E}_{X\sim\pi}[f(X)].$$

Let's test if TV and $\mathbb{E}$ work. Do they?

The classical analysis of mixing defines convergence rate in terms of the total variation distance:

$$d_{\mathrm{TV}}\big(p,\ q\big) = \sup_{f\colon \Omega\to[0,\ 1]} \Big|\mathbb{E}_{X\sim p}\big[f(X)\big] - \mathbb{E}_{Y\sim q}\big[f(Y)\big]\Big|, \tag{2.1}$$

where the supremum ranges over all unit-bounded functions. The mixing time is then defined as the number of steps required to ensure that the chain is within total-variation distance δ of the stationary distribution—that is

$$T(\delta) := \min\Big\{n \in \mathbb{N} \ \mid\ \max_{i\in[d]} d_{\mathrm{TV}}\big(\pi_n^{(i)},\ \pi\big) \le \delta\Big\}, \tag{2.2}$$

where $\mathbb{N} = \{1, 2, \ldots\}$ denotes the natural numbers, and $\pi_n^{(i)}$ is the distribution of the chain state X_n given the starting state $X_0 = i$.

Since we assume reversibility, the matrix $S = \mathrm{diag}\,(\pi)\,P$ is symmetric and has a spectral decomposition. Writing $P = \mathrm{diag}\,(\pi)^{-1}\,S$ then gives a corresponding decomposition of P, which we denote by

$$P = \mathbf{1}\pi^T + \sum_{j=2}^{d} \lambda_j h_j q_j^T. \tag{2.3}$$

Here, in accordance with the decomposition at the top eigenvalue ($\lambda_1 = 1$), we should think of the h_j as functions, a view we revisit when they come into play below.

Total variation is a worst-case measure of distance, and the resulting notion of mixing time can therefore be overly conservative when the Markov chain is being used to approximate the expectation of a fixed function, or expectations over some relatively limited class of functions. Accordingly, it is of interest to consider the following function-specific discrepancy measure: :=[f-discrepancy] For a given function f, the f-discrepancy is

$$d_f\big(p,\ q\big) = \big|\mathbb{E}_{X\sim p}\big[f\big(X\big)\big] - \mathbb{E}_{Y\sim q}\big[f\big(Y\big)\big]\big|. \tag{2.4}$$

The f-discrepancy leads naturally to a function-specific notion of mixing time: :=[f-mixing time] For a given function f, the f-mixing time is

$$T_f(\delta) = \min\Big\{n \in \mathbb{N} \;\mid\; \max_{i\in[d]} d_f\big(\pi_n^{(i)},\ \pi\big) \le \delta\Big\}. \tag{2.5}$$

We sometimes use T_f without an argument either when the argument is obvious from context, or when we want to refer to the quantity generically rather than its evaluation at a specific δ. In the sequel, we also define function-specific notions of the spectral gap of a Markov chain, which can be used to bound the f-mixing time and to obtain function-specific concentration inequalities.

We also use some asymptotic notation, which we now clarify. If $g_1,\ g_2$ are two nonnegative functions of some variable x, we define the notations

$$\begin{aligned}
g_1 \approx g_2 &\iff \exists c,\ c' > 0, \quad cg_1 \le g_2 \le c'g_1,\\
g_1 \asymp g_2 &\iff g_1 \approx g_2,\\
g_1 \lesssim g_2 &\iff \exists c > 0, \quad g_1 \le cg_2,\\
g_1 \gg g_2 &\iff g_1 \lesssim g_2.
\end{aligned}$$

Related work

Mixing times are a classical topic of study in Markov chain theory, and there is a large collection of techniques for their analysis [see, e.g., 2, 24, 63, 68, 73, 86]. These tools and the results based on them, however, generally apply only to worst-case mixing times. Outside of specific examples [22, 25], mixing with respect to individual functions or limited classes of functions has received relatively little attention, and almost none at all in the statistics literature.

One important exception is the recent work by Hayashi and Watanabe [50] and Watanabe and Hayashi [95], who provide asymptotically sharp tail bounds on empirical averages of functions using methods from information geometry. These bounds are of the form

$$\mathbb{P}\bigg(\frac{1}{N}\sum_{n=1}^{N} f(X_n) \ge \mu + \epsilon\bigg) \le \exp\big(NC(\epsilon) + D(\epsilon)\big)$$

where $D(\epsilon)$ is a constant defined explicitly [95] that tends to 0 as $\epsilon \to 0$, and $C(\epsilon)$ is the large-deviation rate, which is to say

$$\begin{aligned}
C(\epsilon) &= (\mu+\epsilon)\phi'^{-1}(\mu+\epsilon) - \phi\big(\phi'^{-1}(\mu+\epsilon)\big)\\
&= \lim_{N\to\infty} \frac{1}{N}\log \mathbb{P}\bigg(\frac{1}{N}\sum_{n=1}^{N} f(X_n) \ge \mu + \epsilon\bigg),
\end{aligned}$$

where $\phi(t)$ denotes the largest eigenvalue of the matrix $P(t)$ defined by $(P(t))_{i,j} = P\cdot e^{tf(i)}$ for $1 \le i, j \le d$. (The largest eigenvalue is a nonnegative real number by the Perron-Frobenius

Theorem.) Watanabe and Hayashi [95] also provide lower bounds that are closely matching these upper bounds. Together these are used to derive classical results in probability theory (large deviations, moderate deviations, and CLT) for Markov chains.

While we do not claim that the inequalities in this chapter are sharper than these results, they are stated in terms of f-mixing times which are much more intuitive and easier to use in practice than the large deviation rates. We provide several results based on spectral methods and coupling arguments that allow us to bound the f-mixing times, and illustrate the quality of our predictions in simulations, a task that appears to be more intensive computationally and algorithmically for the information-geometry bounds.

Other existing bounds are generally uniform over functions, and the rates that are reported include a factor that encodes the global mixing properties of the chain and does not adapt to the function [20, 41, 60, 61, 62, 65, 75, 84]. (A degree of adaptation is possible in that the asymptotic variance of the function f can be accounted for in Bernstein-type bounds, but the key factor does not adapt—see for instance Lezaud [64], Paulin [75].) These factors, which do not appear in classic bounds for independent random variables,[1] are generally either some variant of the spectral gap γ of the transition matrix, or else a mixing time of the chain $T\big(\delta_0\big)$ for some absolute constant $\delta_0 > 0$. For example, the main theorem from [62] shows that for a function $f\colon [d] \to [0,\ 1]$ and a sample $X_0 \sim \pi$ from the stationary distribution, we have

$$\mathbb{P}\big(\big|\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) - \mu\big| \geq \epsilon\big) \leq 2\exp\Big\{-\frac{\gamma_0}{2\big(2-\gamma_0\big)}\cdot\epsilon^2 N\Big\}, \tag{2.6}$$

where the eigenvalues of P are given in decreasing order as $1 > \lambda_2(P) \geq \cdots \geq \lambda_d(P)$, and we denote the spectral gap of P by

$$\gamma_0 := \min\big\{1 - \lambda_2(P),\ 1\big\}.$$

The requirement that the chain start in equilibrium can be relaxed by adding a correction for the burn-in time [75]. Extensions of this and related bounds, including bounded-differences-type inequalities and generalizations to continuous Markov chains and non-Markov mixing processes have also appeared in the literature (e.g., [61, 84]).

The concentration result has an alternative formulation in terms of the mixing time instead of the spectral gap [20]. This version and its variants are weaker, since the mixing time can be lower bounded as

$$T\big(\delta\big) \geq \big(\frac{1}{\gamma_*} - 1\big)\log\big(\frac{1}{2\delta}\big) \geq \big(\frac{1}{\gamma_0} - 1\big)\log\big(\frac{1}{2\delta}\big), \tag{2.7}$$

where we denote the absolute spectral gap [63] by

$$\gamma_* := \min\big(1 - \lambda_2,\ 1 - \big|\lambda_d\big|\big)\ \leq\ \gamma_0.$$

[1] Technically, since independent random variables form a Markov chain with spectral gap 1, the spectral gap does appear, but it appears in a trivial way as a factor of unity.

In terms of the minimum probability $\pi_{\min} := \min_i \pi_i$, the corresponding upper bound is an extra factor of $\log\left(\frac{1}{\pi_{\min}}\right)$ larger, which potentially leads to a significant gap between $\frac{1}{\gamma_0}$ and $T(\delta_0)$, even for a moderate constant such as $\delta_0 = \frac{1}{8}$. Similar distinctions arise in our analysis, and we elaborate on them at the appropriate junctures.

We note that there remains a gap between theoretical work on the convergence of Markov chains, of the kind developed here, and practical applications of the theory to MCMC. Indeed, most current applications of MCMC do not use rigorous bounds of the Hoeffding type; rather, they build variance-based confidence intervals, either via CLT approximations or, more recently, via Chebyshev's inequality [45, 31]. Such bounds are simple to compute and have good *asymptotic* theoretical properties, but they are not valid in a non-asymptotic setting; indeed, they may be over-optimistic and anti-conservative in finite samples. In contrast, Hoeffding-type bounds, including our own, sit at the other end of the spectrum; they come with finite-sample validity built in, but may be difficult to compute due to the dependence on the mixing time, either function-specific or uniform. There is a recent line of work aimed at estimating mixing times from individual sample trajectories [52, 51] that has begun to bridge the gap between the strong theory underlying Hoeffding bounds and their target applications, but this research direction is still nascent. We note that one other promising direction is the connection to analysis of specific statistically-relevant Markov chains [e.g., 19, 83]), which has the potential of yielding numerical bounds on mixing times [59].

Organization of the chapter

In the remainder of the chapter, we develop a formal framework for bounding function-specific mixing times and we apply the framework to the analysis of MCMC algorithms. In Section 2.2, we state some concentration guarantees based on function-specific mixing times, as well as some spectrum-based bounds on f-mixing times, and the spectrum-based Hoeffding bounds they imply. Section 2.4 is devoted to further development of these results in the context of several statistical models. More specifically, in Section 2.4, we show how our concentration guarantees can be used to derive confidence intervals that are superior to those based on uniform Hoeffding bounds and CLT-type bounds. In Section 2.4, we extend these ideas by demonstrating how our concentration results can improve sequential testing procedures in the context of MCMC. In Section 2.5, we show that our mixing time and concentration bounds improve over the non-adaptive bounds in real examples of MCMC from the literature. Finally, the bulk of our proofs are given in Section 2.3, with some more technical aspects of the arguments deferred to Sections 2.6, 2.6, and 2.6.

2.2 Main results

We now present our main technical contributions, starting with a set of "master" Hoeffding bounds with exponents given in terms of f-mixing times. As we explain in Section 2.2, these

mixing time bounds can be converted to spectral quantities that bound the f-mixing time in terms of the spectrum. (We give some techniques for the latter in Section 2.2).

Recall that we use $\mu := \mathbb{E}_\pi[f]$ to denote the mean. Moreover, we follow standard conventions in setting

$$\lambda_* := \max\big\{\lambda_2(P),\ \big|\lambda_d(P)\big|\big\}, \quad \text{and} \quad \lambda_0 := \max\big\{\lambda_2(P),\ 0\big\}.$$

so that the absolute spectral gap and the (truncated) spectral gap introduced earlier are given by $\gamma_* := 1 - \lambda_*$, and $\gamma_0 := 1 - \lambda_0$. In Section 2.2, we define and analyze corresponding function-specific quantities, which we introduce as necessary.

Master Hoeffding bound

In this section, we present a master Hoeffding bound that provides concentration rates that depend on the mixing properties of the chain only through the f-mixing time T_f. The only hypotheses on burn-in time needed for the bounds to hold are that the chain has been run for at least $N \geq T_f\,(\epsilon/2)$ steps—basically, so that thinning is possible—and that the chain was started from a distribution π_0 whose f-discrepancy distance from π is small—so that the expectation of each $f\big(X_n\big)$ iterate is close to μ—even if its total-variation discrepancy from π is large. Note that the latter requirement imposes only a very mild restriction, since it can always be satisfied by first running the chain for a burn-in period of T_f steps and then beginning to record samples. In fact, as we discuss below, it is not really necessary to explicitly discard the first T_f samples, so knowing the (function-specific) mixing time is not actually necessary, as long as N is larger than T_f. The tacit assumption in this theorem and all our concentration results is that f is bounded in $[0,\ 1]$.

Theorem 1. *Given any fixed $\epsilon > 0$ such that $d_f\big(\pi_0,\ \pi\big) \leq \frac{\epsilon}{2}$ and $N \geq T_f\big(\frac{\epsilon}{2}\big)$, we have*

$$\begin{aligned}\mathbb{P}\Big[\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + \epsilon\Big] &\leq \exp\left\{-\frac{\epsilon^2}{8} \cdot \left\lfloor \frac{N}{T_f\big(\frac{\epsilon}{2}\big)} \right\rfloor \right\}. \\ &\leq \exp\left\{-\frac{\epsilon^2 N}{16 T_f\big(\frac{\epsilon}{2}\big)}\right\}\end{aligned} \tag{2.8}$$

Compared to the bounds in earlier work [e.g., 62], the bound (2.8) has several distinguishing features. The primary difference is that the "effective" sample size, that is, the number of samples that would give an equivalent level of concentration if all the samples were i.i.d. from π,

$$N_{\text{eff}} := \big\lfloor \frac{N}{T_f(\epsilon/2)} \big\rfloor, \tag{2.9a}$$

is a function of f, which can lead to significantly sharper bounds on the deviations of empirical means than the earlier uniform bounds can deliver. Further, the result applies when the chain has equilibrated only approximately, and only with respect to f.

The reader might note that if one actually has access to a distribution π_0 that is $\epsilon/2$-close to π in f-discrepancy, then an estimator of μ with tail bounds similar to those guaranteed by Theorem 1 can be obtained as follows: first, draw N i.i.d. samples from π_0, and second, apply the usual Hoeffding inequality for i.i.d. variables. However, it is essential to realize that Theorem 1 does not require that such a π_0 be available to the practitioner. Instead, the theorem statement is meant to apply in the following way: suppose that—starting from *any* initial distribution—we run an algorithm for $N \geq T_f(\epsilon/2)$ steps, and then use the last of $N - T_f(\epsilon/2)$ samples to form an empirical average. Our concentration result then holds with an effective sample size of

$$N_{\text{eff}}^{\text{burnin}} := \big\lfloor \frac{N - T_f(\epsilon/2)}{T_f\,(\epsilon/2)} \big\rfloor \;=\; \big\lfloor \frac{N}{T_f\,(\epsilon/2)} \big\rfloor - 1. \tag{2.9b}$$

In other words, the result can be applied with an arbitrary initial π_0, and accounting for burn-in merely reduces the effective sample size by one. One can take this reasoning further to determine that, as mentioned above, it is not even necessary to explicitly use a burn-in period. Indeed, in order for Theorem 1 to be meaningful, it must be that $\frac{N}{T_f(\epsilon/2)} \gtrsim \frac{1}{\epsilon^2}$, so that, up to constants, the burn-in period is at most an ϵ^2 fraction of the total number of samples. It follows that directly averaging the function values along the trajectory provides a good approximation of the average over the last $N - T_f(\epsilon/2)$ samples, up to an accuracy on the order of $\epsilon^2 \ll \epsilon$.

The appearance of the function-specific mixing time T_f in the bounds comes with both advantages and disadvantages. A notable disadvantage, shared with the mixing time versions of the uniform bounds, is that spectrum-based bounds on the mixing time (including our f-specific ones) introduce a $\log\big(\frac{1}{\pi_{\min}}\big)$ term that can be a significant source of looseness. On the other hand, obtaining rates in terms of mixing times comes with the advantage that any bound on the mixing time translates directly into a version of the concentration bound (with the mixing time replaced by its upper bound). Moreover, since the $\pi_{\min}^{-1}$ term is likely to be an artifact of the spectrum-based approach, and possibly even just of the proof method, it may be possible to turn the T_f-based bound into a stronger spectrum-based bound with a more sophisticated analysis. We go part of the way toward doing this, albeit without completely removing the $\pi_{\min}^{-1}$ term.

An analysis based on mixing time also has the virtue of better capturing the non-asymptotic behavior of the rate. Indeed, as a consequence of the link (2.7) between mixing and spectral graphs (as well as matching upper bounds [63]), for any fixed function f, there exists a function-specific spectral gap $\gamma_f > 0$ such that

$$T_f\big(\frac{\epsilon}{2}\big) \approx \frac{1}{\gamma_f} \log\Big(\frac{1}{\epsilon}\Big) + O\big(1\big), \quad \text{for} \quad \epsilon \ll 1. \tag{2.9c}$$

These asymptotics can be used to turn our aforementioned theorem into a variant of the results of Léon and Perron [62], in which γ_0 is replaced by a value γ_f that (under mild conditions) is at least as large as γ_0. However, as we explore in Section 2.5, such an asymptotic

spectrum-based view loses a great deal of information needed to deal with practical cases, where often $\gamma_f = \gamma_0$ and yet $T_f(\delta) \ll T(\delta)$ even for very small values of $\delta > 0$. For this reason, part of our work is devoted to deriving more fine-grained concentration inequalities that capture this non-asymptotic behavior.

On the other hand, it is important to note that our bounds do not provide optimal rates in the asymptotic setting. Indeed, our results only imply convergence of the sample mean to the true mean at a rate of $\frac{\log N}{\sqrt{N}}$, due to the logarithmic dependence of T_f on the error ϵ. Our lower bound, Proposition 1, shows that the T_f factor cannot be removed in general, so that optimal asymptotic convergence rates do not seem attainable in general with a function-specific analysis. Nonetheless, by interpolating between the function-specific and global bounds, one can obtain the best of both worlds, so we do not believe the asymptotic sub-optimality to be a major concern. This point of view is supported by our experiments, which suggest that at the precisions one typically targets in practice, the seemingly extraneous $\log N$ factor is overcome by the gains of having a larger function-specific spectral gap, and the function-specific bounds end up being superior.

By combining our definition (2.9a) of the effective sample size N_{eff} with the asymptotic expansion (2.9c), we arrive at an intuitive interpretation of Theorem 1: it dictates that the effective sample size scales as $N_{\text{eff}} \approx \frac{\gamma_f N}{\log(1/\epsilon)}$ in terms of the function-specific gap γ_f and tolerance ϵ. This interpretation is backed by the Hoeffding bound derived in Corollary 1 and it is useful as a simple mental model of these bounds. On the other hand, interpreting the theorem this way effectively plugs in the asymptotic behavior of T_f and does not account for the non-asymptotic properties of the mixing time; the latter may actually be more favorable and lead to substantially smaller effective sample sizes than the naive asymptotic interpretation predicts. From this perspective, the master bound has the advantage that any bound on T_f that takes advantage of favorable non-asymptotics translates directly into a stronger version of the Hoeffding bound. We investigate these issues empirically in Section 2.5.

Based on the worst-case Markov Hoeffding bound (2.6), we might hope that the $T_f(\frac{\epsilon}{2})$ term in Theorem 1 is spurious and removable using improved techniques. Unfortunately, it is fundamental. This conclusion becomes less surprising if one notes that even if we start the chain in its stationary distribution and run it for $N < T_f(\epsilon)$ steps, it may still be the case that there is a large set Ω_0 such that for $i \in \Omega_0$ and $1 \leq n \leq N$,

$$|f(X_n) - \mu| \gg \epsilon \quad \text{a.s. if } X_0 = i. \tag{2.10}$$

This behavior is made possible by the fact that large positive and negative deviations associated with different values in Ω_0 can cancel out to ensure that $\mathbb{E}\left[f\left(X_n\right)\right] = \mu$ marginally. However, the lower bound (2.10) guarantees that

$$\begin{aligned}\mathbb{P}\left(\frac{1}{N}\sum_{n=1}^{N} f\left(X_n\right) \geq \mu + \epsilon\right) &\geq \sum_{i \in \Omega_0} \pi_i \cdot \mathbb{P}\left(\frac{1}{N}\sum_{n=1}^{N} f\left(X_n\right) \geq \mu + \epsilon \mid X_0 = i\right) \\ &\geq \pi\left(\Omega_0\right),\end{aligned}$$

so that if $\pi(\Omega_0) \gg 0$, we have no hope of controlling the large-deviation probability unless $N \gtrsim T_f(\epsilon)$.

To make this intuitive idea precise, the basic idea is to start with an arbitrary candidate function $\rho : (0,\ 1) \to (0,\ 1)$, such that $T_f\left(\frac{\epsilon}{2}\right)$ in the denominator of the function-specific Hoeffding bound (2.8) can putatively be replaced by $T_f(\rho(\epsilon))$. We then show that if $\rho(\epsilon) \geq \epsilon$, the replacement is not actually possible. That means that, up to a possible constant factor improvement in the argument to T_f, the dependence of the exponent in Theorem 1 on $T_f(\epsilon/2)$ cannot be eliminated—surprising in light of the fact that *uniform* Hoeffding bounds do not exhibit this behavior in their dependence on the mixing or relaxation times. In this sense, the rate we attain is improvable only in the constants in the exponent of the bound, as claimed above.

We prove Proposition 1 by constructing a Markov chain (which is independent of ϵ) and a function (which depends on both ϵ and ρ) such that the Hoeffding bound is violated for the Markov chain-function pair for some value of N (which in general depends on the chain and ϵ). We defer the proof to Section 2.3.

Proposition 1. *Fix a function $\rho\colon (0,\ 1) \to (0,\ 1)$ with $\rho\big(\epsilon\big) > \epsilon$. For every constant $c_1 > 0$ and $\epsilon \in (0,1)$, there exists a Markov chain P_{c_1}, a number of steps $N = N(c_1, \epsilon)$ and a function $f = f_\epsilon$ such that*

$$\mathbb{P}_\pi\left(\left|\frac{1}{N}\sum_{n=1}^{N} f(X_n) - \frac{1}{2}\right| \geq \epsilon\right) > 2 \cdot \exp\left(-\frac{c_1 N \epsilon^2}{T_f\left(\rho(\epsilon)\right)}\right). \tag{2.11}$$

Bounds on f-mixing times

We generally do not have direct access either to the mixing time $T\big(\delta\big)$ or the f-mixing time $T_f\big(\delta\big)$. Fortunately, any bound on T_f translates directly into a variant of the tail bound (2.8). Accordingly, this section is devoted to methods for bounding these quantities. Since mixing time bounds are equivalent to bounds on d_{TV} and d_f, we frame the results in terms of distances rather than times. These results can then be inverted in order to obtain mixing-time bounds in applications.

The simplest bound is simply a uniform bound on total variation distance, which also yields a bound on the f-discrepancy. In particular, if the chain is started with distribution π_0, then we have

$$d_{\mathrm{TV}}\big(\pi_n,\ \pi\big) \leq \frac{1}{\sqrt{\pi_{\min}}} \cdot \lambda_*^n \cdot d_{\mathrm{TV}}\big(\pi_0,\ \pi\big). \tag{2.12}$$

In order to improve upon this bound, we need to develop function-specific notions of spectrum and spectral gaps. The simplest way to do this is simply to consider the (left) eigenvectors to which the function is not orthogonal and define a spectral gap restricted only to the corresponding eigenvectors.

:=[f-eigenvalues and spectral gaps] For a function $f\colon [d] \to \mathbb{R}$, we define

$$J_f := \Big\{ j \in [d] \;\mid\; \lambda_j \neq 1 \text{ and } q_j^T f \neq 0 \Big\}, \tag{2.13a}$$

where q_j denotes a left eigenvector associated with λ_j. Similarly, we define

$$\lambda_f = \max_{j \in J_f} \big|\lambda_j\big|, \quad \text{and} \quad \gamma_f = 1 - \lambda_f. \tag{2.13b}$$

Using this notation, it is straightforward to show that if the chain is started with the distribution π_0, then

$$d_f\big(\pi_n,\ \pi\big) \le \sqrt{\frac{\mathbb{E}_\pi\big[f^2\big]}{\pi_{\min}}} \cdot \lambda_f^n \cdot d_f\big(\pi_0,\ \pi\big). \tag{2.14}$$

This bound, though useful in many cases, is also rather brittle: it requires f to be exactly orthogonal to the eigenfunctions of the transition matrix. For example, a function f_0 with a good value of λ_f can be perturbed by an arbitrarily small amount in a way that makes the resulting perturbed function f_1 have $\lambda_f = \lambda_*$. More broadly, the bound is of little value for functions with a small but nonzero inner product with the eigenfunctions corresponding to large eigenvalues (which is likely to occur in practice; cf. Section 2.5), or in scenarios where f lacks symmetry (cf. the random function example in Section 2.2).

In order to address these issues, we now derive a more fine-grained bound on d_f. The basic idea is to split the lower f-spectrum J_f into a "bad" piece J, whose eigenvalues are close to 1 but whose eigenvectors are approximately orthogonal to f, and a "good" piece $J_f \setminus J$, whose eigenvalues are far from 1 and which therefore do not require control on the inner products of their eigenvectors with f. More precisely, for a given set $J \subset J_f$, let us define

$$\Delta_J^* := 2\big|J\big| \times \max_{j \in J} \|h_j\|_\infty \times \max_{j \in J} \big|q_j^T f\big|, \qquad \lambda_J := \max\Big\{ \big|\lambda_j\big| \;\mid\; j \in J \Big\}, \quad \text{and}$$
$$\lambda_{-J} := \max\Big\{ \big|\lambda_j\big| \;\mid\; j \in J_f \setminus J \Big\}.$$

Here the h_j are the functions defined in the decomposition of P in (2.3). We obtain the following bound, expressed in terms of λ_{-J} and λ_J, which we generally expect to obey the relation $1 - \lambda_{-J} \ll 1 - \lambda_J$.

Lemma 1 (Sharper f-discrepancy bound)**.** *Given $f\colon [d] \to [0,\ 1]$ and a subset $J \subset J_f$, we have*

$$d_f\big(\pi_n,\ \pi\big) \le \Delta_J^*\, \lambda_J^n \cdot d_{\mathrm{TV}}(\pi_0,\ \pi) + \sqrt{\frac{\mathbb{E}_\pi\big[f^2\big]}{\pi_{\min}}} \cdot \lambda_{-J}^n\, d_f(\pi_0,\ \pi). \tag{2.15}$$

The above bound, while easy to apply and comparatively easy to estimate, can be loose when the first term is a poor estimate of the part of the discrepancy that comes from the J part of the spectrum. We can get a still sharper estimate by instead making use of the following vector quantity that more precisely summarizes the interactions between f and J:

$$h_J(n) := \sum_{j \in J} \left(q_j^T f \cdot \lambda_j^n\right) h_j.$$

This quantity leads to what we refer to as an *oracle-adaptive bound*, because it uses the exact value of the part of the discrepancy coming from the J eigenspaces, while using the same bound as above for the part of the discrepancy coming from $J_f \setminus J$.

Lemma 2 (Oracle f-discrepancy bound)**.** *Given $f\colon [d] \to [0,\ 1]$ and a subset $J \subset J_f$, we have*

$$d_f(\pi_n,\ \pi) \leq \left|(\pi_0 - \pi)^T h_J(n)\right| + \sqrt{\frac{\mathbb{E}_\pi\left[f^2\right]}{\pi_{\min}}} \cdot \lambda_{-J}^n \cdot d_f(\pi_0,\ \pi). \tag{2.16}$$

We emphasize that, although Lemma 2 is stated in terms of the initial distribution π_0, when we apply the bound in the real examples we consider, we replace all quantities that depend on π_0 by their worst-cases values, in order to avoid dependence on initialization; this results in a $\|h_J(n)\|_\infty$ term instead of the dot product in the lemma.

Concentration bounds

The mixing time bounds from Section 2.2 allow us to translate the master Hoeffding bound into a weaker but more interpretable—and in some instances, more directly applicable—concentration bound. The first result we prove along these lines applies meaningfully only to functions f whose absolute f-spectral gap γ_f is larger than the absolute spectral gap γ_*. It is a direct consequence of the master Hoeffding bound and the simple spectral mixing bound (2.14), and it delivers the asymptotics in N and ϵ promised in Section 2.2.

Corollary 1. *Given any $\epsilon > 0$ such that $d_f(\pi_0,\ \pi) \leq \frac{\epsilon}{2}$ and $N \geq T_f\left(\frac{\epsilon}{2}\right)$, we have*

$$\mathbb{P}\left[\frac{1}{N}\sum_{n=1}^N f(X_n) \geq \mu + \epsilon\right] \leq \begin{cases} \exp\left(-\frac{\epsilon^2}{16}\,\frac{\gamma_f N}{\log\left(\frac{2}{\epsilon\sqrt{\pi_{\min}}}\right)}\right) & \text{if } \epsilon \leq \frac{2\lambda_f}{\sqrt{\pi_{\min}}}, \\ \exp\left(-\frac{\epsilon^2 N}{16}\right) & \text{otherwise.} \end{cases}$$

Deriving a Hoeffding bound using the sharper f-mixing bound given in Lemma 1 requires more care, both because of the added complexity of managing two terms in the bound and because one of those terms does not decay, meaning that the bound only holds for sufficiently large deviations $\epsilon > 0$.

The following result represents one way of articulating the bound implied by Lemma 1; it leads to improvements over the previous two results when the contribution from the bad part of the spectrum J—that is, the part of the spectrum that brings γ_f closer to 1 than we would like—is negligible at the scale of interest. Recall that Lemma 1 expresses the contribution of J via the quantity Δ_J^*.

Corollary 2. *Given a triple of positive numbers $(\Delta, \Delta_J, \Delta_J^*)$ such that $\Delta_J \geq \Delta_J^*$, $d_f(\pi_0,\ \pi) \leq \Delta_J + \Delta$, and $N \geq T_f\big(\Delta_J + \Delta\big)$, we have*

$$\mathbb{P}\left[\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + 2\left(\Delta_J + \Delta\right)\right] \leq \begin{cases} \exp\Big(-\frac{\big(\Delta_J+\Delta\big)^2}{4}\,\frac{\big(1-\lambda_{-J}\big)N}{\log\big(\frac{1}{\Delta\sqrt{\pi_{\min}}}\big)}\Big) \ \textit{if } \Delta \leq \frac{\lambda_{-J}}{\sqrt{\pi_{\min}}}, \\ \exp\Big(-\frac{\big(\Delta_J+\Delta\big)^2 N}{4}\Big) \ \textit{if } \Delta > \frac{\lambda_{-J}}{\sqrt{\pi_{\min}}}. \end{cases} \tag{2.17}$$

Similar arguments can be applied to combine the master Hoeffding bounds with the oracle f-mixing bound Lemma 2, but we omit the corresponding result for the sake of brevity. The proofs for both aforementioned corollaries are in Section 2.3.

Example: Lazy random walk on C_{2d}

In order to illustrate the mixing time and Hoeffding bounds from Section 2.2, we analyze their predictions for various classes of functions on the $2d$-cycle C_{2d}, identified with the integers modulo $2d$. In particular, consider the Markov chain corresponding to a lazy random walk on C_{2d}; it has transition matrix

$$P_{uv} = \begin{cases} \frac{1}{2} & \text{if } v = u, \\ \frac{1}{4} & \text{if } v = u+1 \mod 2d, \\ \frac{1}{4} & \text{if } v = u-1 \mod 2d, \\ 0 & \text{otherwise.} \end{cases} \tag{2.18}$$

It is easy to see that the chain is irreducible, aperiodic, and reversible, and its stationary distribution is uniform. It can be shown [63] that its mixing time scales proportionally to d^2. However, as we now show, several interesting classes of functions mix much faster, and in fact, a "typical" function, meaning a randomly chosen one, mixes much faster than the naive mixing bound would predict.

Parity function. The epitome of a rapidly mixing function is the parity function:

$$f_{\text{parity}}(u)) := \begin{cases} 1 & \text{if } u \text{ is odd}, \\ 0 & \text{otherwise.} \end{cases} \tag{2.19}$$

It is easy to see that no matter what the choice of initial distribution π_0 is, we have $\mathbb{E}\big[f_{\text{parity}}(X_1)\big] = \frac{1}{2}$, and thus f_{parity} mixes in a single step.

Periodic functions. A more general class of examples arises from considering the eigenfunctions of P, which are given by $g_j(u) = \cos\left(\frac{\pi j u}{d}\right)$; [see, e.g., 63]. We define a class of functions of varying regularity by setting

$$f_j = \frac{1 + g_j}{2}, \quad \text{for each } j = 0, 1, \ldots, d.$$

Here we have limited j to $0 \leq j \leq d$ because f_j and f_{2d-j} behave analogously. Note that the parity function f_{parity} corresponds to f_d.

Intuitively, one might expect that some of these functions mix well before d^2 steps have elapsed—both because the vectors $\{f_j,\ j \neq 1\}$ are orthogonal to the non-top eigenvectors with eigenvalues close to 1 and because as j gets larger, the periods of f_j become smaller and smaller, meaning that their global behavior can increasingly be well determined by looking at local snapshots, which can be seen in few steps.

Our mixing bounds allow us to make this intuition precise, and our Hoeffding bounds allow us to prove correspondingly improved concentration bounds for the estimation of $\mu = \mathbb{E}_\pi[f_j] = 1/2$. Indeed, we have

$$\gamma_{f_j} = \frac{1 - \cos\left(\frac{\pi j}{d}\right)}{2} \geq \begin{cases} \frac{\pi^2 j^2}{24 d^2} & \text{if } j \leq \frac{d}{2}, \\ \frac{1}{2} & \text{if } \frac{d}{2} < j \leq d. \end{cases} \tag{2.20}$$

Consequently, equation (2.14) predicts that

$$T_{f_j}(\delta) \leq \tilde{T}_{f_j}(\delta) = \begin{cases} \frac{24}{\pi^2}\left[\frac{1}{2}\log 2d + \log\left(\frac{1}{\delta}\right)\right] \cdot \frac{d^2}{j^2} & \text{if } j \leq \frac{d}{2}, \\ \log 2d + 2\log\left(\frac{1}{\delta}\right) & \text{if } \frac{d}{2} < j \leq d, \end{cases} \tag{2.21}$$

where we have used the trivial bound $\mathbb{E}_\pi[f^2] \leq 1$ to simplify the inequalities. Note that this yields an improvement over $\asymp d^2$ for $j \gtrsim \log d$. Moreover, the bound (2.21) can itself be improved, since each f_j is orthogonal to all eigenfunctions other than $\mathbf{1}$ and g_j, so that the $\log d$ factors can all be removed by a more carefully argued form of Lemma 1. It thus follows directly from the bound (2.20) that if we draw $N + \tilde{T}_{f_j}\left(\frac{\epsilon}{2}\right)$ samples, we obtain the tail bound

$$\mathbb{P}\Big[\frac{1}{N_0}\sum_{n=N_{\text{b}}}^{N+N_{\text{b}}} f_j(X_n) \geq \frac{1}{2} + \epsilon\Big] \leq \begin{cases} \exp\Big(-\frac{3d^2}{2\pi^2 j^2} \cdot \frac{\epsilon^2 N}{\log\left(2\sqrt{2}d/\epsilon\right)}\Big) & \text{if } j \leq \frac{d}{2}, \\ \exp\Big(-\frac{\epsilon^2 N}{32\log\left(2\sqrt{2}d/\epsilon\right)}\Big) & \frac{d}{2} < j \leq d, \end{cases} \tag{2.22}$$

where the burn-in time is given by $N_{\text{b}} = \tilde{T}_{f_j}(\epsilon/2)$. Note again that the sharper analysis mentioned above would allow us to remove the $\log 2d$ factors.

Random functions. A more interesting example comes from considering a randomly chosen function $f\colon C_{2d} \to [0,\ 1]$. Indeed, suppose that the function values are sampled i.i.d. from some distribution ν on $[0,\ 1]$ whose mean μ^* is $1/2$:

$$\{f(u),\ u \in C_{2d}\} \overset{\text{i.i.d.}}{\sim} \nu. \tag{2.23}$$

We can then show that for any fixed $\delta^* > 0$, with high probability over the randomness of f, have

$$T_f(\delta) \lesssim \frac{d \log d \big[\log d + \log\big(\frac{1}{\delta}\big)\big]}{\delta^2}, \qquad \text{for all } \delta \in (0, \delta^*]. \tag{2.24}$$

For $\delta \gg \frac{\log d}{\sqrt{d}}$, this scaling is an improvement over the global mixing time of order $d^2 \log(1/\delta)$.

The core idea behind the proof of equation (2.24) is to apply Lemma 1 with

$$J_\delta := \left\{ j \in \mathbb{N} \cap [1, 2d-1] \;\mid\; j \leq 4\delta\sqrt{\frac{d}{\log d}} \;\text{ or }\; j \geq 2d - 4\delta\sqrt{\frac{d}{\log d}} \right\}. \tag{2.25}$$

It can be shown that $\|h_j\|_\infty = 1$ for all $0 \leq j < 2d$ and that with high probability over f, $|q_j^T f| \lesssim \sqrt{\frac{\log d}{d}}$ simultaneously for all $j \in J_\delta$, which suffices to reduce the first part of the sharper f-discrepancy bound to order δ.

In order to estimate the rate of concentration, we proceed as follows. Taking $\delta = c_0 \epsilon$ for a suitably chosen universal constant $c_0 > 0$, we show that $\Delta_J := \frac{\epsilon}{4} \geq \Delta_J^*$. We can then set $\Delta = \frac{\epsilon}{4}$ and observe that with high probability over f, the deviation in Corollary 2 satisfies the bound $2\,(\Delta_J + \Delta) \leq \epsilon$. With δ as above, we have $1 - \lambda_{-J} \geq \frac{c_1 \epsilon^2}{d \log d}$ for another universal constant $c_1 > 0$. Thus, if we are given $N + T_f\big(\epsilon/2\big)$ samples for some $N \geq T_f\big(\frac{\epsilon}{2}\big)$, then we have

$$\mathbb{P}\left[\frac{1}{N} \sum_{n=T_f(\epsilon/2)}^{N+T_f(\epsilon/2)} f(X_n) \geq \mu + \epsilon \right] \leq \exp\left\{ -\frac{c_2 \epsilon^4 N}{d \log d \big[\log\big(\frac{4}{\epsilon}\big) + \log 2d \big]} \right\}, \tag{2.26}$$

for some $c_2 > 0$. Consequently, it suffices for the sample size to be lower bounded by

$$N \gtrsim \frac{d \log d \big[\log\big(1/\epsilon\big) + \log d \big]}{\epsilon^4},$$

in order to achieve an estimation accuracy of ϵ. Notice that this requirement is an improvement over the $\frac{d^2}{\epsilon^2}$ from the uniform Hoeffding bound provided that $\epsilon \gg (\frac{\log^2 d}{d})^{1/2}$. Proofs of all these claims can be found in Section 2.6.

2.3 Proofs of main results

This section is devoted to the proofs of the main results of this chapter.

Proof of Theorem 1

We begin with the proof of the master Hoeffding bound from Theorem 1. At the heart of the proof is the following bound on the moment-generating function (MGF) for the sum of

an appropriately thinned subsequence of the function values $\{f(X_n)\}_{n=1}^{\infty}$. In particular, let us introduce the shorthand notation $\widetilde{X}_{m,t} := X_{(m-1)T_f(\epsilon/2)+t}$ and $N_0 := \lfloor N/T_f(\frac{\epsilon}{2})\rfloor$. With this notation, we have the following auxiliary result:

Lemma 3 (Master MGF bound)**.** *For any scalars $\beta \in \mathbb{R}$, $\epsilon \in (0,1)$, integer $t \in \big[0, T_f(\frac{\epsilon}{2})\big)$, and integer $M \in \Big\{N_0,\ N_0+1\Big\}$, we have*

$$\mathbb{E}\left[\exp\Big(\beta \sum_{m=1}^{M} f\big(\widetilde{X}_{m,t}\big)\Big)\right] \leq \exp\Big\{\big[\frac{1}{2}\beta\epsilon + \beta\mu + \frac{1}{2}\beta^2\big]\cdot M\Big\}. \tag{2.27}$$

See Section 2.3 for the proof of this claim. First write $N = T_f(\epsilon/2)N_0 + r$. Recalling the definition of $\widetilde{X}_{m,t}$, we have

$$\begin{aligned}
e^{-\alpha\big(\mu+\epsilon\big)}\mathbb{E}\left[e^{\alpha\sum_{n=1}^{N} f(X_n)}\right] &= \mathbb{E}\left[\exp\left\{\alpha\sum_{t=1}^{r}\sum_{m=1}^{N_0+1}\big[f(\widetilde{X}_{m,t})-\mu-\epsilon\big] + \alpha\sum_{t=r+1}^{T_f(\epsilon/2)}\sum_{m=1}^{N_0}\big[f(\widetilde{X}_{m,t})-\mu-\epsilon\big]\right\}\right] \\
&= \mathbb{E}\left[\exp\left\{\alpha T_f(\epsilon/2)\cdot\frac{1}{T_f(\epsilon/2)}\cdot\left[\sum_{t=1}^{r}\sum_{m=1}^{N_0+1}\big[f(\widetilde{X}_{m,t})-\mu-\epsilon\big]\right.\right.\right. \\
&\qquad\qquad \left.\left.\left. + \sum_{t=r+1}^{T_f(\epsilon/2)}\sum_{m=1}^{N_0}\big[f(\widetilde{X}_{m,t})-\mu-\epsilon\big]\right]\right\}\right] \\
&\leq \frac{1}{T_f\big(\epsilon/2\big)}\left\{\sum_{t=1}^{r}\mathbb{E}\left[\exp\Big\{\alpha T_f\big(\epsilon/2\big)\sum_{m=1}^{N_0+1} f\big(\widetilde{X}_{m,t}\big)\Big\}\right]e^{-(\mu+\epsilon)(N_0+1)}\right. \\
&\qquad\qquad \left. + \sum_{t=r+1}^{T_f(\epsilon/2)}\mathbb{E}\left[\exp\Big\{\alpha T_f\big(\epsilon/2\big)\sum_{m=1}^{N_0} f\big(\widetilde{X}_{m,t}\big)\Big\}\right]e^{-(\mu+\epsilon)N_0}\right\},
\end{aligned}$$

where the last inequality follows from Jensen's inequality, as applied to the exponential function. On the other hand, applying Lemma 3 with $\beta = \alpha T_f\big(\epsilon/2\big)$, we have

$$\mathbb{E}\left[\exp\Big\{\alpha T_f\big(\epsilon/2\big)\sum_{m=1}^{M} f\big(\widetilde{X}_{m,t}\big)\Big\}\right] \leq \exp\Big\{\big[\frac{1}{2}\alpha T_f\big(\frac{\epsilon}{2}\big)\epsilon + \alpha T_f\big(\frac{\epsilon}{2}\big)\mu + \frac{1}{2}\alpha^2 T_f^2\big(\frac{\epsilon}{2}\big)\big]\cdot M\Big\}$$

for $M \in \{N_0,\ N_0+1\}$ and $\alpha > 0$. By exponentiating and applying Markov's inequality, it

follows that

$$\mathbb{P}\left[\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + \epsilon\right] \leq e^{-\alpha(\mu+\epsilon)}\mathbb{E}\big[e^{\alpha\sum_{n=1}^{N} f(X_n)}\big]$$
$$\leq \frac{r}{T_f(\epsilon/2)} \cdot \exp\left\{\frac{1}{2}\cdot\Big[-\alpha\, T_f(\frac{\epsilon}{2})\epsilon + \alpha^2\, T_f^2(\frac{\epsilon}{2})\Big]\,\big(N_0+1\big)\right\}$$
$$+\left(1-\frac{r}{T_f(\epsilon/2)}\right)\cdot\exp\left\{\frac{1}{2}\cdot\Big[-\alpha\, T_f(\frac{\epsilon}{2})\epsilon + \alpha^2\, T_f^2(\frac{\epsilon}{2})\Big]\, N_0\right\}$$

The proof of Theorem 1 follows by taking $\alpha = \frac{\epsilon}{2T_f\left(\frac{\epsilon}{2}\right)}$ since

$$\mathbb{P}\left[\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + \epsilon\right] \leq \frac{r}{T_f(\epsilon/2)} \cdot \exp\left\{\frac{1}{2}\cdot\Big[-\frac{\epsilon^2}{2}+\frac{\epsilon^2}{4}\Big]\cdot(N_0+1)\right\}$$
$$+\left(1-\frac{r}{T_f(\epsilon/2)}\right)\cdot\exp\left\{\frac{1}{2}\cdot\Big[-\frac{\epsilon^2}{2}+\frac{\epsilon^2}{4}\Big]\cdot N_0\right\}$$
$$\leq \exp\left\{-\frac{\epsilon^2 N_0}{8}\right\}$$
$$= \exp\left\{-\frac{\epsilon^2}{8}\cdot\left\lfloor\frac{N}{T_f\left(\frac{\epsilon}{2}\right)}\right\rfloor\right\}.$$

Since $\lfloor x\rfloor \geq x/2$ for $x \geq 1$ and $\frac{N}{T_f\left(\frac{\epsilon}{2}\right)} \geq 1$ by assumption, the second bound follows immediately.

Proof of Lemma 3

For the purposes of the proof, fix t and let $W_m = \tilde{X}_{m,t}$. For convenience, also define a dummy constant random variable $W_0 := 0$. Now, by assumption, we have

$$\left|\mathbb{E}\left[f\left(W_1\right)\right] - \mu\right| \leq \frac{\epsilon}{2} \quad \text{and} \quad \left|\mathbb{E}\left[f\left(W_{m+1}\right) \mid W_m\right] - \mu\right| \leq \frac{\epsilon}{2}.$$

We therefore have the bound

$$\mathbb{E}\left[e^{\alpha\sum_m f(W_m)}\right] \leq \mathbb{E}\left[\prod_{m=1}^{N_0} e^{\alpha[f(W_m)-\mathbb{E}[f(W_m)\,\mid\,W_{m-1}]]}\right]\cdot e^{\alpha\mu M+\frac{\alpha\epsilon M}{2}}. \tag{2.28}$$

But now observe that the random variables $\Delta_m = f\left(W_m\right) - \mathbb{E}\left[f\left(W_m\right) \mid W_{m-1}\right]$ are deterministically bounded in $[-1,\ 1]$ and zero mean conditional on W_{m-1}. Moreover, by the Markovian property, this implies that the same is true conditional on $W_{<m} := W_{0:(m-1)}$. It follows by standard MGF bounds that

$$\mathbb{E}\left[e^{\alpha\Delta_m} \mid W_{<m}\right] \leq e^{\frac{\alpha^2}{2}}.$$

Combining this bound with inequality (2.28), we conclude that

$$\mathbb{E}\left[e^{\alpha \sum_m f(W_m)}\right] \leq e^{\frac{\alpha^2}{2} \cdot M} \cdot e^{\alpha \mu M + \frac{\alpha \epsilon M}{2}},$$

as claimed.

Proofs of Corollaries 1 and 2

In this section, we prove the derived Hoeffing bounds stated in Corollaries 1 and 2.

Proof of Corollary 1

The proof is a direct application of Theorem 1. Indeed, it suffices to note that if $\epsilon \leq \frac{2\lambda_f}{\sqrt{\pi_{\min}}}$, then

$$T_f\big(\frac{\epsilon}{2}\big) \leq \frac{\log\big(\frac{2}{\epsilon\sqrt{\pi_{\min}}}\big)}{\log\big(\frac{1}{\lambda_f}\big)} = \frac{\log\big(\frac{2}{\epsilon}\big) + \frac{1}{2}\log\big(\frac{1}{\pi_{\min}}\big)}{\log\big(\frac{1}{\lambda_f}\big)},$$

which yields the first bound. Turning to the second bound, note that if $\epsilon > \frac{2\lambda_f}{\sqrt{\pi_{\min}}}$, then equation (2.14) implies that $T_f\big(\frac{\epsilon}{2}\big) = 1$, which establishes the claim.

Proof of Corollary 2

The proof involves combining Theorem 1 with Lemma 1, using the setting $\epsilon = 2\left(\Delta + \Delta_J\right)$. We begin by combining the bounds $\lambda_J \leq 1$, $d_{\mathrm{TV}}\left(\pi_0,\ \pi_n\right) \leq 1$, $d_f\left(\pi_0,\ \pi_n\right) \leq 1$, and $\mathbb{E}_\pi\left[f^2\right] \leq 1$ with the claim of Lemma 1 so as to find that

$$d_f\left(\pi_0,\ \pi_n\right) \leq \Delta_J^* + \frac{\lambda_{-J}^n}{\sqrt{\pi_{\min}}} \ \leq\ \Delta_J + \frac{\lambda_{-J}^n}{\sqrt{\pi_{\min}}}.$$

It follows that

$$T_f(\frac{\epsilon}{2}) = T_f\left(\Delta_J + \Delta\right) \leq \frac{\log\big(\frac{1}{\Delta}\big) + \frac{1}{2}\log\Big(\frac{1}{\pi_{\min}}\Big)}{\log(\frac{1}{\lambda_{-J}})} \qquad \text{whenever } \Delta \leq \frac{\lambda_{-J}}{\sqrt{\pi_{\min}}}.$$

Plugging into Theorem 1 now yields the first part of the bound. On the other hand, if $\Delta > \frac{\lambda_{-J}}{\sqrt{\pi_{\min}}}$, then Lemma 1 implies that $T_f\left(\Delta_J + \Delta\right) = 1$, which proves the bound in the second case.

Proof of Proposition 1

In order to prove the lower bound in Proposition 1, we first require an auxiliary lemma:

Lemma 4. *Fix a function $\rho\colon\ (0,\ 1) \to (0,\ 1)$ with $\rho(\epsilon) > \epsilon$. For every constant $c_0 \geq 1$, there exists a Markov chain P_{c_0} and a function f_ϵ on it such that $\mu = \frac{1}{2}$, yet, for $N = c_0 T_f(\rho(\epsilon))$, and starting the chain from stationarity,*

$$\mathbb{P}_\pi\left(\left|\frac{1}{N}\sum_{n=1}^{N} f_\epsilon(X_n) - \frac{1}{2}\right| \geq \epsilon\right) \geq \frac{1}{3}.$$

Using this lemma, let us now prove Proposition 1. Suppose that we make the choices

$$c_0 := \left\lceil \frac{\log 7}{c_1 \epsilon^2} \right\rceil \geq 1, \quad \text{and} \quad N_{c_1,\epsilon} := c_0 T_f(\epsilon),$$

in Lemma 4. Letting P_{c_0} be the corresponding Markov chain and f_ϵ the function guaranteed by the lemma, we then have

$$\mathbb{P}_\pi\left(\left|\frac{1}{N_{c_1,\epsilon}}\sum_{n=1}^{N_{c_1,\ \epsilon}} f_\epsilon(X_n) - \frac{1}{2}\right| \geq \epsilon\right) \geq \frac{1}{3} \ > \ \frac{2}{7} \ \geq \ 2 \cdot \exp\left(-\frac{c_1 N_{c_1,\epsilon}\epsilon^2}{T_f(\rho(\epsilon))}\right).$$

Proof of Lemma 4: It only remains to prove Lemma 4, which we do by constructing pathological function on a chain graph, and letting our Markov chain be the lazy random walk on this graph. For the proof, fix $\epsilon > 0$, let $\rho = \rho(\epsilon)$ and let $T_f = T_f(\rho)$. Now choose an integer $d > 0$ such that $d > 2c_0$ and let the state space be Ω be the line graph with $2d$ elements with the standard lazy random walk defining P. We then set

$$f(i) = \begin{cases} \frac{1}{2} - \rho & 1 \leq i \leq d, \\ \frac{1}{2} + \rho & d+1 \leq i \leq 2d. \end{cases}$$

It is then clear that $T_f = 1$.

Define the bad event

$$\mathcal{E} = \left\{ X_1 \in \left[0,\ \frac{d}{2}\right] \cup \left[\frac{3d}{2},\ 2d\right] \right\}.$$

When this occurs, we have

$$\left|\frac{1}{N'}\sum_{n=1}^{N'} f(X_n) - \frac{1}{2}\right| \geq \rho > \epsilon \qquad \text{with probability one,}$$

for all $N' < \frac{d}{2}$. Since $N = c_0 < \frac{d}{2}$, we can set $N' = N$.

On the other hand, under π, the probability of $\mathcal{E}$ is $\geq \frac{1}{3}$. (It is actually about $\frac{1}{2}$, but we want to ignore edge cases.) The claim follows immediately.

2.4 Statistical applications

We now consider how our results apply to Markov chain Monte Carlo (MCMC) in various statistical settings. Our investigation proceeds along three connected avenues. We begin by showing, in Section 2.4, how our concentration bounds can be used to provide confidence intervals for stationary expectations that avoid the over-optimism of pure CLT predictions without incurring the prohibitive penalty of the Berry-Esseen correction—or the global mixing rate penalty associated with spectral-gap-based confidence intervals. Then, in Section 2.4, we show how our results allow us to improve on recent sequential hypothesis testing methodologies for MCMC, again replacing the dependence on the spectral gap by a dependence on the f-mixing time. Finally, in Section 2.5, we illustrate the practical significance of function-specific mixing properties by using our framework to analyze three real-world instances of MCMC, basing both the models and datasets chosen on real examples from the literature.

Confidence intervals for posterior expectations

In many applications, a point estimate of $\mathbb{E}_\pi\big[f\big]$ does not suffice; the uncertainty in the estimate must be quantified, for instance by providing $(1-\alpha)$ confidence intervals for some pre-specified constant α. In this section, we discuss how improved concentration bounds can be used to obtain sharper confidence intervals. In all cases, we assume the Markov chain is started from some distribution π_0 that need not be the stationary distribution, meaning that the confidence intervals must account for the burn-in time required to get close to equilibrium.

We first consider a bound that is an immediate consequence of the uniform Hoeffding bound given by [62]. As one would expect, it gives contraction at the usual Hoeffding rate but with an effective sample size of $N_{\text{eff}} \approx \gamma_0(N - T_0)$, where T_0 is the tuneable burn-in parameter. Note that this means that no matter how small T_f is compared to the global mixing time T, the effective size incurs the penalty for a global burn-in and the effective sample size is determined by the global spectral parameter γ_0. In order to make this precise, for a fixed burn-in level $\alpha_0 \in (0, \alpha)$, define

$$\epsilon_N(\alpha,\ \alpha_0) := \sqrt{2\big(2-\gamma_0\big)} \cdot \sqrt{\frac{\log\big(2/\big[\alpha-\alpha_0\big]\big)}{\gamma_0\big[N - T\big(\alpha_0\big)\big]}}. \tag{2.29a}$$

Then the uniform Markov Hoeffding bound [62, Thm. 1] implies that the set

$$I_N^{\text{unif}}\big(\alpha,\ \alpha_0\big) = \left[\frac{1}{N - T\big(\alpha_0/2\big)} \sum_{n=T\left(\alpha_0/2\right)+1}^{N} f\big(X_n\big) \pm \epsilon_N\big(\alpha,\ \alpha_0\big)\right] \tag{2.29b}$$

is a $1-\alpha$ confidence interval. Full details of the proof are given in Section 2.6.

Moreover, given that we have a family of confidence intervals—one for each choice of $\alpha_0 \in (0, \alpha)$—we can obtain the sharpest confidence interval by computing the infimum $\epsilon_N^*(\alpha) := \inf_{0<\alpha_0<\alpha} \epsilon_N(\alpha,\ \alpha_0)$. Equation (2.29b) then implies that

$$I_N^{\text{unif}}(\alpha) = \Big[\frac{1}{N - T(\alpha_0)} \sum_{n=T(\alpha_0/2)+1}^{N} f(X_n) \pm \epsilon_N^*(\alpha)\Big]$$

is a $1-\alpha$ confidence interval for μ.

We now consider one particular application of our Hoeffding bounds to confidence intervals, and find that the resulting interval adapts to the function, both in terms of burn-in time required, which now falls from a global mixing time to an f-specific mixing time, and in terms of rate, which falls from $\frac{1}{\gamma_0}$ to $T_f(\delta)$ for an appropriately chosen $\delta > 0$. We first note that the one-sided tail bound of Theorem 1 can be written as $e^{-r_N(\epsilon)/16}$, where

$$r_N(\epsilon) := \epsilon^2 \left[\frac{N}{T_f\left(\frac{\epsilon}{2}\right)} - 1\right]. \tag{2.30}$$

If we wish for each tail to have probability mass that is at most $\alpha/2$, we need to choose $\epsilon > 0$ so that $r_N(\epsilon) \geq 16 \log \frac{2}{\alpha}$, and conversely any such ϵ corresponds to a valid two-sided $(1-\alpha)$ confidence interval. Let us summarize our conclusions:

Theorem 2. *For any width $\epsilon_N \in r_N^{-1}\big([16 \log(2/\alpha),\ \infty)\big)$, the set*

$$I_N^{\text{func}} := \left[\frac{1}{N - T_f\left(\frac{\epsilon}{2}\right)} \sum_{n=T_f\left(\frac{\epsilon}{2}\right)}^{N} f(X_n) \pm \epsilon_N\right]$$

is a $1-\alpha$ confidence interval for the mean $\mu = \mathbb{E}_\pi[f]$.

In order to make the result more amenable to interpretation, first note that for any $0 < \eta < 1$, we have

$$r_N(\epsilon) \geq \underbrace{\epsilon^2 \left[\frac{N}{T_f\left(\frac{\eta}{2}\right)} - 1\right]}_{r_{N,\eta}(\epsilon)} \quad \text{valid for all } \epsilon \geq \eta. \tag{2.31}$$

Consequently, whenever $r_{N,\eta}(\epsilon_N) \geq 16 \log \frac{2}{\alpha}$ and $\epsilon_N \geq \eta$, we are guaranteed that a symmetric interval of half-width ϵ_N is a valid $(1-\alpha)$-confidence interval. Summarizing more precisely, we have:

Corollary 3. *Fix $\eta > 0$ and let*

$$\epsilon_N = r_{N,\eta}^{-1}\big(16 \log \frac{2}{\alpha}\big) = 4\sqrt{\frac{T_f\big(\frac{\eta}{2}\big) \cdot \log\big(2/\alpha\big)}{N - T_f\big(\frac{\eta}{2}\big)}}.$$

If $N \geq T_f\big(\frac{\eta}{2}\big)$, then I_N^{func} is a $1-\alpha$ confidence interval for $\mu = \mathbb{E}_\pi\big[f\big]$.

Often, we do not have direct access to $T_f\big(\delta\big)$, but we can often obtain an upper bound $\tilde{T}_f\big(\delta\big)$ that is valid for all $\delta > 0$. In Section 2.6, therefore, which contains the proofs for this section, we prove a strengthened form of Theorem 2 and its corollary in that setting.

A popular alternative strategy for building confidence intervals using MCMC depends on the Markov central limit theorem [e.g., 31, 59, 42, 82]. If the Markov CLT held exactly, it would lead to appealingly simple confidence intervals of width

$$\tilde{\epsilon}_N = \sigma_{f,\text{asym}} \sqrt{\frac{\log(2/\alpha)}{N}},$$

where $\sigma_{f,\text{asym}}^2 := \lim_{N\to\infty} \frac{1}{N} \text{Var}_{X_0 \sim \pi}\big[\sum_{n=1}^N f\big(X_n\big)\big]$ is the asymptotic variance of f.

Unfortunately, the CLT does not hold exactly, even after the burn-in period. The amount by which it fails to hold can be quantified using a Berry-Esseen bound for Markov chains, as we now discuss. Let us adopt the compact notation $\tilde{S}_N = \sum_{n=1}^N \big[f\big(X_n\big) - \mu\big]$. We then have the bound [65]:

$$\big|\mathbb{P}\big(\frac{\tilde{S}_N}{\sigma_{f,\text{asym}}\sqrt{N}} \leq s\big) - \Phi\big(s\big)\big| \leq \frac{e^{-\gamma_0 N}}{3\sqrt{\pi_{\min}}} + \frac{13}{\sigma_{f,\text{asym}}\sqrt{\pi_{\min}}} \cdot \frac{1}{\gamma_0\sqrt{N}}, \tag{2.32}$$

where Φ is the standard normal CDF. Note that this bound accounts for both the non-stationarity error and for non-normality error at stationarity. The former decays rapidly at the rate $e^{-\gamma_0 N}$, while the latter decays far more slowly, at the rate $\frac{1}{\gamma_0\sqrt{N}}$.

While the bound (2.32) makes it possible to prove a corrected CLT confidence interval, the resulting bound has two significant drawbacks. The first is that it only holds for extremely large sample sizes, on the order of $\frac{1}{\pi_{\min}\gamma_0^2}$, compared to the order $\frac{\log\big(1/\pi_{\min}\big)}{\gamma_0}$ required by the uniform Hoeffding bound. The second, shared by the uniform Hoeffding bound, is that it is non-adaptive and therefore bottlenecked by the global mixing properties of the chain. For instance, if the sample size is bounded below as

$$N \geq \max\big(\frac{1}{\gamma_0}\log\big(\frac{2}{\sqrt{\pi_{\min}}\alpha}\big)\ ,\ \frac{1}{\gamma_0^2}\frac{6084}{\sigma_{f,\text{asym}}^2\pi_{\min}\alpha^2}\big),$$

then both terms of equation (2.29b) are bounded by 1/6, and the confidence intervals take the form

$$I_N^{\text{BE}} = \left[\frac{1}{N}\sum_{n=1}^N f\big(X_n\big) \pm \sigma_{f,\text{asym}}\sqrt{\frac{2\log\big(6/\alpha\big)}{N}}\right]. \tag{2.33}$$

See Section 2.6 for the justification of this claim.

It is important to note that the width of this confidence interval involves a hidden form of mixing penalty. Indeed, defining the variance $\sigma_f^2 = \mathrm{Var}_\pi\big[f\big(X\big)\big]$ and $\rho_f := \frac{\sigma_f^2}{\sigma_{f,\mathrm{asym}}^2}$, we can rewrite the width as

$$\epsilon_N = \sigma_f \sqrt{\frac{2\log\big(6/\alpha\big)}{\rho_f N}}.$$

Thus, for this bound, the quantity ρ_f captures the penalty due to non-independence, playing the role of γ_0 and γ_f in the other bounds. In this sense, the CLT bound adapts to the function f, but only when it applies, which is at a sample-size scale dictated by the global mixing properties of the chain (i.e., γ_0).

Sequential testing for MCMC

For some applications, full confidence intervals may be unnecessary; instead, a practitioner may merely want to know whether $\mu = \mathbb{E}_\pi[f]$ lies above or below some threshold $0 < r < 1$. In these cases, we would like to develop a procedure for distinguishing between the two possibilities, at a given tolerable level $0 < \alpha < 1$ of combined Type I and II error. The simplest approach is, of course, to choose N so large that the $1-\alpha$ confidence interval built from N MCMC samples lies entirely on one side of r, but it may be possible to do better by using a sequential test. This latter idea was recently investigated in Gyori and Paulin [46], and we consider the same problem settings that they did:

(a) Testing with (known) indifference region, involving a choice between

$$H_0 \colon \mu \geq r + \delta$$
$$H_1 \colon \mu \leq r - \delta;$$

(b) Testing with no indifference region—that is, the same as above but with $\delta = 0$.

For the first setting (a), we always assume $0 < \delta < \nu := \min(\mu, 1-\mu)$, and the algorithm is evaluated on its ability to correctly choose between H_0 and H_1 when one of them holds, but it incurs no penalty for either choice when μ falls in the indifference region $\big(r-\delta,\ r+\delta\big)$. The error of a procedure A can thus be defined as

$$\mathrm{err}\big(\mathrm{A},\ f\big) = \begin{cases} \mathbb{P}\big(\mathrm{A}\big(X_{1:\infty}\big) = H_1\big) & \text{if } \mu \in H_0, \\ \mathbb{P}\big(\mathrm{A}\big(X_{1:\infty}\big) = H_0\big) & \text{if } \mu \in H_1, \\ 0 & \text{otherwise.} \end{cases}$$

To fix ideas, we can formally define a procedure A for this problem as a family of mappings $(\mathrm{A}_n \colon \Omega^n \to \{H_0,\ H_1,\ \textsc{continue}\})_{n=1}^\infty$, with the finite stopping requirement that for

any sample trajectory $X_{1:\infty}$ of the Markov chain, there exist a finite n such that $\mathrm{A}_n(X_{1:n}) \neq$ CONTINUE and the self-consistency requirement that if $\mathrm{A}_n(x_{1:n}) \in \{H_0,\ H_1\}$, then $\mathrm{A}_m(x_{1:m}) = \mathrm{A}_n(x_{1:n})$ for all $m \geq n$. The output of the overall procedure can then formally be defined as

$$\mathrm{A}(X_{1:\infty}) = \mathrm{A}_{n^*}(X_{1:n^*}), \quad n^* = \min\{n > 0 \colon \mathrm{A}_n(X_{1:n}) \neq \textsc{continue}\}.$$

The rest of this subsection is organized as follows. For setting (a), we analyze a baseline procedure $\mathcal{A}_{\text{fixed}}$ that makes a decision after a fixed number $N := N(\alpha)$ of samples. Our main analysis for this problem, however, concerns a sequential procedure $\mathcal{A}_{\text{seq}}$ that chooses whether to reject at a sequence $N_0, \ldots, N_k, \ldots$ of decision times. For the more challenging setting (b), we analyze $\mathcal{A}_{\text{hard}}$, which also rejects at a sequence of decision times. For both $\mathcal{A}_{\text{seq}}$ and $\mathcal{A}_{\text{hard}}$, we calculate the expected stopping times of the procedures. For the sake of clarity, we defer proofs of all results to Section 2.6.

As mentioned above, the simplest procedure $\mathcal{A}_{\text{fixed}}$ would choose a fixed number N of samples to be collected based on the target level α. After collecting N samples, it forms the empirical average $\hat{\mu}_N = \frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big)$ and outputs H_0 if $\hat{\mu}_N \geq r + \delta$, H_1 if $\hat{\mu}_N \leq r - \delta$, and outputs a special indifference symbol, say I, otherwise.

The sequential algorithm $\mathcal{A}_{\text{seq}}$ makes decisions as to whether to output one of the hypotheses or continue testing at a fixed sequence of decision times, say N_k. These times are defined recursively by

$$N_0 = \big\lfloor M \cdot \min\big(\frac{1}{r},\ \frac{1}{1-r}\big)\big\rfloor, \tag{2.34}$$

$$N_k = \big\lfloor N_0\big(1+\xi\big)^k\big\rfloor, \tag{2.35}$$

where $M > 0$ and $0 < \xi < 2/5$ are parameters of the algorithm. At each time N_k for $k \geq 1$, the algorithm $\mathcal{A}_{\text{seq}}$ checks if

$$\hat{\mu}_{N_k} \in \big(r - \frac{M}{N_k}, r + \frac{M}{N_k}\big). \tag{2.36}$$

If the empirical average lies in this interval, then the algorithm continues sampling; otherwise, it outputs H_0 or H_1 accordingly in the natural way.

For the sequential algorithm $\mathcal{A}_{\text{hard}}$, let $N_0 > 0$ be chosen arbitrarily,[2] and let N_k be defined in terms of N_0 as in (2.35). It once again decides at each N_k for $k \geq 1$ whether to output an answer or to continue sampling, depending on whether

$$\hat{\mu}_{N_k} \in \big(r - \epsilon_k\big(\alpha\big),\ r + \epsilon_k\big(\alpha\big)\big).$$

When this inclusion holds, the algorithm continues; when it doesn't hold, the algorithm outputs H_0 or H_1 in the natural way. The following result is restricted to the stationary case; later in the section, we turn to the question of burn-in.

[2] In Gyori and Paulin [46], the authors set $N_0 = \big\lfloor \frac{100}{\gamma_0} \big\rfloor$, but this is inessential.

Theorem 3. *Assume that* $\alpha \le \frac{2}{5}$. *For* $\mathcal{A}_{\text{fixed}}, \mathcal{A}_{\text{seq}}, \mathcal{A}_{\text{hard}}$ *to all satisfy* $\operatorname{err}\big(\mathrm{A},\ f\big) \le \alpha$, *it suffices to (respectively) choose*

$$N = \frac{4T_f\big(\delta\big)\log\big(\frac{1}{\alpha}\big)}{\delta^2}, \tag{2.37}$$

$$M = \frac{16T_f\big(\frac{\delta}{2}\big)\log\big(\frac{2}{\sqrt{\alpha}\xi}\big)}{\delta},\ \textit{and} \tag{2.38}$$

$$\epsilon_k\big(\alpha\big) = \inf\Big\{\epsilon > 0\colon \frac{\epsilon^2}{16T_f\big(\frac{\epsilon}{2}\big)} \ge \frac{\log\big(1/\alpha\big) + 1 + 2\log k}{N_k}\Big\}, \tag{2.39}$$

where we let $\inf\emptyset = \infty$.

Our results differ from those of [46] because the latter implicitly control the worst-case error of the algorithm

$$\operatorname{err}\big(\mathrm{A}\big) = \sup_{f\colon \Omega\to[0,\ 1]} \operatorname{err}\big(\mathrm{A},\ f\big),$$

while our analysis controls $\operatorname{err}\big(\mathrm{A},\ f\big)$ directly. The corresponding choices made in [46] are

$$N = \frac{\log(1/\alpha)}{\gamma_0\delta^2}, M = \frac{\log(\frac{2}{\sqrt{\alpha}\xi})}{\gamma_0\delta},\ \text{and}\ \epsilon_k(\alpha) = \sqrt{\frac{\log(1/\alpha) + 1 + 2\log k}{\gamma_0 N_k}}.$$

Hence, the T_f parameter in our bounds plays the same role that $\frac{1}{\gamma_0}$ plays in their uniform bounds. As a result of this close correspondence, we easily see that our results improve on the uniform result for a fixed function f whenever it converges to its stationary expectation faster than the chain itself converges; i.e., whenever $T_f\big(\delta\big) \le \frac{1}{2\gamma_0}$.

The value of the above tests depends substantially on their sample size requirements. In setting (a), algorithm $\mathcal{A}_{\text{seq}}$ is only valuable if it reduces the number of samples needed compared to $\mathcal{A}_{\text{fixed}}$. In setting (b), algorithm $\mathcal{A}_{\text{hard}}$ is valuable because of its ability to test between hypotheses separated only by a point, but its utility is limited if it takes too long to run. Therefore, we now turn to the question of bounding expected stopping times.

In order to carry out the stopping time analysis, we introduce the true margin $\Delta = |r - \mu|$. First, let us introduce some useful notation. Let $N(\mathrm{A})$ be the number of sampled collected by A. Given a margin schedule $\big(\epsilon_k\big)$, let

$$k_0^*\big(\epsilon_{1:\infty}\big) := \min\big\{k \ge 1\colon \epsilon_k \le \frac{\Delta}{2}\big\},\ \text{and}\ N_0^*\big(\epsilon_{1:\infty}\big) := N_{k_0^*\left(\epsilon_{1:\infty}\right)}.$$

We can bound the expected stopping times of $\mathcal{A}_{\text{seq}}, \mathcal{A}_{\text{hard}}$ in terms of Δ as follows:

Theorem 4. *Assume either* H_0 *or* H_1 *holds. Then,*

$$\mathbb{E}\big[N(\mathcal{A}_{\text{seq}})\big] \le \big(1+\xi\big)\big[\frac{M}{\Delta} + \frac{4}{\Delta}\sqrt{\frac{2T_f\big(\delta/2\big)M}{\Delta} + 8T_f\big(\delta/2\big)} + 1\big]; \tag{2.40}$$

$$\mathbb{E}\big[N(\mathcal{A}_{\text{hard}})\big] \le \big(1+\xi\big)\big(N_0^* + 1\big) + \frac{32\alpha T_f\big(\Delta/4\big)}{\Delta^2}. \tag{2.41}$$

With minor modifications to the proofs in [46], we can bound the expected stopping times of their procedures as

$$\mathbb{E}[N(\mathcal{A}_{\text{seq}})] \leq (1+\xi)\Big\{\frac{M}{\Delta} + \frac{2}{\Delta}\sqrt{\frac{M}{\gamma_0\Delta} + \frac{4}{\gamma_0}} + 1\Big\};$$
$$\mathbb{E}[N(\mathcal{A}_{\text{hard}})] \leq (1+\xi)\big(N_0^* + 1\big) + \frac{4\alpha}{\gamma_0\Delta^2}.$$

In order to see how the uniform and adaptive bounds compare, it is helpful to first note that, under either H_0 or H_1, we have the lower bound $\Delta \geq \delta$. Thus, the dominant term in the expectations in both cases is $(1+\xi)M/\Delta$. Consequently, the ratio between the expected stopping times is approximately equal to the ratio between the M values; viz.,

$$\frac{M_{\text{adapt}}}{M_{\text{unif}}} \approx \gamma_0 T_f\big(\delta/2\big). \tag{2.42}$$

As a result, we should expect a significant improvement in terms of number of samples when the relaxation time $\frac{1}{\gamma_0}$ is significantly larger than the f-mixing time $T_f\big(\delta/2\big)$. Framed in absolute terms, we can write

$$\bar{N}_{\text{unif}}(\mathcal{A}_{\text{seq}}) \approx \frac{\log\big(2/\sqrt{\alpha\xi}\big)}{\gamma_0\delta\Delta} \quad \text{and} \quad \bar{N}_{\text{adapt}}(\mathcal{A}_{\text{seq}}) \approx \frac{T_f\big(\delta/2\big)\log\big(2/\sqrt{\alpha\xi}\big)}{\delta\Delta}.$$

Up to an additive term, the bound for $\mathcal{A}_{\text{hard}}$ is also qualitatively similar to earlier ones, with $\frac{1}{\delta\Delta}$ replaced by $\frac{1}{\Delta^2}$.

2.5 Analyzing mixing in practice

We analyze several examples of MCMC-based Bayesian analysis from our theoretical perspective. These examples demonstrate that convergence in discrepancy can in practice occur much faster than suggested by naive mixing time bounds and that our bounds help narrow the gap between theoretical predictions and observed behavior. Two of these examples appear in the following sections, while a third is relegated to Section 2.5.

Bayesian logistic regression

Our first example is a Bayesian logistic regression problem introduced by Robert and Casella [82]. The data consists of 23 observations of temperatures (in Fahrenheit, but normalized by dividing by 100) and a corresponding binary outcome—failure ($y = 1$) or not ($y = 0$) of a certain component; the aim is to fit a logistic regressor, with parameters $\big(\alpha,\ \beta\big) \in \mathbb{R}^2$, to the data, incorporating a prior and integrating over the model uncertainty to obtain future

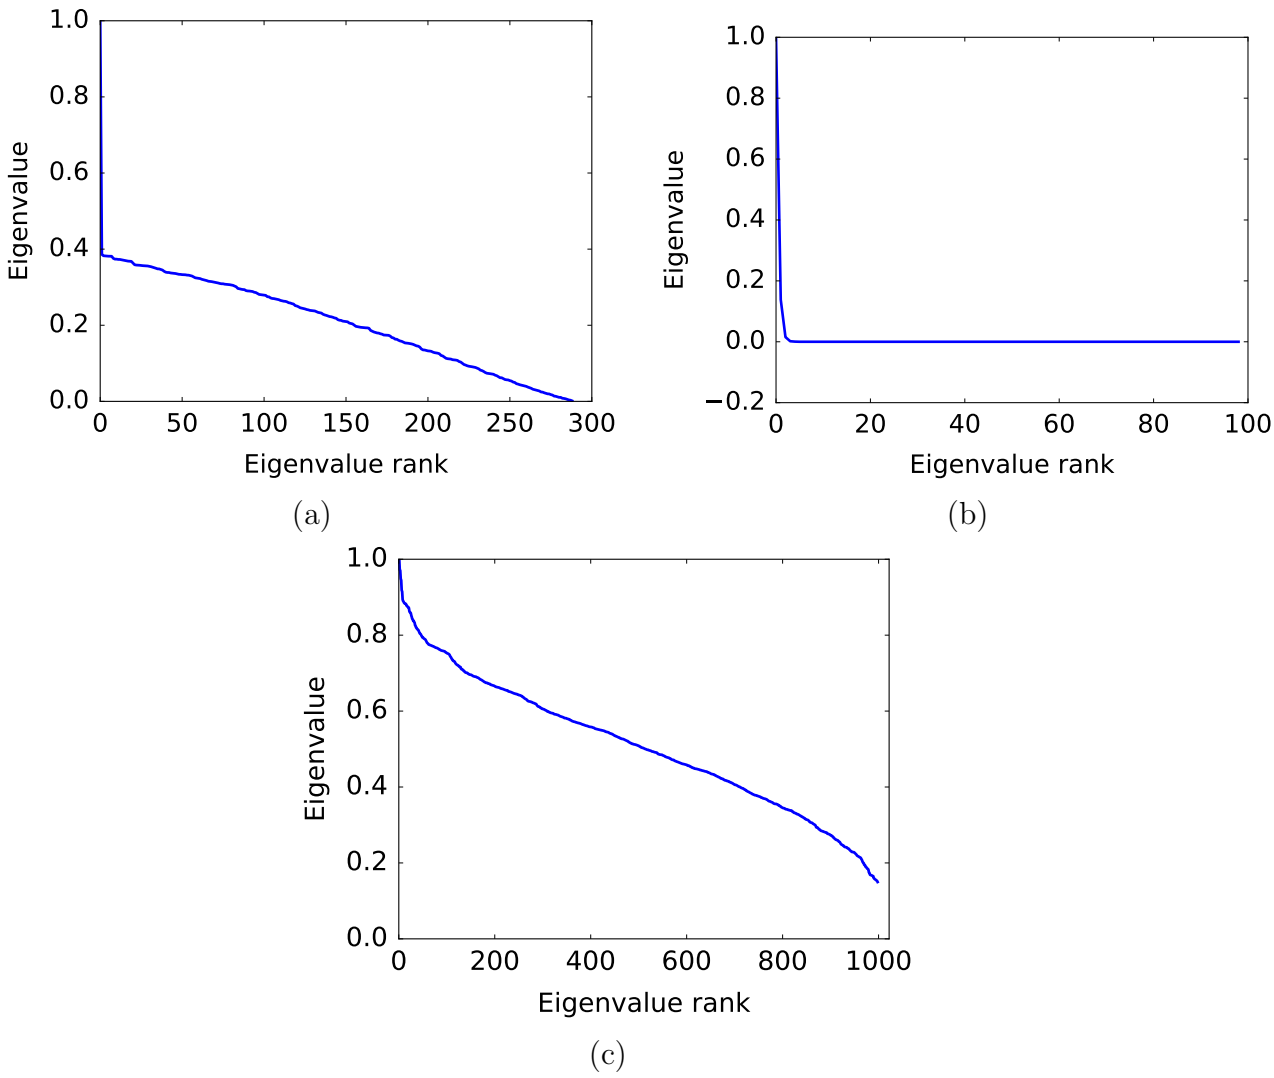

Figure 2.1: Spectra for three example chains: (a) Metropolis-Hastings for Bayesian logistic regression (Section 2.5); (b) collapsed Gibbs sampler for missing data imputation (Section 2.5); and (c) collapsed Gibbs sampler for a mixture model (Section 2.5).

predictions. More explicitly, following the analysis in Gyori and Paulin [45], we consider the following model:

$$p\big(\alpha,\ \beta \mid b\big) = \frac{1}{b} \cdot e^{\alpha} \exp\big(-e^{\alpha}/b\big)$$
$$p\big(y \mid \alpha,\ \beta,\ x\big) \propto \exp\big(\alpha + \beta x\big),$$

which corresponds to an exponential prior on e^{α}, an improper uniform prior on β and a logit link for prediction. As in Gyori and Paulin [45], we target the posterior by run-

ning a Metropolis-Hastings algorithm with a Gaussian proposal with covariance matrix $\Sigma = \begin{pmatrix} 4 & 0 \\ 0 & 10 \end{pmatrix}$. Unlike in their paper, however, we discretize the state space to facilitate exact analysis of the transition matrix and to make our theory directly applicable. The resulting state space is given by

$$\Omega = \Big\{ \big(\hat{\alpha} \pm i \cdot \Delta,\ \hat{\beta} \pm j \cdot \Delta\big) \ \mid\ 0 \le i,\ j \le 8 \Big\},$$

where $\Delta = 0.1$ and $(\hat{\alpha},\ \hat{\beta})$ is the MLE. This space has $d = 17^2 = 289$ elements, resulting in a 289×289 transition matrix that can easily be diagonalized.

Robert and Casella [82] analyze the probability of failure when the temperature x is 65°F; it is specified by the function

$$f_{65}\big(\alpha,\ \beta\big) = \frac{\exp\big(\alpha + 0.65\beta\big)}{1 + \exp\big(\alpha + 0.65\beta\big)}.$$

Note that this function fluctuates significantly under the posterior, as shown in Figure 2.2.

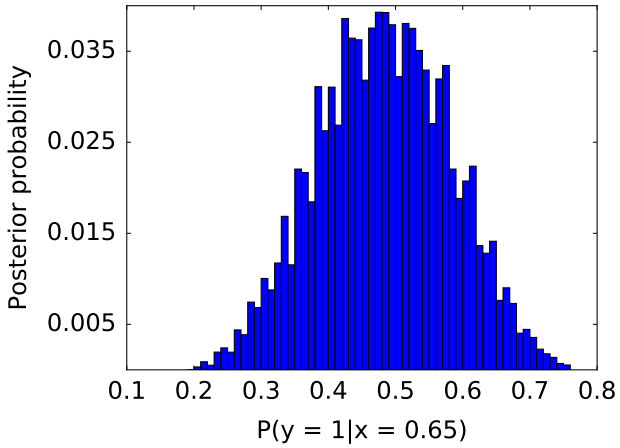

Figure 2.2: Distribution of f_{65} values under the posterior. Despite the discretization and truncation to a square, it generally matches the one displayed in Figure 1.2 in Robert and Casella [82].

We find that this function also happens to exhibit rapid mixing. The discrepancy $d_{f_{65}}$, before entering an asymptotic regime in which it decays exponentially at a rate $1-\gamma^* \approx 0.386$, first drops from about 0.3 to about 0.01 in just two iterations, compared to the predicted ten iterations from the naive bound $d_f\big(n\big) \le d_{\mathrm{TV}}\big(n\big) \le \frac{1}{\sqrt{\pi_{\min}}} \cdot \big(1-\gamma^*\big)^n$. Figure 2.3 demonstrates this on a log scale, comparing the naive bound to a version of the bound in Lemmas 1 and 2. Note that the oracle f-discrepancy bound improves significantly over the uniform baseline, even though the non-oracle version does not. In this calculation, we took $J = \big\{2, \ldots, 140\big\}$ to include the top half of the spectrum excluding 1 and computed $\|h_j\|_\infty$ directly from P

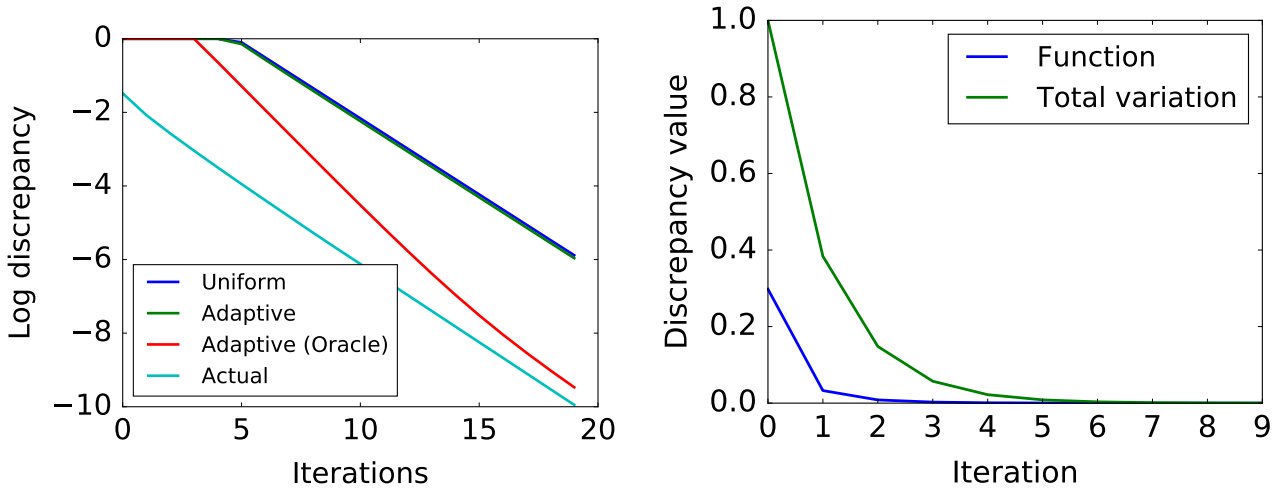

(a) (b)

Figure 2.3: (a) Discrepancies (plotted on log-scale) for f_{65} as a function of iteration number. The prediction of the naive bound is highly pessimistic; the f-discrepancy bound goes part of the way toward closing the gap and the oracle version of the f-discrepancy bound nearly completely closes the gap in the limit and also gets much closer to the right answer for small iteration numbers. (b) Comparison of the function discrepancy $d_{f_{65}}$ and the total variation discrepancy d_{TV}. They both decay fairly quickly due to the large spectral gap, but the function discrepancy still falls much faster.

for $j \in J$ and likewise for $q_j^T f_{65}$. The oracle bound is given by Lemma 2. As shown in panel (b) of Figure 2.3, this decay is also faster than that of the total variation distance.

An important point is that the quality of the f-discrepancy bound depends significantly on the choice of J. In the limiting case where J includes the whole spectrum below the top eigenvalue, the oracle bound becomes exact. Between that and $J = \emptyset$, the oracle bound becomes tighter and tighter, with the rate of tightening depending on how much power the function has in the higher versus lower eigenspaces. Figure 2.4 illustrates this for a few settings of J, showing that although for this function and this chain, a comparatively large J is needed to get a tight bound, the oracle bound is substantially tighter than the uniform and non-oracle f-discrepancy bounds even for small J.

Collapsed Gibbs sampling for mixture models

Due to the ubiquity of clustering problems in applied statistics and machine learning, Bayesian inference for mixture models (and their generalizations) is a widespread application of MCMC [40, 44, 54, 69, 70]. We consider the mixture-of-Gaussians model, applying it to a subset of the schizophrenic reaction time data analyzed in Belin and Rubin [7]. The subset of the data we consider consists of 10 measurements, with 5 coming from healthy

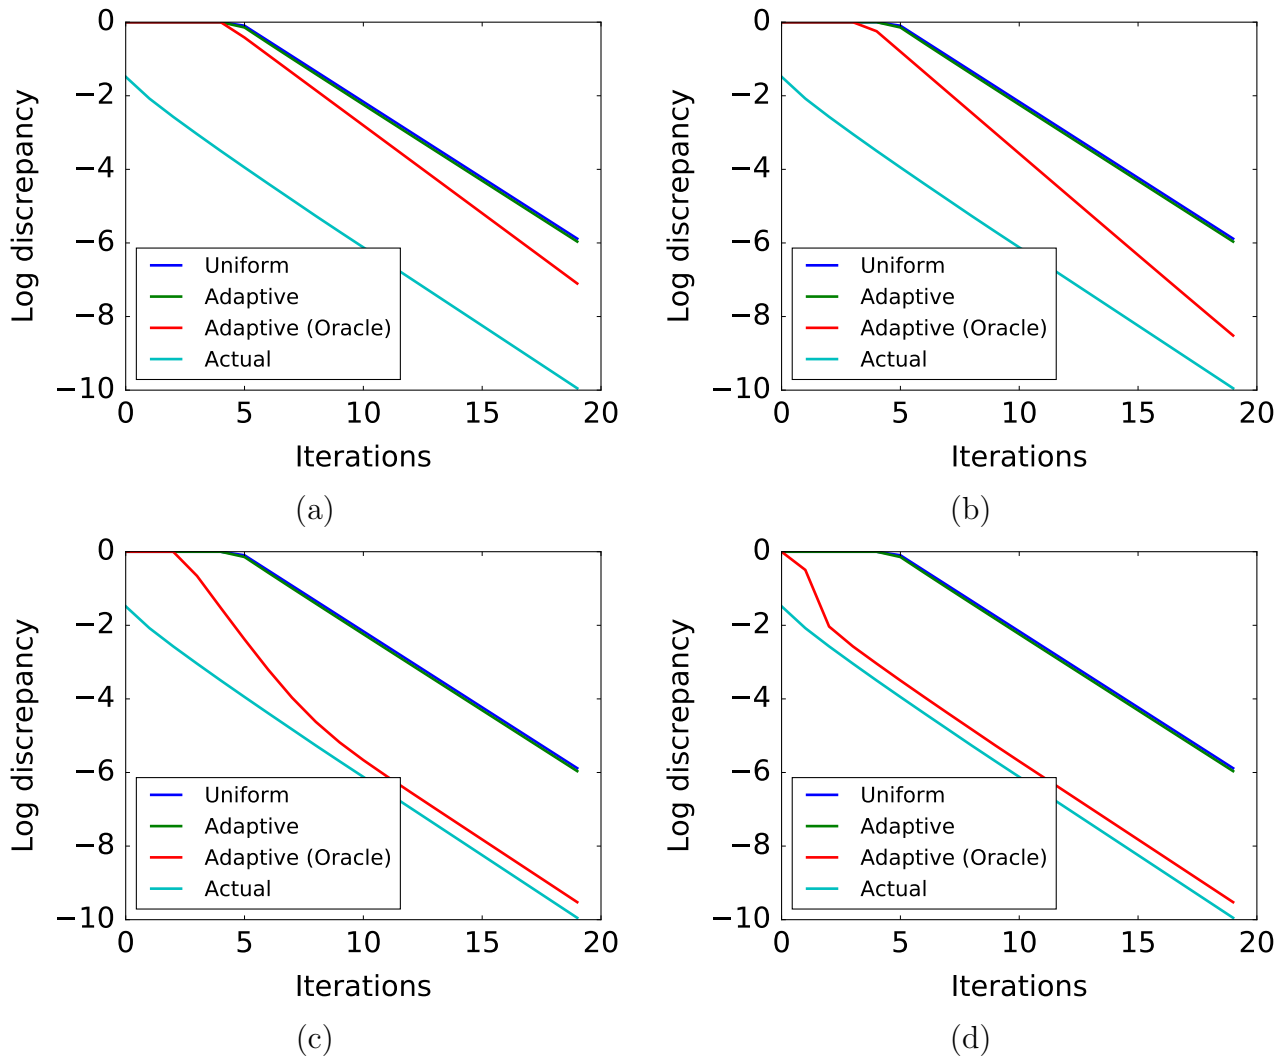

Figure 2.4: Comparisons of the uniform, non-oracle function-specific, and oracle function-specific bounds for various choices of J. In each case, $J = \{2, \ldots, J_{\max}\}$, with $J_{\max} = 50$ in panel (a), $J_{\max} = 100$ in panel (b), $J_{\max} = 200$ in panel (c), and $J_{\max} = 288$ in panel (d). The oracle bound becomes tight in the limit as $J_{\max}$ goes to $d = 289$, but it offers an improvement over the uniform bound across the board.

subjects and 5 from subjects diagnosed with schizophrenia. Since our interest is in contexts where uncertainty is high, we chose the 5 subjects from the healthy group whose reaction times were greatest and the 5 subjects from the schizophrenic group whose reaction times

were smallest. We considered a mixture with $K = 2$ components:

$$\begin{aligned} \mu_b &\sim \mathcal{N}\big(0,\ \rho^2\big),\ b = 0,\ 1, \\ \omega &\sim \mathrm{Be}\big(\alpha_0,\ \alpha_1\big) \\ Z_i \mid \omega &\sim \mathrm{Bern}\big(\omega\big) \\ X_i \mid Z_i = b,\ \mu &\sim \mathcal{N}\big(\mu_b,\ \sigma^2\big). \end{aligned}$$

We chose relatively uninformative priors, setting $\alpha_0 = \alpha_1 = 1$ and $\rho = 237$. Increasing the value chosen in the original analysis [7], we set $\sigma \approx 70$; we found that this was necessary to prevent the posterior from being too highly concentrated, which would be an unrealistic setting for MCMC. We ran collapsed Gibbs on the indicator variables Z_i by analytically integrating out ω and $\mu_{0:1}$.

As Figure 2.1 illustrates, the spectral gap for this chain is small—namely, $\gamma_* \approx 3.83 \times 10^{-4}$—yet the eigenvalues fall off comparatively quickly after λ_2, opening up the possibility for improvement over the uniform γ_*-based bounds. In more detail, define

$$z_b^* := \big(b \quad b \quad b \quad b \quad b \quad 1-b \quad 1-b \quad 1-b \quad 1-b \quad 1-b\big),$$

corresponding to the cluster assignments in which the patient and control groups are perfectly separated (with the control group being assigned label b). We can then define the indicator for exact recovery of the ground truth by

$$f(z) = \mathbf{1}\big(z \in \big\{z_0^*,\ z_1^*\big\}\big).$$

As Figure 2.5 illustrates, convergence in terms of f-discrepancy occurs much faster than convergence in total variation, meaning that predictions of required burn-in times and sample size based on global metrics of convergence drastically overestimate the computational and statistical effort required to estimate the expectation of f accurately using the collapsed Gibbs sampler. This behavior can be explained in terms of the interaction between the function f and the eigenspaces of P. Although the pessimistic constants in the bounds from the uniform bound (2.12) and the non-oracle function-specific bound (Lemma 1) make their predictions overly conservative, the oracle version of the function-specific bound (Lemma 2) begins to make exact predictions after just a hundred iterations when applied with $J = \big\{1, \ldots, 25\big\}$; this corresponds to making exact predictions of $T_f\big(\delta\big)$ for $\delta \leq \delta_0 \approx 0.01$, which is a realistic tolerance for estimation of μ. Panel (b) of Figure 2.5 documents this by plotting the f-discrepancy oracle bound against the actual value of d_f on a log scale.

The mixture setting also provides a good illustration of how the function-specific Hoeffding bounds can substantially improve on the uniform Hoeffding bound. In particular, let us compare the T_f-based Hoeffding bound (Theorem 1) to the uniform Hoeffding bound established by Léon and Perron [62]. At equilibrium, the penalty for non-independence in our bounds is $(2T_f(\epsilon/2))^{-1}$ compared to roughly γ_*^{-1} in the uniform bound. Importantly, however, our concentration bound applies unchanged even when the chain has not equilibrated, provided it has approximately equilibrated with respect to f. As a consequence, our

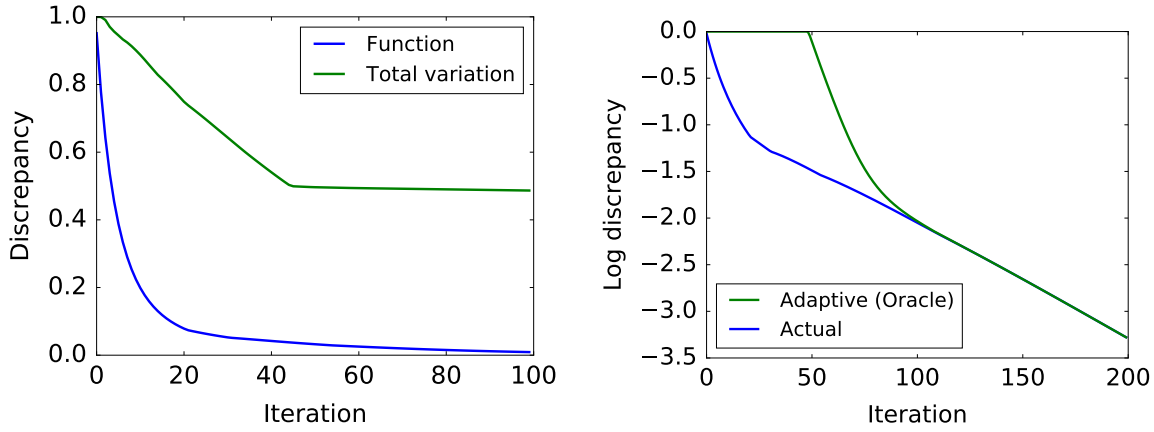

(a) (b)

Figure 2.5: (a) Comparison of the f-discrepancy d_f and the total variation discrepancy d_{TV} over the first 100 iterations of MCMC. Clearly the function mixes much faster than the overall chain. (b) The predicted value of $\log d_f$ (according to the f-discrepancy oracle bound—Lemma 2) plotted against the true value. The predictions are close to sharp throughout and become sharp at around 100 iterations.

Bound type	$T_f(0.01)$	$T_f(10^{-6})$
Uniform	31,253	55,312
Function-Specific	25,374	49,434
Function-Specific (Oracle)	98	409
Actual	96	409

Table 2.1: Comparison of bounds on $T_f(\delta)$ for different values of δ. The uniform bound corresponds to the bound $T_f(\delta) \leq T(\delta)$, the latter of which can be bounded by the total variation bound. The function-specific bounds correspond to Lemmas 1 and 2, respectively. Whereas the uniform and non-oracle f-discrepancy bounds make highly conservative predictions, the oracle f-discrepancy bound is nearly sharp even for δ as large as 0.01.

bound only requires a burn-in of $T_f(\epsilon/2)$, whereas the uniform Hoeffding bound does not directly apply for any finite burn-in. Table 2.1 illustrates the size of these burn-in times in practice. This issue can be addressed using the method of Paulin [75], but at the cost of a burn-in dependent penalty $d_{\mathrm{TV}}(T_0) = \sup_{\pi_0} d_{\mathrm{TV}}(\pi_n,\ \pi)$:

$$\mathbb{P}\Big[\frac{1}{N-T_0}\sum_{n=T_0}^{N} f(X_n) \geq \mu + \epsilon\Big] \leq d_{\mathrm{TV}}(T_0) + \exp\Big\{-\frac{\gamma_0}{2(1-\gamma_0)} \cdot \epsilon^2 [N - T_0]\Big\}, \qquad (2.43)$$

where we have let T_0 denote the burn-in time. Note that a matching bound holds for the lower

tail. For our experiments, we computed the tightest version of the bound (2.43), optimizing T_0 in the range $\left[0,\ 10^5\right]$ for each value of the deviation ϵ. Even given this generosity toward the uniform bound, the function-specific bound still outperforms it substantially, as Figure 2.6 shows.

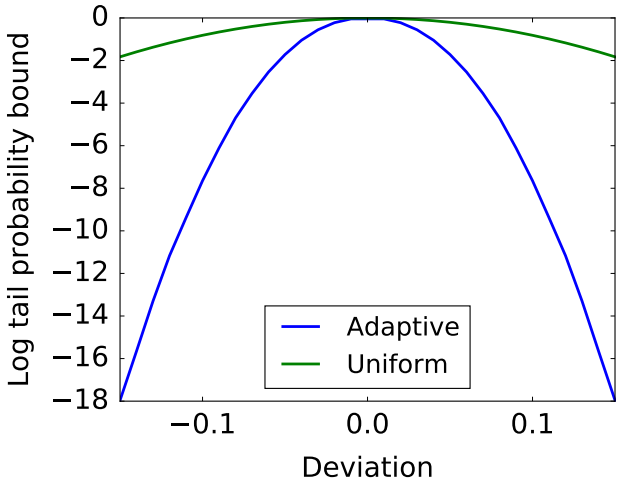

Figure 2.6: Comparison of the (log) tail probability bounds provided by the uniform Hoeffding bound due to [62] with one version of our function-specific Hoeffding bound (Theorem 1). Plots are based on $N = 10^6$ iterations, and choosing the optimal burn-in for the uniform bound and a fixed burn-in of $409 \geq T_f\left(10^{-6}\right)$ iterations for the function-specific bound. The function-specific bound improves over the uniform bound by orders of magnitude.

For the function-specific bound, we used the function-specific oracle bound (Lemma 2) to bound $T_f\left(\frac{\epsilon}{2}\right)$; this nearly coincides with the true value when $\epsilon \approx 0.01$ but deviates slightly for larger values of ϵ.

Bayesian analysis of clinical trials

The problem of missing data often necessitates Bayesian analysis, particularly in settings where uncertainty quantification is important, as in clinical trials. We illustrate how our framework would apply in this context by considering a clinical trials dataset [11, 45].

The dataset consists of $n = 50$ patients, some of whom participated in a trial for a drug and exhibited early indicators (Y_i) of success/failure and final indicators (X_i) of success/failure. Among the 50 patients, both indicator values are available for $n_X = 20$ patients; early indicators are available for $n_Y = 20$ patients; and no indicators are available for $n_0 = 10$ patients. The analysis depends on the following parameterization:

$$\begin{aligned}
\mathbb{P}\big(X_i = 1 \ \big|\ Y_i = 0\big) &= \gamma_0, \\
\mathbb{P}\big(X_i = 1 \ \big|\ Y_i = 1\big) &= \gamma_1, \\
\mathbb{P}\big(X_i = 1 \ \big|\ Y_i \text{ missing}\big) &= p.
\end{aligned}$$

Note that, in contrast to what one might expect, p is to be interpreted as the marginal probability that $X_i = 1$, so that in actuality $p = \mathbb{P}\big(X_i = 1\big)$ unconditionally; we keep the other notation, however, for the sake of consistency with past work [11, 45]. Conjugate uniform (i.e., $\mathrm{Be}\big(1,\ 1\big)$) priors are placed on all the model parameters.

The unknown variables include the parameter triple $\big(\gamma_0,\ \gamma_1,\ p\big)$ and the unobserved X_i values for $n_Y + n_0 = 30$ patients, and the full sample space is therefore $\tilde{\Omega} = [0,\ 1]^3 \times \big\{0,\ 1\big\}^{30}$. We cannot estimate the transition matrix for this chain, even with a discretization with as coarse a mesh as $\Delta = 0.1$, since the number of states would be $d = 10^3 \times 2^{30} \sim 10^{12}$. We therefore make two changes to the original MCMC procedure. First, we collapse out the X_i variables to bring the state space down to $[0,\ 1]^3$; while analytically collapsing out the discrete variables is impossible, we can estimate the transition probabilities for the collapsed chain analytically by sampling the X_i variables conditional on the parameter values and forming a Monte Carlo estimate of the collapsed transition probabilities. Second, since the function of interest in the original work—namely, $f\big(\gamma_0,\ \gamma_1,\ p\big) = \mathbf{1}\big(p > 0.5\big)$—depends only on p, we fix γ_0 and γ_1 to their MLE values and sample only p, restricted to the unit interval discretized with mesh $\Delta = 0.01$.

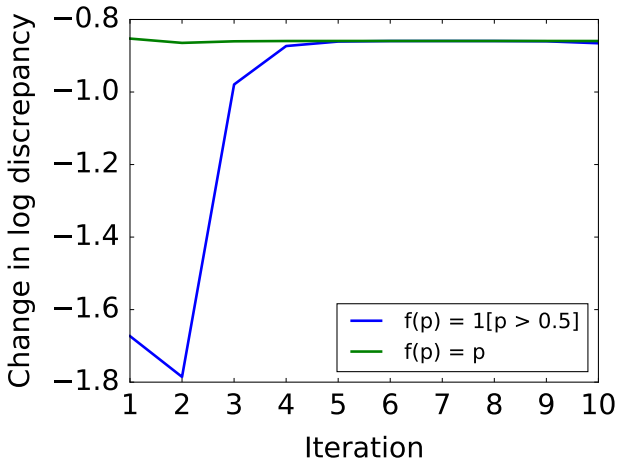

Figure 2.7: Change in log discrepancy for the two functions $f(p) = \mathbf{1}\big(p \geq 0.5\big)$ and $f(p) = p$ considered above. Whereas $f(p) = p$ always changes at the constant rate dictated by the spectral gap, the indicator discrepancy decays more quickly in the first few iterations.

As Figure 2.1 shows, eigenvalue decay occurs rapidly for this sampler, with $\gamma^* \approx 0.86$. Mixing thus occurs so quickly that none of the bounds—uniform or function-specific—get close to the truth, due to the presence of the constant terms (and specifically the large term $\frac{1}{\sqrt{\pi_{\min}}} \approx 2.14 \times 10^{33}$). Nonetheless, this example still illustrates how in actual fact, the choice of target function can make a big difference in the number of iterations required for accurate estimation; indeed, if we consider the two functions

$$f_1(p) := \mathbf{1}\big(p > 0.5\big), \quad \text{and} \quad f_2(p) := p,$$

we see in Figure 2.7 that the mixing behavior differs significantly between them: whereas the discrepancy for the second decays at the asymptotic exponential rate from the outset, the discrepancy for the first decreases faster (by about an order of magnitude) for the first few iterations, before reaching the asymptotic rate dictated by the spectral gap.

2.6 Additional proofs

Proofs for Section 2.2

In this section, we gather the proofs of the mixing time bounds from Section 2.2, namely equations (2.12) and (2.14) and Lemmas 1 and 2.

Proof of the bound (2.12)

Recall that

$$d_f(p,\ q) = \sup_{f\colon [d]\to[0,\ 1]} \left|\mathbb{E}_p\big[f\big(X\big)\big] - \mathbb{E}_q\big[f(Y)\big]\right|.$$

It follows from equation (2.14) that

$$\begin{aligned} d_{\mathrm{TV}}\big(\pi_n,\ \pi\big) &= \sup_{f\colon [d]\to[0,\ 1]} d_f\big(\pi_n,\ \pi\big) \\ &\leq \sup_{f\colon [d]\to[0,\ 1]} \big[\frac{\lambda_f^n}{\sqrt{\pi_{\min}}} \cdot d_f\big(\pi_0,\ \pi\big)\big] \\ &= \frac{1}{\sqrt{\pi_{\min}}} \cdot \lambda_*^n \cdot d_{\mathrm{TV}}\big(\pi_0,\ \pi\big), \end{aligned}$$

as claimed.

Proof of equation (2.14)

Let $D = \mathrm{diag}\big(\sqrt{\pi}\big)$. Then the matrix $A = DPD^{-1}$ is symmetric and so has an eigendecomposition of the form $A = \gamma_1\gamma_1^T + \sum_{j=2}^d \lambda_j\gamma_j\gamma_j^T$. Using this decomposition, we have

$$P = \mathbf{1}\pi^T + \sum_{j=2}^d \lambda_j h_j q_j^T,$$

where $h_j := D^{-1}\gamma_j$ and $q_j := D\gamma_j$. Note that the vectors $\{q_j\}_{j=2}^d$ correspond to the left eigenvectors associated with the eigenvalues $\{\lambda_j\}_{j=2}^d$.

Now, if we let π_0 be an arbitrary distribution over $[d]$, we have

$$d_f(\pi_n,\ \pi) = \left|\pi_0^T P^n f - \pi^T P^n f\right| \ \leq\ \left|(\pi_0 - \pi)^T P^n f\right|.$$

Defining $P_f := \mathbf{1}\pi^T + \sum_{j\in J_f} \lambda_j h_j q_j^T$, we have $P^n f = P_f^n f$. Moreover, if we define $\tilde{P}_f := \sum_{j\in J_f} \lambda_j h_j q_j^T$, and correspondingly $\tilde{A}_f := D\tilde{P}_f D^{-1}$, we then have the relation $\big(\pi_0 - \pi\big)^T \tilde{P}_f = \big(\pi_0 - \pi\big)^T P_f$. Consequently, by the definition of the operator norm and sub-multiplicativity, we have

$$\begin{aligned} d_f\big(\pi_n,\ \pi\big) &\leq \big|\big(\pi_0 - \pi\big)^T \tilde{P}_f^n f\big| \\ &\leq |\!|\!|\tilde{A}_f|\!|\!|_{\mathrm{op}}^n \|Df\|_2 \big\|D^{-1}\big(\pi_0 - \pi\big)\big\|_2 \\ &= \sqrt{\mathbb{E}_\pi\big[f^2\big] \cdot \sum_{i\in[d]} \frac{\big(\pi_{0,i} - \pi_i\big)^2}{\pi_i}} \cdot \lambda_f^n d_f\big(\pi_0,\ \pi\big). \end{aligned}$$

In order to complete the proof, let $Z \in \{0,\ 1\}^d$ denote the indicator vector $Z_j = \mathbf{1}\big(X_0 = j\big)$. Observe that the function

$$r(z) := \sum_{i\in[d]} \frac{\big(z_i - \pi_i\big)^2}{\pi_i}$$

is convex in terms of z. Thus, Jensen's inequality implies that

$$\mathbb{E}_{\pi_0}\big[r(Z)\big] \geq r\big(\mathbb{E}_{\pi_0}\big[Z\big]\big) = r\big(\pi_0\big) = \sum_{i\in[d]} \frac{\big(\pi_{0,i} - \pi_i\big)^2}{\pi_i}.$$

On the other hand, for any fixed value $X_0 = j$, corresponding to $Z = e_j$, we have

$$r\big(z\big) = r\big(e_j\big) = \frac{\big(1 - \pi_j\big)^2}{\pi_j} + \sum_{i\neq j} \pi_i = \frac{1 - \pi_j}{\pi_j} \leq \frac{1}{\pi_{\min}}.$$

We deduce that $d_f(\pi_n,\ \pi) \leq \sqrt{\frac{\mathbb{E}_\pi[f^2]}{\pi_{\min}}} \cdot \lambda_f^n \cdot d_f(\pi_0,\ \pi)$, as claimed.

Proof of Lemma 1

We observe that

$$\begin{aligned} \big|(\pi_0 - \pi)^T h_J(n)\big| &\leq \|\pi_0 - \pi\|_1\ \|h_J(n)\|_\infty \\ &= 2d_{\mathrm{TV}}\big(\pi_0,\ \pi\big)\ \|h_J(n)\|_\infty \\ &\leq 2d_{\mathrm{TV}}\big(\pi_0,\ \pi\big)\ \Big\{\sum_{j\in J} \big|\lambda_j\big|^n \cdot \big|q_j^T f\big| \cdot \|h_j\|_\infty\Big\} \\ &\leq 2d_{\mathrm{TV}}\big(\pi_0,\ \pi\big)\ \Big\{2\big|J\big| \cdot \max_{j\in J} \big|q_j^T f\big| \cdot \max_{j\in J} \|h_j\|_\infty\Big\}, \end{aligned}$$

as claimed.

Proof of Lemma 2

We proceed in a similar fashion as in the proof of equation (2.14). Begin with the identity proved there, viz.

$$d_f\big(\pi_n,\ \pi\big) = \big|(\pi_0 - \pi)^T \tilde{P}_f^n f\big|,$$

where $\tilde{P}_f = \sum_{j \in J_f} \lambda_j h_j q_j^T$. Now decompose $\tilde{P}_f$ further into

$$P_J = \sum_{j \in J} \lambda_j h_j q_j^T \quad \text{and} \quad P_{-J} = \sum_{j \in J_f \setminus J} \lambda_j h_j q_j^T.$$

Note also that $\tilde{P}_f^n = P_J^n + P_{-J}^n$. We thus find that

$$d_f(\pi_n,\ \pi) \leq \big|(\pi_0 - \pi)^T P_J^n f\big| + \big|(\pi_0 - \pi)^T P_{-J}^n f\big|.$$

Now observe that $P_J^n f = h_J(n)$, so $\big|(\pi_0 - \pi)^T P_J^n f\big| = \big|(\pi_0 - \pi)^T h_J(n)|$. On the other hand, the second term can be bounded using the argument from the proof of equation (2.14) to obtain

$$\big|(\pi_0 - \pi)^T P_{-J}^n f\big| \leq \sqrt{\frac{\mathbb{E}_\pi\big[f^2\big]}{\pi_{\min}}} \cdot \lambda_{-J_\delta}^n \cdot d_f(\pi_0,\ \pi),$$

as claimed.

Proofs for Section 2.2

In this section, we provide detailed proofs of the bound (2.26), as well as the other claims about the random function example on C_{2d}.

Proposition 2. *Let $f\colon [d] \to [0,\ 1]$ with $f(i) \sim \tau$ iid from some distribution on $[0,\ 1]$. There exists a universal constant $c_0 > 0$ such that with probability $\geq 1 - \frac{\delta^*}{128\sqrt{d \log d}}$ over the randomness f, we have*

$$T_f\big(\delta\big) \leq \frac{c_0 d \log d \log \frac{128d}{\delta}}{\delta^2} \qquad \textit{for all } 0 < \delta \leq \delta^*.$$

Proof. We proceed by defining a "good event" $\mathcal{D}_\delta$, and then showing that the stated bound on $T_f(\delta)$ holds conditioned on this event. The final step is to show that $\mathbb{P}[\mathcal{D}_\delta]$ is suitably close to one, as claimed.

The event $\mathcal{D}_\delta$ is defined in terms of the interaction between f and the eigenspaces of P corresponding to eigenvalues close to 1. More precisely, denote the indices of these eigenvalues by

$$J_\delta := \Big\{ j \in \{1, \ldots, 2d-1\} \ \mid \ j \leq 4\delta \sqrt{\frac{d}{\log d}} \ \text{ or } \ j \geq 2d - 4\delta \sqrt{\frac{d}{\log d}} \Big\}.$$

The good event $\mathcal{D}_\delta$ occurs when f has small inner product with all the corresponding eigenfunctions—that is

$$\mathcal{D}_\delta := \Big\{ \max_{j \in J_\delta} |q_j^T f| \leq 2\sqrt{\frac{10 \log d}{d}} \Big\}.$$

Viewed as family of events indexed by δ, these events form a decreasing sequence. (In particular, the associated sequence of sets J_δ is increasing in δ, in that whenever $\delta \leq \delta^*$, we are guaranteed that $J_\delta \subset J_{\delta^*}$.)

Establishing the bound conditionally on $\mathcal{D}_\delta$: We now exploit the spectral properties of the transition matrix to bound T_f conditionally on the event $\mathcal{D}_\delta$. Recall that the lazy random walk on C_{2d} has eigenvalues $\lambda_j = \frac{1}{2}\big(1 + \cos(\frac{\pi j}{d})\big)$ for $j \in [d]$, with corresponding unit eigenvectors

$$v_j^T = \frac{1}{\sqrt{2d}} \begin{pmatrix} 1 & \omega_j & \cdots & \omega_j^{2d-1} \end{pmatrix}, \quad \omega_j := e^{\frac{\pi i j}{d}}.$$

(See [63] for details.) We note that this diagonalization allows us to write $P = \mathbf{1}\pi^T + \sum_{j=1}^{2d-1} \lambda_j h_j q_j^T$, where $h_j = \sqrt{2d} \cdot v_j$ and $q_j = \frac{v_j}{\sqrt{2d}}$, where we have used the fact that $\operatorname{diag}\big(\sqrt{\pi}\big) = \frac{1}{\sqrt{2d}} \cdot I$. Note that $\|h_j\|_\infty = 1$.

Combining Lemma 1 with the bounds $\lambda_{J_\delta} \leq 1$, $\|h_j\|_\infty \leq 1$, and $\big|J_\delta\big| \leq 8\delta\sqrt{\frac{d}{\log d}}$, we find that

$$d_f(\pi_n,\ \pi) \leq 16\delta\sqrt{\frac{d}{\log d}} \cdot \max_{j \in J} \big|q_j^T f\big| + \sqrt{d} \cdot \lambda_{-J_\delta}^n.$$

Therefore, when the event $\mathcal{D}_\delta$ holds, we have

$$d_f(\pi_n,\ \pi) \leq 32\sqrt{10} \cdot \delta + \sqrt{d} \cdot \lambda_{-J_\delta}^n. \tag{2.44}$$

In order to conclude the argument, we use the fact that

$$\lambda_{-J_\delta} = \frac{1 + \max_{j \in J_f \setminus J_\delta} \cos\big(\frac{\pi j}{d}\big)}{2} \leq \frac{1 + \cos\big(\frac{\pi j_0}{d}\big)}{2},$$

where $j_0 = 4\delta\sqrt{\frac{d}{\log d}}$. On the other hand, we also have

$$\cos\big(\pi x\big) \leq 1 - \frac{\pi^2 x^2}{2} + \frac{\pi^4 x^4}{24} \leq 1 - \frac{\pi^2 x^2}{12}, \qquad \text{for all } |x| \leq 1,$$

which implies that

$$\lambda_{-J_\delta} \leq 1 - \frac{2\pi^2\delta^2}{3d \log d} \leq \exp\big(-\frac{2\pi^2\delta^2}{3d \log d}\big).$$

Together with equation (2.44), this bound implies that for $n \geq \frac{3d \log d \log \frac{d}{\delta}}{2\pi^2 \delta^2}$, we have $\sqrt{d}\lambda^n_{-J_\delta} \leq \delta$, whence

$$d_f(\pi_n,\ \pi) \leq (32\sqrt{10}+1)\,\delta \leq 128\delta.$$

Replacing δ by $\frac{\delta}{128}$ throughout, we conclude that for

$$n \geq \frac{3\,(128)^2 d \log d \log \frac{128d}{\delta}}{2\pi^2\delta^2} = \frac{3 \cdot 2^{13}}{\pi^2} \cdot \frac{d \log d \log \frac{128d}{\delta}}{\delta^2},$$

we have $d_f(\pi_n,\ \pi) \leq \delta$ with probability at least $\mathbb{P}\left(\mathcal{D}_{\delta/128}\right)$.

Controlling the probability of $\mathcal{D}_\delta$: It now suffices to prove $\mathbb{P}(\mathcal{D}_\delta) \geq 1 - \frac{\delta}{\sqrt{d \log d}}$, since this implies that $\mathbb{P}\left(\mathcal{D}_{\delta/128}\right) \geq 1 - \frac{\delta}{128\sqrt{d \log d}}$, as required. In order to do so, observe that the vectors $\{q_j\}_{j=1}^d$ are rescaled versions of an orthonormal collection of eigenvectors, and hence

$$\mathbb{E}\big[q_j^T f\big] = \mathbb{E}_\nu\,[\mu] \cdot q_j^T \mathbf{1} = 0.$$

We can write the inner product as $q_j^T f = A_j + iB_j$, where (A_j, B_j) are a pair of real numbers. The triangle inequality then guarantees that $|q_j^T f| \leq |A_j| + |B_j|$, so that it suffices to control these two absolute values.

By definition, we have

$$A_j = \frac{1}{2d} \sum_{\ell=0}^{2d-1} f\big(\ell\big) \cdot \cos\big(\frac{\pi j \ell}{d}\big),$$

showing that it is the sum of sub-Gaussian random variables with parameters $\sigma^2_{\ell,j} = \cos^2\left(\frac{\pi j \ell}{d}\right) \leq 1$. Thus, the variable A_j is sub-Gaussian with parameter at most $\sigma_j^2 \leq \frac{1}{2d}$. A parallel argument applies to the scalar B_j, showing that it is also sub-Gaussian with parameter at most σ_j^2.

By the triangle inequality, we have $|q_j^T f| \leq |A_j| + |B_j|$, so it suffices to bound $\big|A_j\big|$ and $|B_j|$ separately. In order to do so, we use sub-Gaussianity to obtain

$$\mathbb{P}\big(\max_{j \in J} |A_j| \geq r\big) \leq |J| \cdot e^{-\frac{r^2}{2}} \leq 8\delta\sqrt{\frac{d}{\log d}} \cdot e^{-\frac{dr^2}{2}}.$$

With $r := \sqrt{\frac{2\log 16d}{d}}$, we have

$$\mathbb{P}\big(\max_{j \in J_\delta} \big|A_j\big| \geq \sqrt{\frac{2\log 16d}{d}}\big) \leq \frac{\delta}{2\sqrt{d \log d}}.$$

Applying a similar argument to B_j and taking a union bound, we find that

$$\mathbb{P}\big(\max_{j\in J_\delta}\big|q_j^T f\big| \geq 2\sqrt{\frac{2\log 16d}{d}}\big) \leq \frac{\delta}{\sqrt{d\log d}}.$$

Since $2\sqrt{\frac{2\log 16d}{d}} \leq 2\sqrt{\frac{10\log d}{d}}$ for $d \geq 2$, we deduce that

$$1-\mathbb{P}\left(\mathcal{D}_\delta\right) = \mathbb{P}\left(\max_{j\in J_\delta}\big|q_j^T f\big| \geq 2\sqrt{\frac{10\log d}{d}}\right) \leq \frac{\delta}{\sqrt{d\log d}},$$

as required. □

The concentration result now follows.

Proposition 3. *The random function f on C_{2d} defined in equation* (2.23) *satisfies the mixing time and tail bounds*

$$T_f\big(\frac{\epsilon}{2}\big) \leq \frac{c_0 d\log d\big[\log d + \log\big(\frac{1}{\epsilon^2}\big)\big]}{\epsilon^2},$$

and

$$\mathbb{P}\left[\frac{1}{N}\sum_{n=T_f\left(\epsilon/2\right)}^{N+T_f(\epsilon/2)} f\big(X_n\big) \geq \mu+\epsilon\right] \leq \exp\big(-\frac{c_1\epsilon^4 N}{d\log d\big[\log(\frac{1}{\epsilon})+\log d\big]}\big).$$

with probability at least $1-\frac{c_2\epsilon^2}{\sqrt{d\log d}}$ over the randomness of f provided $\epsilon \geq c_3\big(\frac{\log d}{d}\big)^{1/2}$, where $c_0, c_1, c_2, c_3 > 0$ are universal constants.

Proof. We first note that from the proof of Proposition 2, we have the lower bound $1-\lambda_{-J_\delta} \geq \frac{c_4\delta^2}{d\log d}$, valid for all $\delta \in (0,1)$. The proof of the previous proposition guarantees that $\Delta_J^* \leq 32\sqrt{10}\delta$, so setting $\delta = \frac{\epsilon}{128\sqrt{10}}$ yields

$$\frac{\epsilon}{4} = 32\sqrt{10}\delta \geq \Delta_J^*, \quad \text{and} \quad 1-\lambda_{-J_\delta} \geq \frac{c_4'\epsilon^2}{d\log d}.$$

Now, by Proposition 2, there is a universal constant $c_5 > 0$ such that, with probability at least $1-\frac{\delta}{128\sqrt{d\log d}}$, we have

$$T_f\big(\delta'\big) \leq \frac{c_5 d\log d\log d/\delta'}{\big(\delta'\big)^2} \qquad \text{for all } \delta' \geq \delta.$$

In particular, we have

$$T_f\big(\frac{\epsilon}{2}\big) \leq \frac{c_2' d \log d \log d/\epsilon}{\epsilon^2}$$

with this same probability. Thus, we have this bound on T_f with the high probability claimed in the statement of the proposition.

We now finish by taking $\Delta = \frac{\epsilon}{4}$ in Corollary 2. Noting that $\Delta_J + \Delta = \frac{\epsilon}{2}$ and $1 - \lambda_{-J} \geq \frac{c_4' \epsilon^2}{d \log d}$ completes the proof. □

Proofs for Section 2.4

We begin with the proof of the confidence interval corresponding to our bound (Theorem 2). Proofs of the claims (2.29b) and (2.33) can be found in Sections 2.6 and 2.6 respectively.

Proof of Theorem 2

As discussed in Section 2.4, we actually prove a somewhat stronger form of Theorem 2, in order to guarantee that the confidence interval can be straightforwardly built using an upper bound $\tilde{T}_f$ on the f-mixing time rather than the true value. Setting $\tilde{T}_f = T_f$ recovers the original theorem.

Specifically, suppose $\tilde{T}_f \colon \mathbb{N} \to \mathbb{R}_+$ is an upper bound on T_f and note that the corresponding tail bound becomes $e^{-\tilde{r}_N(\epsilon)/8}$, where

$$\tilde{r}_N(\epsilon) = \epsilon^2 \left[\frac{N}{\tilde{T}_f\big(\frac{\epsilon}{2}\big)} - 1 \right].$$

This means that, just as before we wanted to make the rate r_N in equation (2.30) at least as large as $8 \log \frac{2}{\alpha}$, we now wish to do the same with $\tilde{r}_N$, which means choosing ϵ_N with $\tilde{r}_N\big(\epsilon_N\big) \geq 8 \log \frac{2}{\alpha}$. We therefore have the following result.

Proposition 4. *For any width $\epsilon_N \in \tilde{r}_N^{-1}\big(\big[8 \log\big(2/\alpha\big),\ \infty\big)\big)$, the set*

$$I_N^{\text{func}} = \left[\frac{1}{N - T_f\big(\frac{\epsilon}{2}\big)} \sum_{n=\tilde{T}_f\left(\frac{\epsilon}{2}\right)}^{N} f\big(X_n\big) \pm \epsilon_N \right]$$

is a $1-\alpha$ confidence interval for $\mu = \mathbb{E}_\pi\big[f\big]$.

Proof. For notational economy, let us introduce the shorthands $\tau_f(\epsilon) = T_f\big(\frac{\epsilon}{2}\big)$ and $\tilde{\tau}_f\big(\epsilon\big) = \tilde{T}_f\big(\frac{\epsilon}{2}\big)$. Theorem 1 then implies

$$\begin{aligned}\mathbb{P}\left[\frac{1}{N-\tilde{\tau}_f}\sum_{n=\tilde{\tau}_f}^{N} f\big(X_n\big) \geq \mu + \epsilon\right] &\leq \exp\big(-\frac{N-\tau_f}{4\tau_f}\cdot\epsilon^2\big)\\ &\leq \exp\big(-\frac{N-\tilde{\tau}_f}{4\tilde{\tau}_f}\cdot\epsilon^2\big)\\ &= \exp\big(-\frac{\tilde{r}_N\big(\epsilon\big)}{4}\big).\end{aligned}$$

Setting $\epsilon = \epsilon_N$ yields

$$\mathbb{P}\left[\frac{1}{N-\tilde{\tau}_f}\sum_{n=\tilde{\tau}_f}^{N} f\big(X_n\big) \geq \mu + \epsilon_N\right] \leq \frac{\alpha}{2}.$$

The corresponding lower bound leads to an analogous bound on the lower tail. □

As we did with Corollary 3, we can derive a more concrete, though slightly weaker, form of this result that is more amenable to interpretation. We derive the corollary from the specialized bound by setting $\tilde{T}_f = T_f$.

To obtain this bound, define the following lower bound, in parallel with equation (2.31):

$$\tilde{r}_N\big(\epsilon\big) \geq \tilde{r}_{N,\eta}\big(\epsilon\big) := \epsilon^2\big[\frac{N}{\tilde{T}_f\big(\frac{\eta}{2}\big)} - 1\big],\ \epsilon \geq \eta.$$

Since this is a lower bound, we see that whenever $\epsilon_N \geq \eta$ and $\tilde{r}_{N,\eta}\big(\epsilon_N\big) \geq 8\log\frac{2}{\alpha}$, ϵ_N is a valid half-width for a $\big(1-\alpha\big)$-confidence interval for the stationary mean centered at the empirical mean. More formally, we have the following:

Proposition 5. *Fix $\eta > 0$ and let*

$$\epsilon_N = \tilde{r}^{-1}_{N,\eta}\big(8\log\frac{2}{\alpha}\big) = 2\sqrt{2}\sqrt{\frac{\tilde{T}_f\big(\frac{\eta}{2}\big)\cdot\log\big(2/\alpha\big)}{N-\tilde{T}_f\big(\frac{\eta}{2}\big)}}.$$

If $N \geq \tilde{T}_f\big(\frac{\eta}{2}\big)$, then I_N^{func} is a $1-\alpha$ confidence interval for $\mu = \mathbb{E}_\pi\big[f\big]$.

Proof. By assumption, we have

$$\eta \leq \epsilon_N\big(\eta\big) = 2\sqrt{\frac{\tilde{T}_f\big(\frac{\eta}{2}\big)\cdot\log\big(2/\alpha\big)}{N-\tilde{T}_f\big(\frac{\eta}{2}\big)}}.$$

This implies $\tilde{T}_f\big(\frac{\epsilon_N}{2}\big) \geq \tilde{T}_f\big(\frac{\eta}{2}\big)$, which yields

$$\tilde{r}_N\big(\epsilon_N\big) = \epsilon_N^2\big[\frac{N}{\tilde{T}_f\big(\frac{\epsilon_N}{2}\big)} - 1\big] \geq \epsilon_N^2\big[\frac{N}{\tilde{T}_f\big(\frac{\eta}{2}\big)} - 1\big] = 8\log\big(2/\alpha\big).$$

But now Proposition 4 applies, so that we are done. □

We now prove correctness of the confidence intervals based on the uniform Hoeffding bound (2.6), and the Berry-Esseen bound (2.32).

Proof of claim (2.29b)

This claim follows directly from a modified uniform Hoeffding bound, due to [75]. In particular, for any integer $T_0 \geq 0$, let $d_{\mathrm{TV}}(T_0) = \sup_{\pi_0} d_{\mathrm{TV}}(\pi_0 P^{T_0},\ \pi)$ be the worst-case total variation distance from stationarity after T_0 steps. Using this notation, [75] shows that for any starting distribution π_0 and any bounded function $f\colon [d] \to [0,\ 1]$, we have

$$\mathbb{P}\big(\big|\frac{1}{N-T_0}\sum_{n=T_0+1}^{N} f\big(X_n\big) - \mu\big| \geq \epsilon\big) \leq 2\exp\big(-\frac{\gamma_0}{2\big(2-\gamma_0\big)}\cdot\epsilon^2 N\big) + 2d_{\mathrm{TV}}\big(T_0\big). \tag{2.45}$$

We now use the bound (2.45) to prove our claim (2.29b). Recall that we have chosen T_0 so that $d_{\mathrm{TV}}\big(T_0\big) \leq \alpha_0/2$. Therefore, the bound (2.45) implies that

$$\begin{aligned}\mathbb{P}\left[\big|\frac{1}{N-T_0}\sum_{n=T_0+1}^{N} f(X_n) - \mu\big| \geq \epsilon_N\right] &\leq 2\exp\Big\{-\frac{\gamma_0}{2\big(2-\gamma_0\big)}\cdot\epsilon_N^2 N\Big\} + \alpha_0 \\ &\leq 2\cdot\frac{\alpha-\alpha_0}{2} + \alpha_0 = \alpha,\end{aligned}$$

as required.

Proof of the claim (2.33)

We now use the result (2.32) to prove the claim (2.33).

By the lower bound on N, we have

$$\frac{e^{-\gamma_0 N}}{3\sqrt{\pi_{\min}}} \leq \frac{\alpha}{6} \quad \text{and} \quad \frac{13}{\sigma_{f,\mathrm{asym}}\sqrt{\pi_{\min}}}\cdot\frac{1}{\gamma_0\sqrt{N}} \leq \frac{\alpha}{6}.$$

It follows from equation (2.32) that

$$\begin{aligned}\mathbb{P}\left[\frac{1}{\sigma_{f,\mathrm{asym}}N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + \epsilon_N\right] &\leq \Phi\big(\epsilon_N\sqrt{N}\big) + \frac{\alpha}{3} \\ &\leq \exp\big(-\frac{N}{2}\cdot\epsilon_N^2\big) + \frac{\alpha}{3} \\ &= \frac{\alpha}{2},\end{aligned}$$

and since a matching bound holds for the lower tail, we get the desired result.

Proofs for Section 2.4

In this section, we collect various proofs associated with our analysis of the sequential testing problem.

Proof of Theorem 3 for $\mathcal{A}_{\text{fixed}}$

We provide a detailed proof when H_1 is true, in which case we have $\mu \leq r-\delta$; the proof for the other case is analogous. When H_1 is true, we need to control the probability $\mathbb{P}\big(\mathcal{A}_{\text{fixed}}\big(X_{1:N}\big) = H_0\big)$. In order to do so, note that Theorem 1 implies that

$$\begin{aligned}\mathbb{P}\big(\mathcal{A}_{\text{fixed}}\big(X_{1:N}\big) = H_0\big) &= \mathbb{P}\big(\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq r + \delta\big) \\ &\leq \mathbb{P}\big(\frac{1}{N}\sum_{n=1}^{N} f\big(X_n\big) \geq \mu + 2\delta\big) \\ &\leq \exp\big(-\frac{\delta^2 N}{2T_f\big(\delta\big)}\big).\end{aligned}$$

Setting $N = \frac{2T_f(\delta)\log\left(\frac{1}{\alpha}\right)}{\delta^2}$ yields the bound $\mathbb{P}\big(\mathcal{A}_{\text{fixed}}\big(X_{1:N}\big) = H_0\big) \leq \alpha$, as claimed.

Proof of Theorem 3 for $\mathcal{A}_{\text{hard}}$

We may assume that H_1 holds, as the other case is analogous. Under H_1, letting k_0 be the smallest k such that $\epsilon_k < \infty$, we have

$$\text{err}(\mathcal{A}_{\text{hard}},\ f) \leq \sum_{k=k_0}^{\infty} \mathbb{P}\left(\hat{\mu}_{N_k} \geq r + \epsilon_k\right) \ \leq \ \sum_{k=k_0}^{\infty} \mathbb{P}(\hat{\mu}_{N_k} \geq \mu + 2\epsilon_k).$$

By Theorem 1, and the definition of ϵ_k, we thus have

$$\begin{aligned}\text{err}\left(\mathcal{A}_{\text{hard}},\ f\right) \leq \sum_{k=k_0}^{\infty} \exp\left(-\frac{N_k\epsilon_k^2}{8T_f\left(\frac{\epsilon_k}{2}\right)}\right) &\leq \frac{\alpha}{2}\sum_{k=k_0}^{\infty}\frac{1}{k^2} \\ &= \frac{\pi^2}{12}\,\alpha \ < \ \alpha,\end{aligned}$$

as claimed.

Proof of Theorem 4 for $\mathcal{A}_{\text{hard}}$

For concreteness, we may assume H_1 holds, as the H_0 case is symmetric. We now have that

$$\mathbb{P}\left[N \geq N_k\right] \leq \mathbb{P}\left[\left|\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) - r\right| \leq \epsilon_k\right] \ \leq \ \mathbb{P}\left[\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \geq \mu + \Delta - \epsilon_k\right].$$

For convenience, let us introduce the shorthand

$$T_{f,k}^{+} := \begin{cases} T_f\big(\frac{\Delta-\epsilon_k}{2}\big) & \text{if } \epsilon_k \leq \Delta, \\ 1 & \text{otherwise.} \end{cases}$$

Applying the Hoeffding bound from Theorem 1, we then find that

$$\mathbb{P}\left[N \geq N_k\right] \leq \exp\left\{-\frac{N_k}{8T_{f,k}^{+}} \cdot \big(\Delta - \epsilon_k\big)_{+}^{2}\right\}.$$

Observe further that

$$\begin{aligned} \mathbb{E}[N] &= N_1 + \sum_{k=1}^{\infty} \big(N_{k+1} - N_k\big)\mathbb{P}\big(N > N_k\big) \\ &\leq N_{k_0^*+1} + \sum_{k=k_0^*+1}^{\infty} \big(N_{k+1} - N_k\big)\mathbb{P}\big(N > N_k\big) \\ &\leq \big(1+\xi\big)\big(N_0^* + 1\big) + \sum_{k=k_0^*+1}^{\infty} \big(N_{k+1} - N_k\big)\mathbb{P}\big(N > N_k\big). \end{aligned}$$

Combining the pieces yields

$$\mathbb{E}[N] \leq \big(1+\xi\big)\big(N_0^* + 1\big) + \sum_{k=k_0^*+1}^{\infty} \big(N_{k+1} - N_k\big) \exp\big(-\frac{N_k}{8T_{f,k}^{+}} \cdot \big(\Delta - \epsilon_k\big)_{+}^{2}\big). \tag{2.46}$$

The crux of the proof is a bound on the infinite sum, which we pull out as a lemma for clarity.

Lemma 5. *The infinite sum* (2.46) *is upper bounded by*

$$\sum_{k=k_0^*+1}^{\infty} \big(N_{k+1} - N_k\big) \exp\big(-\frac{N_k}{8T_{f,k}^{+}} \cdot \big(\Delta - \epsilon_k\big)_{+}^{2}\big) \leq \alpha \cdot \sum_{m=1}^{\infty} \exp\big(-m \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big).$$

See Section 2.6 for the proof of this claim.

Lemma 5 then implies that

$$\begin{aligned} \sum_{k=k_0^*+1}^{\infty} \big(N_{k+1} - N_k\big) \exp\big(-\frac{N_k}{T_f\big(\frac{\Delta}{4}\big)} \cdot \frac{\Delta^2}{32}\big) &\leq \alpha \cdot \sum_{m=1}^{\infty} \exp\big(-m \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big) \\ &= \frac{\alpha \exp\big(-\frac{\Delta^2}{32T_f\left(\frac{\Delta}{4}\right)}\big)}{1 - \exp\big(\frac{\Delta^2}{32T_f\left(\frac{\Delta}{4}\right)}\big)} \\ &\leq \frac{32\alpha T_f\big(\frac{\Delta}{4}\big)}{\Delta^2}. \end{aligned}$$

The claim now follows from equation (2.46).

Proof of Theorem 3 for $\mathcal{A}_{\text{seq}}$

The proof is nearly identical to that given by [46], with $T_f\big(\delta/2\big)$ replacing $\frac{1}{\gamma_0}$. We again assume that H_1 holds, so $\mu \leq r - \delta$. In this case, it is certainly true that

$$\begin{aligned}
\text{err}\big(\mathcal{A}_{\text{seq}},\ f\big) &= \mathbb{P}\big(\exists k\colon \mathcal{A}_{\text{seq}}\big(X_{1:N_k}\big) = H_0\big) \\
&= \mathbb{P}\big(\exists k\colon \frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \geq r + \frac{M}{N_k}\big) \\
&\leq \sum_{k=1}^{\infty} \mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \geq r + \frac{M}{N_k}\big).
\end{aligned}$$

It follows by Theorem 1, with $\epsilon_k = \delta + \frac{M}{N_k}$, that

$$\begin{aligned}
\mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \geq r + \frac{M}{N_k}\big) &\leq \mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \geq \mu + \delta + \frac{M}{N_k}\big) \\
&\leq \exp\big(-\frac{\epsilon_k^2 N_k}{16 T_f\big(\frac{\epsilon_k}{2}\big)}\big) \\
&\leq \exp\big(-\frac{\epsilon_k^2 N_k}{16 T_f\big(\frac{\delta}{2}\big)}\big).
\end{aligned}$$

In order to simplify notation, for the remainder of the proof, we define $\tau := 8T_f(\delta/2)$, $\beta := \frac{\sqrt{\alpha\xi}}{2}$, and $\zeta_k := \frac{\delta^2 N_k}{2\tau \log(1/\beta)}$. In terms of this notation, we have $M = \frac{2\tau\log(1/\beta)}{\delta}$, and hence that

$$\begin{aligned}
\exp\big(-\frac{\epsilon_k^2 N_k}{2\tau}\big) &= \exp\big(-\frac{1}{2\tau}\cdot\big(\delta^2 N_k + 2\delta M + \frac{M^2}{N_k}\big)\big) \\
&= \exp\big(-\big[\frac{\delta^2 N_k}{2\tau} + \log\big(1/\beta\big) + \frac{2\tau\log^2\big(1/\beta\big)}{\delta^2 N_k}\big]\big) \\
&= \exp\big(-\log\big(1/\beta\big)\big[1 + \zeta_k + \zeta_k^{-1}\big]\big) \\
&= \beta\cdot\exp\big(-\log\big(1/\beta\big)\big[\zeta_k + \zeta_k^{-1}\big]\big).
\end{aligned}$$

It follows that the error probability is at most

$$\beta\sum_{k=1}^{\infty}\exp\big(-\log\big(1/\beta\big)\big[\zeta_k + \zeta_k^{-1}\big]\big).$$

We now finish the proof using two small technical lemmas, whose proofs we defer to Section 2.6.

Lemma 6. *In the above notation, we have*

$$\sum_{k=1}^{\infty}\exp\Big\{-\log(1/\beta)\big[\zeta_k + \zeta_k^{-1}\big]\Big\} \leq 4\sum_{\ell=0}^{\infty}\exp\Big\{-\log(1/\beta)\Big[\big(1+\xi\big)^{\ell} + \big(1+\xi\big)^{-\ell}\Big]\Big\}.$$

Lemma 7. *For any integer $c \geq 0$, we have*

$$(1+\xi)^{\ell} + (1+\xi)^{-\ell} \geq 2(c+1) \quad \textit{for all } \ell \in \Big[\tfrac{9c}{5\xi}, \tfrac{9(c+1)}{5\xi}\Big).$$

Using this bound, and grouping together terms in blocks of size $\frac{9}{5\xi}$, we find that the error is at most

$$4\sum_{\ell=0}^{\infty} \exp\big(-\log\big(1/\beta\big)\big[\big(1+\xi\big)^{\ell} + \big(1+\xi\big)^{-\ell}\big]\big) \leq \frac{36}{5\xi} \cdot \sum_{c=0}^{\infty} \beta^{2\left(c+1\right)}.$$

Since both α and ξ are at most $\frac{2}{5}$, we have $\beta = \frac{\sqrt{\alpha\xi}}{2} \leq \frac{1}{5}$, and hence the error probability is bounded as

$$\frac{36\beta}{5\xi}\sum_{c=0}^{\infty}\beta^{2\left(c+1\right)} \leq \frac{36\beta^3}{5\xi\big(1-\beta^2\big)} \leq \frac{36\beta^2}{25\xi\big(1-\beta^2\big)} \leq \frac{3\beta^2}{2\xi} = \frac{3\alpha}{4} < \alpha.$$

Proof of Theorem 4 for $\mathcal{A}_{\text{seq}}$

We may assume H_1 holds; the other case is analogous. Note that

$$\begin{aligned}
\mathbb{E}[N] &\leq N_1 + \sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\mathbb{P}\big(N > N_k\big) \\
&\leq N_1 + \sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) \in \big(r - \frac{M}{N_k},\ r + \frac{M}{N_k}\big)\big) \\
&\leq N_1 + \sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) > r - \frac{M}{N_k}\big) \\
&= N_1 + \sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\mathbb{P}\big(\frac{1}{N_k}\sum_{n=1}^{N_k} f\big(X_n\big) > \mu + \Delta - \frac{M}{N_k}\big) \\
&\leq N_1 + \sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\exp\left\{-\frac{\big(\Delta N_k - M\big)_+^2}{16T_f\big(\delta/2\big)N_k}\right\}.
\end{aligned}$$

Our proof depends on the following simple technical lemma, whose proof we defer to Section 2.6.

Lemma 8. *Under the conditions of Theorem 4, we have*

$$\sum_{k=1}^{\infty}\big(N_{k+1}-N_k\big)\exp\left\{-\frac{\big(\Delta N_k - M\big)_+^2}{16T_f\big(\delta/2\big)N_k}\right\} \leq \big(1+\xi\big)\big[1 + \int_{N_1}^{\infty} h(s)\mathrm{d}s\big], \tag{2.47}$$

where $h(s) := \exp\Big\{-\frac{(\Delta s - M)_+^2}{16T_f(\delta/2)s}\Big\}$.

Given this lemma, we then follow the argument of Gyori and Paulin [46] in order to bound the integral. We have

$$\int_{N_1}^{\infty} h\big(s\big)\mathrm{d}s \leq \frac{4}{\Delta}\sqrt{\frac{4T_f\big(\delta/2\big)M}{\Delta} + 16T_f\big(\delta/2\big)}.$$

To conclude, note that either $r \geq \Delta$ or $1-r \geq \Delta$, since $0 < \mu < 1$, so that $\min\big(\frac{1}{r},\ \frac{1}{1-r}\big) \leq \frac{1}{\Delta}$. It follows that

$$N_1 \leq \big(1+\xi\big)N_0 \leq \frac{(1+\xi)M}{\Delta}.$$

Combining the bounds yields the desired result.

Proofs of Lemmas for sequential testing

In this section, we gather the proofs of Lemmas 6–5.

Proof of Lemma 6: Observe that the function

$$g(\zeta) := \exp\Big\{ -\log(1/\beta)\,(\zeta + \zeta^{-1})\Big\}$$

is increasing on $(0,\ 1]$ and decreasing on $[1,\ \infty)$. Therefore, bringing ζ closer to 1 can only increase the value of the function.

Now, for fixed $k \geq 1$, define

$$\ell_k := \begin{cases} \min\big\{\ell\colon \big(1+\xi\big)^{\ell} \geq \zeta_k\big\} & \text{if } \zeta_k \leq 1, \\ \max\big\{\ell\colon \big(1+\xi\big)^{\ell} \leq \zeta_k\big\} & \text{otherwise.} \end{cases}$$

In words, the quantity ℓ_k is either the smallest integer such that $\big(1+\xi\big)^{\ell}$ is bigger than ζ_k (if $\zeta_k \leq 1$) or the largest integer such that $\big(1+\xi\big)^{\ell}$ is smaller than ζ_k (if $\zeta_k \geq 1$).

With this definition, we see that $\big(1+\xi\big)^{\ell_k}$ always lies between ζ_k and 1, so that we are guaranteed that $g\big(\big(1+\xi\big)^{\ell_k}\big) \geq g\big(\zeta_k\big)$, and hence

$$\sum_{k=1}^{\infty} g\big(\zeta_k\big) \leq \sum_{k=1}^{\infty} g\big((1+\xi)^{\ell_k}\big).$$

Thus, it suffices to show that at most two distinct values of k map to a single ℓ_k. Indeed, when this mapping condition holds, we have

$$\sum_{k=1}^{\infty} g\big(\zeta_k\big) \leq 2\sum_{\ell=-\infty}^{\infty} g\big(\big(1+\xi\big)^{\ell}\big) \leq 4\sum_{\ell=0}^{\infty} g\big(\big(1+\xi\big)^{\ell}\big).$$

In order to prove the stated mapping condition, note first that ℓ_k is clearly nondecreasing in k, so that we need to prove that $\ell_{k+2} > \ell_k$ for all $k \geq 1$. It is sufficent to show that $\zeta_{k+2} \geq \big(1+\xi\big)\zeta_k$, since this inequality implies that $\ell_{k+2} \geq \ell_k + 1$.

We now exploit the fact that $\zeta_k = an_k$ for some absolute constant a, where $n_k = \lfloor n_0\big(1+\xi\big)^k \rfloor$. For this, let $b = n_0\big(1+\xi\big)^k$, so that $n_k = \lfloor b \rfloor$. Since $n_{k+1} > n_k$, we have $\big(1+\xi\big)b \geq \lfloor \big(1+\xi\big)b \rfloor \geq \lfloor b \rfloor + 1$, and hence

$$\begin{aligned} \frac{n_{k+2}}{n_k} \;=\; \frac{\lfloor \big(1+\xi\big)^2 b \rfloor}{\lfloor b \rfloor} &\geq \frac{\big(1+\xi\big)^2 b - 1}{\lfloor b \rfloor} \\ &\geq \frac{\big(1+\xi\big)\big[\lfloor b \rfloor + 1\big] - 1}{\lfloor b \rfloor} \\ &\geq 1 + \xi, \end{aligned}$$

as required.[3]

Proof of Lemma 7: When $c = 0$ and $\ell = 0$, we note that the claim obviously holds with equality. On the other hand, the left hand side is increasing in ℓ, so that the $c = 0$ case follows immediately.

Turning to the case $c > 0$, we first note that it is equivalent to show that

$$(1+\xi)^{2\ell} - 2(c+1)(1+\xi)^{\ell} + 1 \geq 0 \qquad \text{for all } \ell \in (\tfrac{9c}{5\xi}, \tfrac{9(c+1)}{5\xi}).$$

It suffices to show that $(1+\xi)^{\ell}$ is at least as large as the largest root of the the quadratic equation $z^2 - 2\big(c+1\big)z + 1 = 0$. This largest root is given by

$$z^* = c + 1 + \sqrt{c\,(c+2)} \leq 2(c+1).$$

Consequently, it is enough to show that $\ell \geq \frac{\log 2(c+1)}{\log(1+\xi)}$. Since $\frac{9c}{5\xi}$ is a lower bound on ℓ, we need to verify that

$$\frac{9c}{5\xi} \geq \frac{\log 2\,(c+1)}{\log(1+\xi)}.$$

In order to verify this claim, note first that since $\xi \leq \frac{2}{5}$, we have $\log(1+\xi) \geq \xi - \frac{1}{2}\xi^2 \geq \frac{4}{5}\xi$, whence

$$\frac{\log 2\,(c+1)}{\log(1+\xi)} \leq \frac{5 \log 2\,(c+1)}{4\xi}.$$

Differentiating the upper bound in c, we find that its derivative is

$$\frac{5}{4\,(c+1)\,\xi} \leq \frac{5}{8\xi} \leq \frac{9}{5\xi},$$

so it actually suffices to verify the claim for $c = 1$, which can be done by checking numerically that $\frac{5 \log 4}{4} \leq \frac{9}{5}$.

[3]We thank Daniel Paulin for suggesting this argument as an elaboration on the shorter proof in Gyori and Paulin [46].

Proof of Lemma 8: Our strategy is to split the infinite sum into two parts: one corresponding to the range of s where h is constant and equal to 1 and the other to the range of s where h is decreasing. In terms of the N_k, these two parts are obtained by splitting the sum into terms with $k < k_0$ and $k \geq k_0$, where $k_0 \geq 1$ is minimal such that $M \leq \Delta N_k$ for $k \geq k_0$.

For convenience in what follows, let us introduce the convenient shorthand

$$T_k := \exp\big(-\frac{\big(\Delta N_k - M\big)_+^2}{2\tau_f\big(\delta/2\big)N_k}\big).$$

Now, if $k_0 = 1$, we note that h must then be decreasing for $s \geq N_1$, so that

$$\sum_{k=1}^{\infty}\big(N_{k+1} - N_k\big)T_k \leq \int_{N_1}^{\infty} h\big(s\big)\mathrm{d}s.$$

Otherwise, if $k_0 > 1$, we have

$$\sum_{k=k_0}^{\infty}\big(N_{k+1} - N_k\big)T_k \leq \int_{N_{k_0}}^{\infty} h\big(s\big)\mathrm{d}s.$$

For $k < k_0$, we have $T_k = 1$, so that when $k < k_0 - 1$, we have

$$\big(N_{k+1} - N_k\big)\exp\big(-\frac{\big(\Delta N_k - M\big)_+^2}{2\tau_f\big(\delta/2\big)N_k}\big) = \int_{N_k}^{N_{k+1}} h\big(s\big)\ \mathrm{d}s.$$

Thus

$$\sum_{k=1}^{k_0-1}\big(N_{k+1} - N_k\big)\exp\big(-\frac{\big(\Delta N_k - M\big)_+^2}{2\tau_f\big(\delta/2\big)N_k}\big) = \int_{N_1}^{N_{k_0-1}} h\big(s\big)\mathrm{d}s.$$

Note that this implies

$$\int_{N_1}^{\infty}\exp\big(-\frac{\big(\Delta s - M\big)_+^2}{2\tau_f\big(\delta/2\big)s}\big)\ \mathrm{d}s \geq N_{k_0-1}.$$

Finally, we observe that $N_{k+1} \leq \big(1+\xi\big)N_k + 1 + \xi$, so that $N_{k_0} - N_{k_0-1} \leq \xi N_{k_0-1} + 1 + \xi$. Putting together the pieces, we have

$$\big(N_{k_0} - N_{k_0-1}\big)\exp\big(-\frac{\big(\Delta N_{k_0-1} - M\big)_+^2}{2\tau_f\big(\delta/2\big)N_{k_0-1}}\big) \leq 1 + \xi + \xi\int_{N_1}^{\infty}\exp\big(-\frac{\big(\Delta s - M\big)_+^2}{2\tau_f\big(\delta/2\big)s}\big)\ \mathrm{d}s,$$

and hence

$$\sum_{k=1}^{\infty}\big(N_{k+1} - N_k\big)\exp\big(-\frac{\big(\Delta N_k - M\big)_+^2}{2\tau_f\big(\delta/2\big)N_k}\big) \leq 1 + \xi + \big(1+\xi\big)\int_{N_1}^{\infty} h\big(s\big)\mathrm{d}s.$$

Proof of Lemma 5: Observe that for $k > k_0^*$, we have $\Delta - \epsilon_k \geq \frac{\Delta}{2}$. It follows that for $k > k_0^*$, we have $T_{f,k}^+ \leq T_f\big(\frac{\Delta}{4}\big)$. Thus, we can bound each term in the sum by

$$\big(N_{k+1} - N_k\big) \exp\big(-\frac{N_k}{8T_{f,k}^+} \cdot \big(\Delta - \epsilon_k\big)_+^2\big) \leq \underbrace{\big(N_{k+1} - N_k\big) \exp\big(-\frac{N_k}{T_f\big(\frac{\Delta}{4}\big)} \cdot \frac{\Delta^2}{32}\big)}_{F_k}.$$

Furthermore, the exponential in the definition of F_k is a decreasing function of N_k, so we further bound the overall sum as

$$\begin{aligned}
\sum_{k=k_0^*+1}^{\infty} F_k &\leq \sum_{n=N_0^*+1}^{\infty} \exp\big(-n \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big) \\
&= \exp\big(-N_0^* \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big) \times \sum_{m=1}^{\infty} \exp\big(-m \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big) \\
&= \exp\big(-\frac{N_0^*}{8T_f\big(\frac{(\Delta/2)}{2}\big)} \cdot \big(\frac{\Delta}{2}\big)^2\big) \times \sum_{m=1}^{\infty} \exp\big(-m \cdot \frac{\Delta^2}{32T_f\big(\frac{\Delta}{4}\big)}\big).
\end{aligned}$$

On the other hand, by the definition of N_0^*, $\epsilon_{k_0^*} \leq \frac{\Delta}{2}$, so

$$T_f\big(\frac{(\Delta/2)}{2}\big) \leq T_f\big(\frac{\epsilon_{k_0^*}}{2}\big).$$

By the definition of ϵ_k, however, we know that

$$\frac{\epsilon_k^2}{8T_f\big(\frac{\epsilon_k}{2}\big)} \geq \frac{\log\big(1/\alpha\big) + 1 + 2\log k}{N_k} \geq \frac{\log\big(1/\alpha\big)}{N_k},$$

which implies that $(\Delta/2)^2 N_0^* \geq \log\big(1/\alpha\big)\ 8T_f\big(\frac{(\Delta/2)}{2}\big)$. Re-arranging yields the claim.

Chapter 3

Optimal Error Tradeoffs for Gaussian-Like Multiple Testing

3.1 Overview

The problem of multiple comparisons has been a central topic in statistics ever since Tukey's influential book [91]. In broad terms, suppose we observe a sequence of n independent random variables, $X_1, \ldots, X_n$, of which some unknown subset are drawn from a null distribution, corresponding to the absence of a signal or effect. The remainder are drawn from a non-null distribution, corresponding to signals or effects. Within this framework, we can pose three problems of increasing hardness: the *detection* problem, testing whether or not there is at least one signal; the *localization* problem, identifying the positions of the nulls and signals; and the *estimation* problem, which returns estimates of the means and/or distributions of the observations. Note that these problems form a hierarchy of difficulty: identifying the signals implies that we know whether at least one exists, and estimating each mean implies we know which are zero and which are not. This chapter focuses on the problem of localization.

There are a variety of ways of measuring type-I errors for the localization problem. For example, the *family-wise error rate* measures the probability of incorrectly rejecting at least one null, and the *false discovery rate* (FDR) is the expected ratio of incorrect rejections to total rejections. An extensive body of literature has developed around both metrics, resulting in algorithms geared towards controlling one or the other. Our focus is the FDR metric. Although the FDR has been widely studied, relatively little is known about the behavior of existing algorithms in terms of the corresponding Type-II error concept, namely the *false nondiscovery rate* (FNR).[1] Indeed, it is only recently that Arias-Castro and Chen [4], working with a version of the sparse generalized Gaussian sequence model, established asymptotic consistency for the FDR-FNR localization problem. Informally, in this framework, we receive n independent observations, $X_1, \ldots, X_n$, of which n^{1-s_n} are non-

[1] We follow Arias-Castro and Chen [4] in defining the FNR as the ratio of undiscovered to total non-nulls, which differs from the definition of Genovese and Wasserman [38].

nulls. The remainder are nulls. The $n - n^{1-s_n}$ null variables are drawn from a centered distribution with tails decaying as $\exp\big(-\frac{|x|^\gamma}{\gamma}\big)$, whereas the non-nulls are drawn from the same distribution, shifted by $(\gamma r_n \log n)^{1/\gamma}$. Using this notation, Arias-Castro and Chen [4] considered the setting with fixed problem parameters $r_n = r$ and $s_n = s$, and showed that when $r < s < 1$, all procedures must have risk FDR+FNR $\to 1$. They also showed that in the achievable regime $r > s > 0$, the Benjamini–Hochberg (BH) procedure is consistent, meaning that FDR + FNR $\to 0$. Finally, they proposed a new "distribution-free" method inspired by the knockoff procedure of Barber and Candès (2015), and showed that the resulting procedure is also consistent in the achievable regime.

These existing consistency results are asymptotic. However, to date, no studies have examined the important nonasymptotic questions that are of interest in comparing procedures. For instance, for a given FDR level, what is the best achievable FNR? What is the best nonasymptotic behavior of the risk FDR + FNR attainable in finite samples? In addition, and perhaps most importantly, nonasymptotic questions on whether procedures such as BC and BH are *rate-optimal* for the FDR+FNR risk, remain unanswered. The main contributions of this work are to develop techniques to address such questions, and then to use these techniques to solve the problems in the context of the sparse generalized Gaussians model.

Specifically, we establish the trade-off between the FDR and FNR in finite samples (and, hence, also asymptotically), and we use the trade-off to determine the best attainable rate for the FDR + FNR risk. Our theory is sufficiently general to accommodate sequences of parameters (r_n, s_n), and enabling it to reveal new phenomena that arise when $r_n - s_n = o(1)$. For a fixed pair of parameters (r, s) in the achievable regime $r > s$, our theory leads to an explicit expression for the optimal rate at which FDR+FNR can decay. In particular, defining the γ-"distance" $D_\gamma(a, b) := \big|a^{1/\gamma} - b^{1/\gamma}\big|^\gamma$ between pairs of positive numbers, we show that the equation

$$\kappa = D_\gamma(s + \kappa, r)$$

has a unique solution κ_*. Moreover, the combined risk of any threshold-based multiple testing procedure $\mathcal{I}$ is lower bounded as $\mathcal{R}_n(\mathcal{I}) \gtrsim n^{-\kappa_*}$. Furthermore, using a direct analysis, we prove that the Benjamini–Hochberg (BH) and the Barber–Candès (BC) algorithms both attain this optimal rate.

At the core of our analysis is a simple comparison principle. The flexibility of the resulting proof strategy allows us to identify a new critical regime in which $r_n - s_n = o(1)$. However, in this regime, the problem is infeasible, which means that if the FDR is driven to zero, then the FNR must remain bounded away from zero. Moreover, we are able to study challenging settings in which the fraction of signals is a constant $\pi_1 \in (0, 1)$ and not asymptotically vanishing. This corresponds to the setting $s_n = \frac{\log(1/\pi_1)}{\log n}$, such that $s_n \to 0$. Perhaps surprisingly, even in these regimes, the BH and BC algorithms continue to be optimal, although the best rate can weaken from polynomial to subpolynomial in the number of hypotheses n.

Related work

As described above, our work provides a nonasymptotic generalization of the recent work by Arias-Castro and Chen [4] on asymptotic consistency in localization, using FDR + FNR as the notion of risk. Note that this notion of risk is distinct from the asymptotic Bayes optimality under sparsity (ABOS) studied by Bogdan et al. [13] for Gaussian sequences, and more recently, by Neuvial and Roquain [71] for binary classification with extreme class imbalance. The ABOS results concern a risk derived from the probability of incorrectly rejecting a single null sample (false positive; FP) and the probability of incorrectly failing to reject a single non-null sample (false negative; FN). Specifically, we have $\mathcal{R}_n^{\text{ABOS}} = w_1 \cdot \text{FP} + w_2 \cdot \text{FN}$ for some pair of positive weights (w_1, w_2), which need not be equal. Because this risk is based on the error probability for a single sample, it is much closer to a misclassification risk or a single-testing risk than it is to the ratio-based FDR + FNR risk examined here.

Using our notation, the work of Neuvial and Roquain [71] can be understood as focusing on the setting $r = s$, a regime the authors refer to as the "verge of detectability." Furthermore, their performance metric is given by the Bayes classification risk, rather than the combination of FDR and FNR studied here. In comparison, our results provide additional insight into models that are close to the verge of detectability. Even when $s_n = s$ is fixed, we can provide quantitative lower and upper bounds on the FDR/FNR ratio as $r_n \to s$ from above. Moreover, these bounds depend on how quickly r_n approaches s. A further transition in rates occurs when $r = s$ exactly, for all n; however, we do not explore this case in depth. We suspect that our methods may offer sufficient precision to answer the nonasymptotic minimaxity questions posed by Neuvial and Roquain [71] on whether any threshold-based procedure can match the Bayes optimal classification error rate, up to an additive error $\ll \frac{1}{\log n}$.

For the special case of $\gamma = 2$, Ji and Jin [56] and Ji and Zhao [57] prove bounds for localization that are closely related to, but distinct from, our bounds on the overall risk. Both deal with a sparse high-dimensional regression. The former work proposes a new method for variable selection, called UPS, that has advantages over the lasso and subset selection methods in certain settings. The latter builds on the first to prove upper and lower bounds for multiple testing, using the so-called mFNR and mFDR. These metrics replace the expected ratio in the definitions of FDR and FNR (see definition (3.3) below) with a ratio of expectations—a modification that should lead to qualitatively similar behavior as n becomes large. The resulting bounds in both works can be used to recover our bounds up to polylogarithmic factors in the special case where $\gamma = 2$. The main advantage of their work, relative to ours, lies in how they handle the dependence between the p-values. Unlike our work, however, they do not establish the trade-off between the FDR and FNR when both quantities can decay to zero at different rates; in addition, as mentioned, they only consider the case of $\gamma = 2$. Nor do they consider regimes where the sparsity and signal strength vary with n. Our results can handle this more general setting, which encompasses dense regimes with qualitatively different behavior from the more commonly investigated sparse one.

The above line of work is complementary to the well-known asymptotic results of Donoho

and Jin (2004, 2015) on phase transitions in detectability using Tukey's higher-criticism statistic, which employs standard type-I and type-II errors for testing of the single global null hypothesis. Note that Donoho and Jin use the generalized Gaussian assumption directly on the PDFs, whereas our assumption (3.5) is on the survival function. Just as in Arias-Castro and Chen [4], Donoho and Jin consider the asymptotic setting with $r_n = r$ and $s_n = s$, which they sometimes call the RW (rare and weak) model. We are not aware of any nonasymptotic results for detection that are similar to those proposed here for localization.

This chapter also complements work on estimation, the most notable result being the asymptotic minimax optimality of BH-derivedI thresholding for denoising an approximately sparse high-dimensional vector [1, 28]. The relevance of our results to the minimaxity of BH for approximately sparse denoising problems lies primarily in the use of deterministic thresholds as a useful proxy for BH, as well as other procedures that determine their threshold in a manner that has complex dependence on the input data [28]. Unlike the strategy of Donoho and Jin [28], which depends on establishing the concentration of the empirical threshold around the population-level value, we use a more flexible comparison principle. Deterministic approximations to optimal FDR thresholds are also studied by Chi [18] and Genovese et al. [39]. Other related papers are discussed in Section 3.5, when discussing directions for future work.

The remainder of this paper is organized as follows. In Section 3.2, we provide some background on the multiple testing problem, as well as the particular model we consider. In Section 3.3, we provide an overview of our main results: the optimal trade-offs between the FDR and FNR, which imply lower bounds on the FDR + FNR risk, and optimality guarantees for the BH and BC algorithms. In Section 3.4, we prove our main results. We first focus on the lower bounds, and then provide matching upper bounds for the well-known and popular BH and BC algorithms for multiple testing with FDR control. The proofs of some technical lemmas are given in Sections 3.6-3.6.

3.2 Problem formulation

In this section, we provide background and a precise formulation of the problem under study.

Multiple testing and the FDR

Suppose that we observe a real-valued sequence $X_1^n := \{X_1, \dots, X_n\}$ of n independent random variables. When the null hypothesis is true, X_i is assumed to have a zero mean; otherwise, it is assumed that the mean of X_i is equal to some unknown number $\mu_n > 0$. The binary labels $\{H_1, \dots, H_n\}$ indicate whether the null hypothesis holds for each observation; the setting $H_i = 0$ indicates that the null hypothesis holds. We define

$$\mathcal{H}_0 := \{i \in [n] \mid H_i = 0\}, \quad \text{and} \quad \mathcal{H}_1 := \{i \in [n] \mid H_i = 1\}, \tag{3.1}$$

corresponding to the *nulls* and *signals*, respectively. Our task is to identify a subset of indices that contains as many signals as possible, while not containing too many nulls.

More formally, a testing rule $\mathcal{I} : \mathbb{R}^n \to 2^{[n]}$ is a measurable mapping of the observation sequence X_1^n to a set $\mathcal{I}(X_1^n) \subseteq [n]$ of *discoveries*, where the subset $\mathcal{I}(X_1^n)$ contains those indices for which the procedure rejects the null hypothesis. There is no single unique measure of performance for a testing rule for the localization problem. We employ the FDR and FNR for this purpose. These can be viewed as generalizations of the type-I and type-II errors for single hypothesis testing.

We begin by defining the false discovery proportion (FDP) and the false nondiscovery proportion (FNP), respectively, as

$$\mathrm{FDP}_n(\mathcal{I}) := \frac{\mathrm{card}(\mathcal{I}(X_1^n) \cap \mathcal{H}_0)}{\mathrm{card}(\mathcal{I}(X_1^n)) \vee 1}, \quad \text{and} \quad \mathrm{FNP}_n(\mathcal{I}) := \frac{\mathrm{card}(\mathcal{I}(X_1^n)^c \cap \mathcal{H}_1)}{\mathrm{card}(\mathcal{H}_1)}. \tag{3.2}$$

Because the output $\mathcal{I}(X_1^n)$ of the testing procedure is random, both quantities are random variables. The FDR and FNR are given by taking the expectations of these random quantities; that is,

$$\mathrm{FDR}_n(\mathcal{I}) := \mathbb{E}\Big[\frac{\mathrm{card}(\mathcal{I}(X_1^n) \cap \mathcal{H}_0)}{\mathrm{card}(\mathcal{I}(X_1^n)) \vee 1}\Big], \quad \text{and} \quad \mathrm{FNR}_n(\mathcal{I}) := \mathbb{E}\Big[\frac{\mathrm{card}(\mathcal{I}(X_1^n)^c \cap \mathcal{H}_1)}{\mathrm{card}(\mathcal{H}_1)}\Big], \tag{3.3}$$

where the expectation is taken over the random samples X_1^n.

Note that our definitions of the FNP and FNR, which follow those of Arias-Castro and Chen [4], differ from an alternative definition of the FNR_{alt}, where the denominator is set to the number of nonrejections. In general, however, the number of nonrejections will be close to n for any procedure with low FDR. Thus, in the sparse regime, the FNR_{alt} would trivially go to zero for any procedure that controls the FDR at any level strictly below one. Our definition is therefore better suited to studying transitions in difficulty in the multiple testing problem.

We measure the overall performance of a procedure in terms of its *combined risk*,

$$\mathcal{R}_n(\mathcal{I}) := \mathrm{FDR}_n(\mathcal{I}) + \mathrm{FNR}_n(\mathcal{I}). \tag{3.4}$$

Finally, when the testing rule $\mathcal{I}$ is clear from the context, we frequently omit an explicit reference to the dependence on the testing rule.

Tail generalized Gaussian model

In this chapter, we describe the distribution of the observations for both nulls and non-nulls in terms of a *tail generalized Gaussian model*. Our model is a variant of the generalized Gaussian sequence model, studied in Arias-Castro and Chen [4] and Donoho and Jin [27]; the only difference is that whereas a γ-generalized Gaussian has a density proportional to $\exp\big(-\frac{|x|^\gamma}{\gamma}\big)$, we focus on distributions with tails proportional to $\exp\big(-\frac{|x|^\gamma}{\gamma}\big)$. This alteration

is in line with the asymptotically generalized Gaussian (AGG) distributions studied by Arias-Castro and Chen [4], with the important caveat that our assumptions are imposed in a nonasymptotic fashion.

For a given degree $\gamma \geq 1$, a γ-tail generalized Gaussian random variable with mean zero, written as $G \sim \mathrm{tGG}_\gamma(0)$, has a survival function $\Psi(t) := \mathbb{P}\big(G \geq t\big)$ that satisfies the bounds

$$\frac{e^{\frac{-|t|^\gamma}{\gamma}}}{Z_\ell} \;\leq\; \min\{\Psi(t), 1-\Psi(t)\} \;\leq\; \frac{e^{\frac{-|t|^\gamma}{\gamma}}}{Z_u}, \quad t \in \mathbb{R}, \tag{3.5}$$

for some constants $Z_\ell > Z_u > 0$. (Note that $t \mapsto \Psi(t)$ is a decreasing function, and becomes smaller than $1 - \Psi(t)$ at the origin.) As a concrete example, a γ-tail generalized Gaussian with $Z_\ell = Z_u = 1$ can be generated by sampling a standard exponential random variable E and a Rademacher random variable ε, and then letting $G = \varepsilon\big(\gamma E\big)^{1/\gamma}$. We use the terminology "tail generalized Gaussian" because the survival function of a two-tail Gaussian random variable is of the order of $\exp(-|x|^2/2)$, whereas that of a Gaussian is of the order of $\frac{1}{\mathrm{poly}(x)}\exp(-x^2/2)$. In particular, this observation implies that a tGG_2 random variable has tails that are equivalent to those of a Gaussian in terms of their exponential decay rates.

In terms of this notation, we assume that each observation X_i is distributed as

$$X_i \sim \begin{cases} \mathrm{tGG}_\gamma(0) & \text{if } i \in \mathcal{H}_0 \\ \mathrm{tGG}_\gamma(0) + \mu_n & \text{if } i \in \mathcal{H}_1, \end{cases} \tag{3.6}$$

where our notation reflects the fact that the mean shift μ_n is permitted to vary with the number of observations n. See Section 3.3 for further discussion of the scaling of the mean shift.

Threshold-based procedures

Following prior work [4, 27], we restrict our attention to testing procedures of the form

$$\mathcal{I}(X_1^n) = \big\{i \in [n] \mid X_i \geq T_n(X_1^n)\big\}, \tag{3.7}$$

where $T_n(X_1^n) \in \mathbb{R}_+$ is a data-dependent threshold. We refer to such methods as *threshold-based procedures*. The BH and BC procedures both belong to this class. Moreover, from an intuitive standpoint, the observations are exchangeable in the absence of prior information, and we are testing between a single unimodal null distribution and a single positive shift of that distribution. In this setting, it is difficult to conceive of reasonable procedures that would reject the hypothesis corresponding to one observation, while rejecting a hypothesis with a smaller observation value.

In particular, as part of our argument, it will be important to analyze the performance metrics associated with rules of the form

$$\mathcal{I}_t(X_1^n) = \big\{i \in [n] \mid X_i \geq t\big\}, \tag{3.8}$$

where $t > 0$ is a prespecified (fixed, nonrandom) threshold. In this case, we adopt the notation $\mathrm{FDR}_n(t)$, $\mathrm{FNR}_n(t)$, and $\mathcal{R}_n(t)$ to denote the metrics associated with the rule $X_1^n \mapsto \mathcal{I}_t(X_1^n)$.

The BH and BC procedures

Arguably the most popular threshold-based procedure that provably controls FDR at a user-specified level α_n is the BH procedure. More recently, Arias-Castro and Chen [4] proposed a method that we refer to as the BC procedure. Both algorithms are based on estimating the FDP_n that would be incurred at a range of possible thresholds, and then choosing the largest one possible (maximizing discoveries), while satisfying an upper bound linked to α_n (controlling FDR_n). Furthermore, they both only consider thresholds that coincide with one of the values X_1^n, which we denote as the set $\mathcal{X}_n = \{X_1, \dots, X_n\}$. The data-dependent threshold for both can be written as

$$t_n(X_1, \dots, X_n) = \min\big\{t \in \mathcal{X}_n \colon \widehat{\mathrm{FDP}}_n(t) \le \alpha_n\big\}. \tag{3.9}$$

The two algorithms differ in the estimator $\widehat{\mathrm{FDP}}_n(t)$ they use. The BH procedure assumes access to the true null distribution through its survival function Ψ and sets

$$\widehat{\mathrm{FDP}}_n^{\mathrm{BH}}(t) = \frac{\Psi(t)}{\#(X_i \ge t)/n}, \quad \text{for } t \in \mathcal{X}_n. \tag{3.10}$$

The BC procedure instead estimates the survival function $\Psi(t)$ from the data and, therefore, does not need to know the null distribution. This approach is viable when $\#(X_i \le -t)/n$ is a good proxy for $\Psi(t)$, which our upper and lower tail bounds guarantee; more typically, the BC procedure is applicable when the null distribution is (nearly) symmetric, and the signals are shifted by a positive amount (as they are in our case). Then, the BC estimator is given by

$$\widehat{\mathrm{FDP}}_n^{\mathrm{BC}}(t) = \frac{\big[\#(X_i \le -t) + 1\big]/n}{\#(X_i \ge t)/n}, \quad \text{for } t \in \mathcal{X}_n. \tag{3.11}$$

With these definitions in place, we are now ready to describe our main results.

3.3 Main results

We now state our main results and examine their consequences. Our first main result (Theorem 5) characterizes the optimal trade-off between the FDR and FNR for any testing procedure. By optimizing this trade-off, we obtain a lower bound on the combined FDR and FNR of any testing procedure (Corollary 4). Our second main result (Theorem 6), shows that the BH procedure achieves the optimal FDR-FNR trade-off up to constants, and that the BC procedure almost achieves optimality. In particular, our result implies that, with the proper choice of target FDR, the BH and BC procedures can both achieve the optimal combined FDR-FNR rate (Corollary 5).

Scaling of sparsity and mean shifts

We study a sparse instance of the multiple testing problem, in which the number of signals is assumed to be small relative to the total number of hypotheses. In particular, motivated by related works on multiple hypothesis testing [4, 27, 29, 58], we assume that the number of signals scales as

$$\text{card}(\mathcal{H}_1) = m_n = n^{1-s_n} \quad \text{for some } s_n \in (0,1). \tag{3.12}$$

Note that, to the best of our knowledge, all previous results in the literature assume that $s_n = s$ is actually independent of n. In this case, the sparsity assumption (3.12) implies that all but a polynomially vanishing fraction of the hypotheses are null. In contrast, as indicated by our choice of notation, our setup allows for a sequence of parameters β_n that can vary with the number of hypotheses n. As a result, our framework is flexible enough to handle relatively dense regimes (e.g., those with $\frac{n}{\log n}$ or even $\mathcal{O}(n)$ signals).

The non-null hypotheses are distinguished by a positively shifted mean $\mu_n > 0$. It is natural to parameterize this mean shift in terms of a quantity $r_n > 0$ via the relation

$$\mu_n = \big(\gamma r_n \log n\big)^{1/\gamma}. \tag{3.13}$$

As shown by Arias-Castro and Chen [4], when the pair (s, r) are fixed such that $r < s$, the problem is asymptotically infeasible, meaning that there is no procedure such that $\mathcal{R}_n(\mathcal{I}) \to 0$ as $n \to \infty$. Accordingly, we focus on sequences (s_n, r_n), for which $r_n > s_n$. Furthermore, even though the asymptotic consistency boundary of $r < s$ versus $r > s$ is apparently independent of γ, we find that the rate at which the risk decays to zero is determined jointly by r, s, and γ.

Lower bound on any threshold-based procedure

In this section, we assume:

$$s_n \overset{(i)}{\geq} \frac{\log 2}{\log n} \quad \Longleftrightarrow \quad n^{1-s_n} \leq n/2, \quad \text{and} \tag{3.14a}$$

$$\max\{\beta_n, \frac{1}{\log^{\frac{\gamma-1/2}{\gamma}} n}\} \overset{(ii)}{<} r_n \overset{(iii)}{<} r_{\max} \quad \text{for some constant } r_{\max} < 1. \tag{3.14b}$$

Condition (i) requires that the proportion π_1 of non-nulls is at most $1/2$. Condition (ii) asserts that the natural requirement of $r_n > s_n$ is not sufficient, but further insists that r_n cannot approach zero too fast. The constants $\log 2$ and $\frac{\gamma-1/2}{\gamma}$ are somewhat arbitrary and can be replaced, respectively, by $\log \frac{1}{\pi_{\max}}$ for any $0 < \pi_{\max} < 1$ and $\frac{\gamma-1+\rho}{\gamma}$ for any $\rho > 0$. However, we fix their values in order not to introduce unnecessary extra parameters. With regard to condition (iii), although the assumption $r_n < 1$ is imposed because the problem becomes qualitatively easy for $r_n \geq 1$, the assumption that it is bounded away from one is a technical convenience that simplifies some of our proofs.

Our analysis shows that the FNR behaves differently depending on the closeness of the parameter r_n to the boundary of feasiblity given by s_n. In order to characterize this closeness, we define

$$r_{\min} = r_{\min}(\kappa_n) := \begin{cases} s_n + \kappa_n + \frac{\log \frac{1}{6Z_\ell}}{\log n} & \text{if } \kappa_n \leq 1 - s_n - \frac{\log \frac{3}{\log 16}}{\log n}, \\ 1 + \frac{\log \frac{1}{24Z_\ell}}{\log n} & \text{otherwise.} \end{cases} \tag{3.15}$$

Here ,κ_n is interpreted as the "exponent" of a target FDR rate α_n, in the sense that $\alpha_n = n^{-\kappa_n}$. The rate α_n may differ from the actual achieved FDR_n, but it is nonetheless useful for parameterizing the quantities in our analysis. When we need to move between α_n and κ_n, we shall write $\kappa_n = \kappa_n(\alpha_n) = \frac{\log(1/\alpha_n)}{\log n}$ and $\alpha_n = \alpha_n(\kappa_n) = n^{-\kappa_n}$. For mathematical convenience, we wish to have the target FDR α_n be bounded away from one; therefore, we impose one further technical, but inessential assumption in this section:

$$\alpha_n \leq \min\left\{\frac{1}{24},\ \frac{1}{6Z_\ell}\right\} \quad \Longleftrightarrow \quad \kappa_n \geq \frac{\log \max\left\{24,\ 6Z_\ell\right\}}{\log n}. \tag{3.16}$$

The theorem that follows applies to all sample sizes $n > n_{\min,\ell}$ (the subscript ℓ denotes lower), where

$$n_{\min,\ell} := \min\left\{n \in \mathbb{N} : \exp\left(-\frac{n^{1-r_{\max}}}{24(Z_\ell \vee 1)}\right) \leq \frac{1}{4}\right\} \tag{3.17}$$

$$= \left\lfloor [24(Z_\ell \vee 1)\log 4]^{\frac{1}{1-r_{\max}}} \right\rfloor, \tag{3.18}$$

which is an explicit known function of the problem parameters, and can therefore be computed whenever the problem setting is fixed.

Finally, for $\gamma \in [1, \infty)$ and nonnegative numbers $a, b > 0$, we define the associated γ-"distance" as follows:

$$D_\gamma(a, b) := \left|a^{1/\gamma} - b^{1/\gamma}\right|^\gamma. \tag{3.19}$$

Our first main theorem states that for $r_n > r_{\min}(\kappa_n)$, the FNR decays as a power of $1/n$, with the exponent specified by the γ-distance.

Theorem 5. *Consider the γ-tail generalized Gaussian testing problem with sparsity β_n and signal level r_n, satisfying conditions* (3.14a), *and* (3.14b) *and with sample size $n > n_{\min,\ell}$, from definition* (3.18). *Then, for any choice of exponent $\kappa_n \in (0, 1)$ satisfying condition* (3.16), *there exists a minimum signal strength $r_{\min}(\kappa_n)$ from definition* (3.15), *such that any threshold-based procedure $\mathcal{I}$ that satisfies $\mathrm{FDR}_n(\mathcal{I}) \leq n^{-\kappa_n}$ must have its* FNR *lower bounded as*

$$\mathrm{FNR}_n(\mathcal{I}) \geq \begin{cases} \frac{1}{32} & \text{if } r_n \in \left[\beta_n, r_{\min}\right] \\ c(\beta_n, \gamma)\ n^{-D_\gamma(\beta_n + \kappa_n, r_n)} & \text{otherwise,} \end{cases} \tag{3.20}$$

where $c(\beta_n, \gamma) := c_0 \exp\left(c_1 \beta_n^{\frac{1-\gamma}{\gamma}}\right)$, with (c_0, c_1) being positive constants depending only on (Z_ℓ, Z_u, γ).

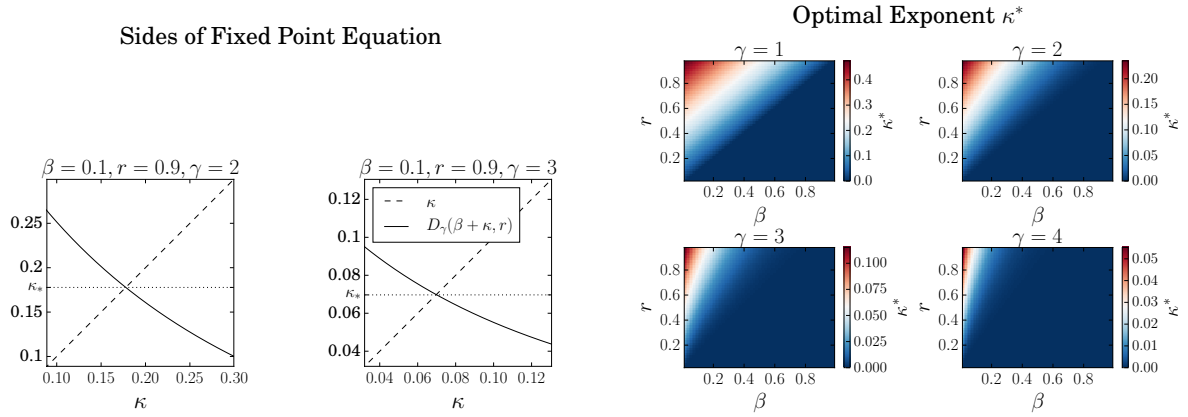

(a) (b)

Figure 3.1: Visualizations of the fixed-point equation (3.21). (a) Plots comparing the left- and right-hand sides of the fixed-point equation. (b) The optimal exponent κ^* as a function of r and s.

The proof of this theorem is provided in Section 3.4. Note that the theorem holds for any choice of $\kappa_n \in (0, 1)$. In the special case of constant pairs (s, r), this choice can be optimized to achieve the best possible lower bound on the risk $\mathcal{R}_n(\mathcal{I}) = \mathrm{FDR}_n(\mathcal{I}) + \mathrm{FNR}_n(\mathcal{I})$. Because we obtain this lower bound by optimizing the sum of the FDR and FNR lower bounds from Theorem 5, we want to balance the contributions from these two bounds. Doing so requires that we set the FDR rate κ equal to the corresponding FNR rate $D_\gamma(s + \kappa, r)$, which leads to a fixed-point equation for the overall rate, as summarized below.

Corollary 4. *When $r > s$, let $\kappa_* = \kappa_*(s, r, \gamma) > 0$ be the unique solution to the equation*

$$\kappa = D_\gamma(s + \kappa, r). \tag{3.21}$$

Then, the combined risk of any threshold-based multiple testing procedure $\mathcal{I}$ is lower bounded as

$$\mathcal{R}_n(\mathcal{I}) \gtrsim n^{-\kappa_*}, \tag{3.22}$$

where $\gtrsim$ denotes inequality up to a prefactor independent of n.

The proof of this corollary is provided in the Section 3.6. Figure 3.1 shows the predictions in Corollary 4. In particular, panel (a) shows how the unique solution κ_* to equation (3.21) is determined for varying settings of the triple (r, s, γ). Panel (b) shows how κ_* varies over the interval $(0, 0.5)$, again for different settings of the triple (r, s, γ). As would be expected, the fixed point κ_* increases as a function of the difference $r - s > 0$.

Upper bounds for some specific procedures

Thus far, we have provided general lower bounds that can be applied to any threshold procedure. We now turn to the complementary question—how do these lower bounds compare to the results achievable by the BH and BC algorithms introduced in Section 3.2? Remarkably, we find that up to the constants defining the prefactor, both procedures achieve the minimax lower bound of Theorem 5.

We state these achievable results in terms of the fixed point κ_* from equation (3.21). Moreover, they apply to all problems with sample size $n > n_{\min,u}$ (the subscript u denotes upper), where

$$\begin{aligned} n_{\min,u} &:= \min\left\{n \in \mathbb{N} : \exp\left(-\frac{n^{1-r_{\max}}}{24}\right) \leq \frac{1}{Z_u n}\right\} \\ &= \min\left\{n \in \mathbb{N} : n \geq [24\log(Z_u n)]^{\frac{1}{1-r_{\max}}}\right\}. \end{aligned} \tag{3.23}$$

As in the case of (3.18), this lower bound on n is explicitly computable from the problem parameters.

In order to state our results cleanly, let us introduce the constants

$$c_{\mathrm{BH}} := \frac{Z_u}{36 Z_\ell}, \quad c_{\mathrm{BC}} := \frac{Z_u}{48 Z_\ell}, \quad \text{and} \quad \zeta := \max\left\{6Z_\ell, \frac{1}{6Z_\ell}\right\}, \tag{3.24}$$

and require in particular that $r_n \geq r_{\min}(\kappa_n(c_{\mathrm{A}}\alpha_n))$ for algorithm $\mathrm{A} \in \{\mathrm{BH}, \mathrm{BC}\}$. Note that $c_{\mathrm{A}} < 1$ because $Z_\ell \geq Z_u$, by definition, and that the introduction of c_{A} into the argument of $r_{\min}$ only changes the minimum allowed value of r_n by a conceptually negligible amount of $\mathcal{O}\big(\frac{1}{\log n}\big)$.

Lastly, note that the BC procedure requires an additional mild condition that the number of non-nulls n^{1-s_n} is large relative to the target FDR $\alpha_n = n^{-\kappa_n}$ (otherwise, in some sense, the problem is too hard if there are too few non-nulls and a very strict target FDR). Specifically, we need that both quantities cannot simultaneously be too small, formalized by the assumption:

$$\exists n_{\min,\mathrm{BC}}, \quad \text{such that, for all} \quad n \geq n_{\min,\mathrm{BC}}, \quad \text{we have} \quad \frac{3c_{\mathrm{BC}}}{4} \cdot \frac{q_n}{\log\frac{1}{\alpha_n}} \cdot n^{1-s_n} \geq 1. \tag{3.25}$$

Note that when $r_n = r$ and $s_n = s$ are constants, this decay condition is satisfied by $\alpha_n = n^{-\kappa_*}$.

Our second main theorem delivers an optimality result for the BH and BC procedures, showing that under some regularity conditions, their performance achieves the lower bounds in Theorem 5, up to constant factors.

Theorem 6. *Consider the β_n-sparse γ-tail generalized Gaussian testing problem with target FDR level α_n, upper bounded as in condition* (3.16)*.*

(a) Guarantee for BH procedure: Given a signal strength $r_n \geq r_{\min}(\kappa_n(c_{\mathrm{BH}}\alpha_n))$ *and sample size* $n > n_{\min,u}$*, as in condition* (3.23)*, the BH procedure satisfies the bounds*

$$\mathrm{FDR}_n \leq \alpha_n \quad \textit{and} \quad \mathrm{FNR}_n \leq \frac{2\zeta_{\mathrm{BH}}^{2s_n^{\frac{1-\gamma}{\gamma}}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r_n)}, \qquad \textit{where } \zeta_{\mathrm{BH}} := \tfrac{\zeta}{c_{\mathrm{BH}}}. \tag{3.26}$$

(b) Guarantee for BC procedure: Given a signal strength $r_n \geq r_{\min}(\kappa_n(c_{\mathrm{BC}}\alpha_n))$ *and sample size* $n > \max\{n_{\min,\mathrm{BC}}, n_{\min,u}\}$*, as in condition* (3.25)*, the BC procedure satisfies the bounds*

$$\mathrm{FDR}_n \leq \alpha_n \quad \textit{and} \quad \mathrm{FNR}_n \leq \frac{2\zeta_{\mathrm{BC}}^{2s_n^{\frac{1-\gamma}{\gamma}}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r_n)} + \alpha_n, \qquad \textit{where } \zeta_{\mathrm{BC}} := \tfrac{\zeta}{c_{\mathrm{BC}}}. \tag{3.27}$$

The proof of the theorem can be found in Section 3.4. For constant pairs (r, s), Theorem 6 can be applied with a target FDR proportional to $n^{-\kappa_*}$ to show that the BH and BC procedures both achieve the optimal decay of the combined FDR-FNR up to constant factors, as stated formally below.

Corollary 5. *For* $s < r$ *and* $\alpha_* = c_* n^{-\kappa_*}$*, with* $0 < c_* \leq \min\left\{\frac{1}{24}, \frac{1}{6Z_\ell}\right\}$*, the BH and BC procedures with target FDR* α_* *satisfy*

$$\mathcal{R}_n \lesssim n^{-\kappa_*}. \tag{3.28}$$

The proof of this corollary is given in Section 3.6. To help visualize the result of the corollary, Figure 3.2 displays the results of simulations of the BH procedure that show the correspondence between its performance and the theoretically predicted rate of $n^{-\kappa_*}$.

Despite the optimality, Figures 3.1 and 3.2 suggest that the methods may not be practical. Although asymptotic consistency can be achieved when $r > s$, the convergence of the risk to zero can be extremely slow, exhibiting "nonparametric" rates far slower than $n^{-1/2}$. Figure 3.2 shows in particular that the decay to zero may be barely evident, even for sample sizes as large as $n = 250,000$, and with comparatively strong signals. The "nonparametric" nature may arise because the dimensionality of the decision space increases linearly with the sample size. Thus, asymptotically, the advantage of having a greater amount of data seems to only *just* overcome the disadvantage of having to make an increasing number of decisions. However, nonasymptotically, one cannot hope to drive both the FDR and the FNR to zero at any practical sample size in this general setting, at least when the mean signal lies below the maximum of the nulls (i.e., $r_n < 1$).

Intuition for the γ-distance. The distance D_γ plays a crucial role because of the scaling of order statistics under the tGG_γ model. If $W_1, \ldots, W_n$ are independent and identically

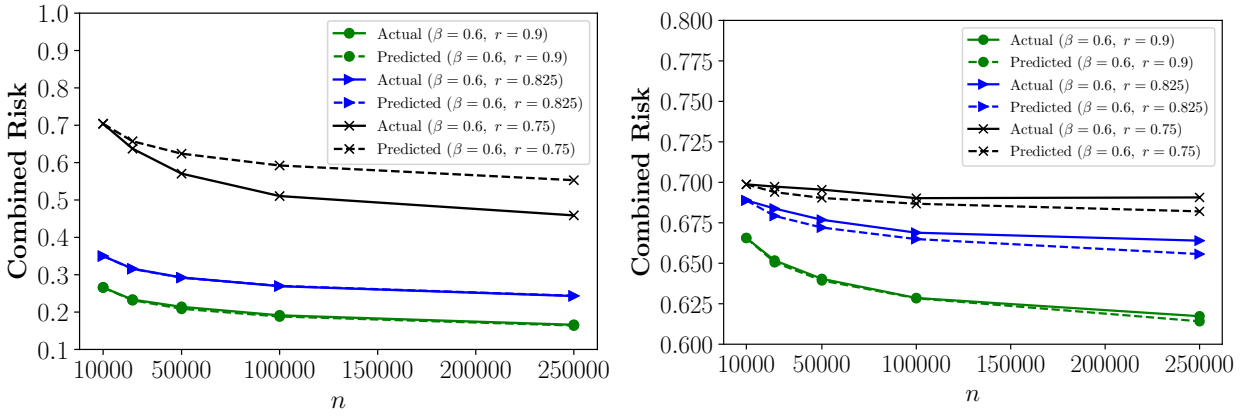

(a) $\gamma = 1$ (b) $\gamma = 2$

Figure 3.2: Results of simulations comparing the predicted combined risk with the actual, experimentally observed risk for the BH procedure. Agreement is good and improves as the gap $(r - s)$ increases, which we believe occurs because the sampling error becomes a smaller fraction of the risk as the separation increases.

distributed (i.i.d.) from a $\mathrm{tGG}_\gamma\big(0\big)$ model, then—ignoring constants inside the logarithm—we expect the ith-largest order statistic $W_{(i)}$ to be around $\big(\gamma i \log n\big)^{1/\gamma}$ if $i \ll n/2$, and around $-\big(\gamma i \log n\big)^{1/\gamma}$ if $n - i \ll n/2$. If an algorithm is to achieve an FNR on the order of $n^{-\kappa'}$, it must successfully identify all but the smallest $n^{-\kappa'}$ fraction of true signals. Thus, the algorithm's cutoff for rejection must exceed the $m - n^{-\kappa'}m$ order statistic of the signals, which is approximately

$$\mu - \big(\gamma \log \frac{m}{n^{-\kappa'}m}\big)^{1/\gamma} = \big(\gamma r \log n\big)^{1/\gamma} - \big(\gamma \kappa' \log n\big)^{1/\gamma}. \tag{3.29}$$

If we suppose that the FDR is also vanishing at a rate $n^{-\kappa}$, then first of all the algorithm must identify about $(1 \pm o(1))m$ indices as signals, because otherwise either the FDR or FNR would fail to vanish. Second, it must be that the $n^{-\kappa}m$th or, equivalently, the $n^{1-s-\kappa}$th largest null is of the order of the quantity in (3.29). Combining these insights, we obtain the relation

$$\big(\gamma\big(s + \kappa\big) = \big(\gamma r \log n\big)^{1/\gamma} - \big(\gamma \kappa' \log n\big)^{1/\gamma},$$

which, after rearranging, yields the heuristic

$$\kappa' = \left(r^{1/\gamma} - \big(s + \kappa\big)^{1/\gamma}\right)^{1/\gamma}. \tag{3.30}$$

The theorems and corollaries in this chapter together show that this intuition is exactly right.

Regime of linear sparsity: We turn to the regime of *linear sparsity*—that is, when the number of signals scales as $\pi_1 n$, for some scalar $\pi_1 \in (0,1)$. Recalling that we have parameterized the number of signals as $n^{1-\beta_n}$, some algebra leads to $s_n = \frac{\log \frac{1}{\pi_1}}{\log n}$; thus, both Theorem 5 and Theorem 6 predict an upper and lower bound on the risk of the form

$$c_0 \exp\left(c_1 \left[\frac{\log n}{\log \frac{1}{\pi_1}}\right]^{\frac{\gamma-1}{\gamma}}\right) \cdot n^{-\kappa_*}. \tag{3.31}$$

Note that here we overload the exponent κ_* to the case when it is nonconstant. In order to interpret this result, observe that if $r_n = r$ is constant, then $\kappa_* = \frac{r}{2^\gamma} - o(1)$; thus, the rate is $n^{-r/2^\gamma}$ up to subpolynomial factors in n. On the other hand, if $r_n = \frac{1}{\log^{\frac{\gamma-1/2}{\gamma}} n}$ is at the extreme lower limit permitted by the lower bound (ii) in (3.14b), then it is not difficult to see that $\kappa_* \approx \log^{-\frac{\gamma-1/2}{\gamma}} n$, which ensures that $n^{\kappa_*} \gg \exp\big(\log^{\frac{\gamma-1}{\gamma}} n\big)$, so that the risk (3.31) still approaches zero asymptotically, albeit subpolynomially, in n.

3.4 Proofs

We now turn to the proofs of our main results, namely Theorems 5 and 6. The proofs of the associated corollaries can be found in Sections 3.6 and 3.6.

Proof of Theorem 5

The main idea of the proof is to reduce the problem of lower bounding the FNR_n of threshold-based procedures that use random, data-dependent thresholds T_n to the easier problem of lower bounding the FNR_n of threshold-based procedures that use a deterministic, data-independent threshold t_n. We refer to the latter class of procedures as *fixed-threshold procedures*, and we parameterize them by their target FDR $\alpha_n = n^{-\kappa_n}$. Specifically, we define the *critical threshold*, derived from the critical regime boundary $r_{\min}$ from equation (3.15), by

$$\tau_{\min}(\kappa_n) := \big(\gamma r_{\min}(\kappa_n) \log n\big)^{1/\gamma} \quad \equiv \quad \tau_{\min}(\alpha_n) := \left(\gamma r_{\min}\left(\frac{\log(1/\alpha_n)}{\log n}\right) \log n\right)^{1/\gamma}. \tag{3.32}$$

Here, and throughout the proof, we express $\tau_{\min}$ and $r_{\min}$ as functions of α_n rather than κ_n; this formulation makes certain calculations in the proof simpler to express.

From data-dependent threshold to fixed threshold: Our first step is to reduce the analysis from data-dependent to fixed-threshold procedures. In particular, consider a threshold procedure, using a possibly random threshold T_n, that satisfies the FDR uppper bound $\mathrm{FDR}_n(T_n) \le \alpha_n$. We claim that the FNR of any such procedure must be lower bounded as

$$\mathbb{E}[\mathrm{FNP}_n(T_n)] \;\ge\; \frac{\mathrm{FNR}_n\big(\tau_{\min}(4\alpha_n)\big)}{16}. \tag{3.33}$$

This lower bound is crucial, because it reduces the study of random threshold procedures (LHS) to the study of fixed-threshold procedures (RHS). Its proof can be found in Section 3.6.

Our next step is to lower bound the FNR for choices of the threshold $t \geq \tau_{\min}(\alpha_n)$:

Lemma 9. *For any $t \geq \tau_{\min}(\alpha_n)$, we have*

$$\mathrm{FNR}_n(t) \geq \begin{cases} \frac{\zeta^{2s_n^{\frac{1-\gamma}{\gamma}}}}{Z_\ell} \cdot n^{-D_\gamma(s_n+\kappa_n,r)} & \text{if } r > r_{\min}\big(\kappa_n(\alpha_n)\big), \\ \frac{1}{2} & \text{otherwise,} \end{cases} \tag{3.34}$$

where ζ is defined as in (3.24).

The proof of this lemma can be found in Section 3.6. Using Lemma 9 and the lower bound (3.33), we can now complete the proof of Theorem 5. We split the argument into two cases:

Case 1: First, suppose that $r \leq r_{\min}(\kappa_n(4\alpha_n))$. In this case, we have

$$\mathrm{FNR}_n(T_n) \overset{(i)}{\geq} \frac{\mathrm{FNR}_n\big(\tau_{\min}\big(4\alpha_n\big)\big)}{16} \overset{(ii)}{\geq} \frac{1}{32},$$

where step (i) follows from the lower bound (3.33), and step (ii) follows by lower bounding the FNR by 1/2, as is guaranteed by Lemma 9 in the regime $r \leq r_{\min}(\kappa_n(4\alpha_n))$.

Case 2: Otherwise, we may assume that $r > r_{\min}(4\alpha_n)$. In this case, we have

$$\mathrm{FNR}_n(T_n) \overset{(i)}{\geq} \frac{\mathrm{FNR}_n\big(\tau_{\min}\big(4\alpha_n\big)\big)}{16} \overset{(ii)}{\geq} \frac{\big(4\zeta\big)^{2s_n^{\frac{1-\gamma}{\gamma}}}}{Z_\ell} \cdot n^{-D_\gamma(s_n+\kappa_n,r)}.$$

Here, step (i) follows from the lower bound (3.33), whereas step (ii) follows from applying Lemma 9 in the regime $r > r_{\min}(\kappa_n(4\alpha_n))$. With some further algebra, we find that

$$\mathrm{FNR}_n(T_n) \geq \frac{1}{Z_\ell} \exp\big(2\log\big(4\zeta\big) \cdot s_n^{\frac{1-\gamma}{\gamma}}\big)\, n^{-D_\gamma(s+\kappa_n,r)} \;=\; c_0 \exp\big(c_1 s_n^{\frac{1-\gamma}{\gamma}}\big)\, n^{-D_\gamma(s+\kappa_n,r)},$$

where $c_0 := \frac{1}{Z_\ell}$ and $c_1 := 2\log\big(4\zeta\big)$. Note that because $Z_\ell > 0$ and $\zeta \geq 1$, the constants c_0 and c_1 are both positive, as claimed in the theorem statement.

Proof of Theorem 6

We now sketch the proof of Theorem 6, which states that the BH and BC algorithms achieve the minimax rate (3.20) when $r_n > r_{\min}(\kappa_n(c_{\mathrm{A}}\alpha_n))$, where $\mathrm{A} \in \{\mathrm{BH},\ \mathrm{BC}\}$ and c_{A} is the algorithm-dependent constant defined in (3.24). The details are relegated to Section 3.6.

The proof strategy for both algorithms is essentially the same. Given a target FDR rate α_n, we apply each algorithm with α_n as the target FDR level, and then prove that the resulting threshold satisfies $t_{\mathrm{A}} \leq \tau_{\min}(c_{\mathrm{A}}\alpha_n)$ with high probability. Letting $\tau_{\min,\mathrm{A}} = \tau_{\min}\big(c_{\mathrm{A}}\alpha_n\big)$, we can formulate the specific claims we seek as:

$$\mathbb{P}\big(t_{\mathrm{BH}} > \tau_{\min,\mathrm{BH}}\big) \leq \exp\big(-\frac{n^{1-r_{\max}}}{24}\big) \tag{3.35}$$

and

$$\mathbb{P}\left(t_{\mathrm{BC}} > \tau_{\min,\mathrm{BC}}\right) \leq \alpha_n + \exp\big(-\frac{n^{1-r_{\max}}}{24}\big). \tag{3.36}$$

The known properties of the algorithms guarantee the required FDR bounds [as studied by 4, 33, 8]. The following converse to Lemma 9, coupled with the probabilistic upper bounds (3.35) and (3.36), provides the requisite upper bounds on the FNR.

Lemma 10. *If $r_n > r_{\min}\big(c\alpha_n\big)$ and $t \leq \tau_{\min}(c\alpha_n)$, for some $c > 0$, then we have*

$$\mathrm{FNR}_n\big(t\big) \leq \frac{\big(\max\big\{c,\ 1/c\big\} \cdot \zeta\big)^{2s_n^{\frac{1-\gamma}{\gamma}}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r)},$$

where constant ζ is defined in (3.24).

3.5 Discussion

3.6 Additional proofs

Proof of (3.33)

In order to establish the claim (3.33), define the events

$$\mathcal{E}_1 := \big\{T_n \geq \tau_{\min}\big(4\alpha_n\big)\big\}, \quad \text{and} \quad \mathcal{E}_2 := \left\{\mathrm{FNP}_n\big(\tau_{\min}\big(4\alpha_n\big)\big) \geq \frac{\mathrm{FNR}_n\big(\tau_{\min}\big(4\alpha_n\big)\big)}{2}\right\}.$$

The following lemma guarantees that both of these events have a non-vanishing probability:

Lemma 11. *For any threshold T_n such that $\mathrm{FDR}_n(T_n) \leq \alpha_n$, we have*

$$\mathbb{P}[\mathcal{E}_1] \overset{(a)}{\geq} 3/8, \quad \textit{and} \quad \mathbb{P}[\mathcal{E}_2] \overset{(b)}{\geq} 3/4. \tag{3.37a}$$

We prove this lemma below. Using it, we can complete the proof of claim (3.33). Define the event

$$\mathcal{E} := \left\{ \mathrm{FNP}_n(T_n) \geq \frac{\mathrm{FNR}_n\big(\tau_{\min}(4\alpha_n)\big)}{2} \right\}.$$

The monotonicity of the function $t \mapsto \mathrm{FNP}_n(t)$ ensures that the inclusion $\mathcal{E} \supseteq \mathcal{E}_1 \cap \mathcal{E}_2$ must hold. Consequently, we have

$$\mathbb{P}[\mathcal{E}] \geq \mathbb{P}[\mathcal{E}_1 \cap \mathcal{E}_2] \;\geq\; \mathbb{P}[\mathcal{E}_2] - \mathbb{P}[\mathcal{E}_1^c] \;\overset{(i)}{\geq}\; \frac{3}{4} - \frac{5}{8} \;=\; 1/8,$$

where step (i) follows by applying the probability bounds from Lemma 11.

Finally, by Markov's inequality, we have

$$\mathrm{FNR}_n(T_n) \;=\; \mathbb{E}[\mathrm{FNP}_n(T_n)] \geq \mathbb{P}[\mathcal{E}]\, \frac{\mathrm{FNR}_n\big(\tau_{\min}(4\alpha_n)\big)}{2} \;\geq\; \frac{\mathrm{FNR}_n\big(\tau_{\min}(4\alpha_n)\big)}{16},$$

which establishes the claim (3.33).

Proof of Lemma 11

Our proof makes use of the following auxiliary lemma:

Lemma 12. *For $\alpha_n \in (0, 1/24)$, we have*

$$\mathbb{P}\Big[\mathrm{FDP}_n\big(t\big) \geq 8\alpha_n \quad \textit{for all } t \in [0, \tau_{\min}(4\alpha_n)]\Big] \geq \frac{1}{2}. \tag{3.38}$$

We return to prove this claim in Section 3.6. For the moment, we take it as given and complete the proof of Lemma 11.

Control of $\mathcal{E}_1$: Let us now prove the first bound in Lemma 11, namely that $\mathbb{P}[\mathcal{E}_1] \geq \frac{3}{8}$ where $\mathcal{E}_1 := \big\{T_n \geq \tau_{\min}(4\alpha_n)\big\}$. So as to simplify notation, let us define the event

$$\mathcal{D} := \Big\{\mathrm{FDP}_n\big(t\big) \geq 8\alpha_n \quad \text{for all } t \in [0, \tau_{\min}(4\alpha_n)]\Big\}. \tag{3.39a}$$

Now observe that

$$\mathbb{P}\Big[T_n \geq \tau_{\min}\big(4\alpha_n\big)\Big] \;\geq\; \mathbb{P}\Big[T_n \geq \tau_{\min}\big(4\alpha_n\big) \text{ and } \mathcal{D}\Big] \;=\; \mathbb{P}[\mathcal{D}] - \mathbb{P}\big[T_n \leq \tau_{\min}\big(4\alpha_n\big) \text{ and } \mathcal{D}\big]. \tag{3.39b}$$

Now by the definition (3.39a) of the event $\mathcal{D}$, we have the inclusion

$$\Big\{T_n \leq \tau_{\min}\big(4\alpha_n\big) \text{ and } \mathcal{D}\Big\} \subseteq \Big\{\mathrm{FDP}_n(T_n) \geq 8\alpha_n\Big\}.$$

Combining with our earlier bound (3.39b), we see that

$$\mathbb{P}\Big[T_n \geq \tau_{\min}\big(4\alpha_n\big)\Big] \geq \mathbb{P}[\mathcal{D}] - \mathbb{P}\big[\mathrm{FDP}_n(T_n) \geq 8\alpha_n\big].$$

It remains to control the two probabilities on the right-hand side of this bound. Applying Lemma 12 guarantees that $\mathbb{P}[\mathcal{D}] \geq \frac{1}{2}$. On the other hand, by Markov's inequality, the assumed lower bound $\mathrm{FDR}_n\big(T_n\big) \leq \alpha_n$ implies that $\mathbb{P}\Big[\mathrm{FDP}_n\big(T_n\big) \geq 8\alpha_n\Big] \leq \frac{1}{8}$. Putting together the pieces, we conclude that

$$\mathbb{P}[\mathcal{E}_1] \;=\; \mathbb{P}\Big[T_n \geq \tau_{\min}\big(4\alpha_n\big)\Big] \;\geq\; \frac{1}{2} - \frac{1}{8} \;=\; \frac{3}{8},$$

as claimed.

Control of $\mathcal{E}_2$: Let us now prove the lower bound $\mathbb{P}[\mathcal{E}_2] \geq 3/4$. We split our analysis into two cases.

Case 1: First, suppose that $r_n > r_{\min}$. In this case, we can write

$$\mathrm{FNP}_n(t) = \frac{F_n(t)}{n^{1-s_n}}, \qquad \text{where } F_n\big(t\big) \sim \mathrm{Bin}\big(1 - \Psi\big(t-\mu\big),\; n^{1-s_n}\big).$$

Since $r_{\min} > s_n$, we have $\big|t-\mu\big| = \mu - \tau_{\min}$ and

$$\mu - \tau_{\min} \leq \big(\gamma \log n\big)^{1/\gamma}\big[r_n^{1/\gamma} - s_n^{1/\gamma}\big] \leq \big(\gamma \log n\big)^{1/\gamma} \cdot \big(r_n - s_n\big)^{1/\gamma},$$

from which it follows that

$$\frac{\mathbb{E}[F_n]}{n^{1-s_n}} \;=\; 1 - \Psi\big(t-\mu\big) \geq \frac{n^{s_n - r_n}}{Z_\ell}. \tag{3.40}$$

Now by applying the Bernstein bound to the binomial random variable F_n, we have

$$\begin{aligned} 1 - \mathbb{P}[\mathcal{E}_2] = \mathbb{P}\left[F_n \leq \frac{\mathbb{E}[F_n]}{2}\right] &\leq \exp\big(-\frac{\mathbb{E}[F_n]}{12}\big) \\ &\overset{(i)}{\leq} \exp\big(-\frac{n^{1-r_n}}{12Z_\ell}\big) \\ &\overset{(ii)}{\leq} \exp\big(-\frac{n^{1-r_{\max}}}{12Z_\ell}\big), \end{aligned} \tag{3.41}$$

where step (i) follows from the lower bound (3.40), and step (ii) follows since $r_n < r_{\max}$ by assumption.

Case 2: Otherwise, we may assume that $r_n \in \big(s_n, r_{\min}\big)$. In this regime, we have the lower bound $\tau_{\min} - \mu \geq 0$, so that the binomial random variable F_n stochastically dominates a second binomial distributed as $\widetilde{F}_n \sim \text{Bin}\big(\frac{1}{2},\, n^{1-s_n}\big)$. By this stochastic domination condition, it follows that

$$1 - \mathbb{P}\big[\mathcal{E}_2\big] \;\leq\; \mathbb{P}\Big[F_n \leq \frac{n^{1-s_n}}{4}\Big] \leq \mathbb{P}\Big[\widetilde{F}_n \leq \frac{\mathbb{E}\left[\widetilde{F}_n\right]}{2}\Big].$$

By applying the Bernstein bound to $\widetilde{F}_n$, we find that

$$1 - \mathbb{P}\big[\mathcal{E}_2\big] \leq \exp\big(-\frac{n^{1-s_n}}{24}\big) \;\leq \exp\big(-\frac{n^{1-r_{\max}}}{24}\big), \tag{3.42}$$

where the final step follows since $r_{\max} > s_n$.

Putting together the two bounds (3.41) and (3.42), we conclude that $\mathbb{P}[\mathcal{E}_2] \geq \frac{3}{4}$ for all sample sizes n large enough to ensure that

$$\max\big\{\exp\big(-\frac{n^{1-r_{\max}}}{24Z_\ell}\big),\;\; \exp\big(-\frac{n^{1-r_{\max}}}{24}\big)\big\} \leq \frac{1}{4}, \tag{3.43}$$

as was claimed. Note that condition (3.43) is identical to condition (3.18), so that our definition of $n_{\min}$ guarantees that (3.43) is satisfied. This completes the proof.

Proof of Lemma 12

It remains to prove our auxiliary result stated in Lemma 12. For notational economy, let $\tau = \tau_{\min}(s)$ and let $s = s_n$. The FDP at a threshold t can be expressed in terms of two binomial random variables

$$L_n\big(t\big) = \sum_{i \in \mathcal{H}_0} \mathbf{1}\big(X_i \geq t\big) \;\;\text{and}\;\; W_n\big(t\big) = \sum_{i \notin \mathcal{H}_0} \mathbf{1}\big(X_i \geq t\big) \leq n^{1-s}.$$

Here $L_n(t)$ and $W_n(t)$ correspond (respectively) to the number of nulls, and the number of signals that exceed the threshold t. In terms of these two binomial random variables, we have the expression

$$\text{FDP}_n(t) = \frac{L_n(t)}{L_n(t) + W_n(t)} \geq \frac{L_n(t)}{L_n(t) + n^{1-s}}.$$

Note that the inequality here follows by replacing $W_n(t)$ by the potentially very loose upper bound n^{1-s}; doing so allows us to reduce the problem of bounding the FDP to control of $L_n(t)$ uniformly for $t \in [0, \tau]$. By definition of $L_n(t)$, we have the lower bound

$$\frac{L_n(t)}{L_n(t) + n^{1-s}} \geq \frac{L_n(\tau)}{L_n(\tau) + n^{1-s}} \qquad \text{for all } t \in [0, \tau].$$

Moreover, observe that

$$\frac{3s}{1+3s} \geq \frac{12}{5}s \geq 2s \qquad \text{for all } s \in (0, 1/6).$$

Combining these bounds, we find that

$$\begin{aligned}\mathbb{P}\Big[\mathrm{FDP}_n(t) \geq 2s \quad \text{for all } t \in [0,\tau]\Big] &\geq \mathbb{P}\Big[\frac{L_n(\tau)}{L_n(\tau)+n^{1-s}} \geq \frac{3s}{1+3s}\Big] \\ &= \mathbb{P}\big[L_n(\tau) \geq 3sn^{1-s}\big].\end{aligned}$$

Consequently, the remainder of our proof is devoted to proving that

$$\mathbb{P}\big[L_n(\tau) \geq 3sn^{1-s}\big] \geq 1/2, \tag{3.44}$$

where $s = 4\alpha_n \in \big(0,\ 1/6\big)$ by assumption. We split our analysis into two cases:

Case 1: First, suppose that $\alpha_n \geq \frac{2\log 4}{3n^{1-s}}$ In this case, we have

$$\eta \colon = \Psi\big(\tau\big) \geq \frac{6\xi}{n^s} > \frac{16\log 4}{n}. \tag{3.45}$$

A simple calculation based on this inequality yields

$$\eta n - 3\xi n^{1-s} \geq \frac{\eta n}{2} \geq \sqrt{\big(4\log 4\big)\alpha\big(1-\eta\big)n} := a\sigma, \tag{3.46}$$

where $a = \sqrt{4\log 4}$ and $\sigma = \sqrt{\eta\big(1-\eta\big)n}$. Notice that $\sigma^2 = \mathrm{Var}\,[L_n(\tau)]$.

We now apply the Bernstein inequality to $L_n(\tau)$ to obtain

$$\begin{aligned}\mathbb{P}\big(L_n \leq 3\xi n^{1-s}\big) &\leq \mathbb{P}\big(L_n \leq \eta n - a\sigma\big) \\ &\leq 2\cdot\exp\big(-\frac{a^2\sigma^2}{2\big[\sigma^2 + a\sigma\big]}\big) \\ &\leq 2\cdot\exp\big(-\frac{a^2}{2\big(1+\frac{a}{\sigma}\big)}\big). \\ &\leq \exp\big(-\frac{a^2}{4}\big) = \frac{1}{4},\end{aligned}$$

where we have used the fact that $a < \sigma$. We conclude that

$$\mathbb{P}\big(L_n \geq 3\xi n^{1-s}\big) \geq \frac{1}{2},$$

as desired.

Case 2: Otherwise, we may assume that $\alpha_n < \frac{2\log 4}{3n^{1-s}}$. The definition of τ implies that

$$\eta \geq \frac{24}{n} \quad \text{and} \quad 3\xi n^{1-s} \leq 8\log 4.$$

It follows that $\mathbb{E}[L_n(\tau)] \geq 24$. On the other hand, given that $8\log 4 < 12$, it suffices to prove that

$$P\big[L_n(\tau) \leq 12\big] \leq \frac{1}{2}.$$

This is straightforward, however, since Bernstein's inequality gives

$$\begin{aligned}\mathbb{P}\big[L_n(\tau) \leq 12\big] \;=\; \mathbb{P}\big[L_n(\tau) \leq \frac{\mathbb{E}[L_n(\tau)]}{2}\big] &\leq \exp\big(-\frac{24}{12}\big) \\ &= e^{-2} \\ &< \frac{1}{2},\end{aligned}$$

which completes the proof.

Proof of Lemmas 9 and 10

This section is devoted to the proofs of Lemmas 9 and 10 from earlier in the chapter. We combine the proofs, since these two lemmas provide lower and upper bounds, respectively, on the because they are matching lower and upper bounds, respectively, on the FNR for a fixed threshold procedure, and their proofs involve extremely similar calculations.

So as to simplify notation, we make use of the convenient shorthands let $\tau = \tau_{\min}(\alpha_n)$, $s = s_n$, and $\mu = \mu_n$ throughout the proof. Recall that the FNP can be written as the ratio $\mathrm{FNP}_n(t) = \frac{F_n(t)}{n^{1-s}}$, where

$$F_n(t) = \sum_{i \notin \mathcal{H}_0} \mathbf{1}\big(X_i \leq t\big) \sim \mathrm{Bin}\Big(1 - \Psi\big(t-\mu\big),\; n^{1-s}\Big) \tag{3.47}$$

is a binomial random variable. We split the remainder of the analysis into two cases.

Case 1: First, suppose that $\tau \geq \mu$. In this case, we only seek to prove a lower bound. For this, observe that $\Psi(\tau - \mu) \leq \Psi(0) = \frac{1}{2}$, so $1 - \Psi(\tau - \mu) \geq \frac{1}{2}$. Thus,

$$\mathrm{FNR}_n(\tau) = \frac{\mathbb{E}[F_n]}{n^{1-s}} = 1 - \Psi(\tau - \mu) \geq \frac{1}{2},$$

as claimed.

Case 2: Otherwise, we may assume that $\mu > \tau$. Recall the parameterization (3.13) of μ in terms of r, the definition (3.15) of $r_{\min}$, and the definition (3.19) of the D_γ distance. In terms of these quantities, we have

$$\begin{aligned}\mu - \tau &= \big(\gamma \log n\big)^{1/\gamma} \left\{r^{1/\gamma} - r_{\min}(\kappa_n)\right\}^{1/\gamma} \\ &= \Big\{\gamma D_\gamma\big(r_{\min}(\kappa_n), r\big) \ \log n\Big\}^{1/\gamma} \\ &= \big[\gamma D_\gamma\Big(s + \kappa_n + \frac{\log \frac{1}{6Z_\ell}}{\log n}, r\Big) \ \log n\big]^{1/\gamma},\end{aligned}$$

which shows how the quantity D_γ determines the rate. In order to complete the proof, we need to show that the additional order of $\frac{1}{\log n}$ term inside D_γ can be removed.

More precisely, it suffices to establish the sandwich relation

$$\frac{\zeta^{2s^{\frac{1-\gamma}{\gamma}}}}{Z_u} \cdot n^{-D_\gamma(s+\kappa_n, r)} \geq 1 - \Psi\big(\tau - \mu\big) \geq \frac{\zeta^{2s^{\frac{1-\gamma}{\gamma}}}}{Z_\ell} \cdot n^{-D_\gamma(s+\kappa_n, r)},$$

where $\zeta = \max\left\{6Z_\ell, \ \frac{1}{6Z_\ell}\right\}$ as in (3.24). But now note that

$$\tau - \mu = \big(\gamma \log n\big)^{1/\gamma}\big[\big(r_{\min}\big)^{1/\gamma} - r\big] = -\big[\gamma D_\gamma\left(r_{\min}, r\right) \log n\big]^{1/\gamma},$$

allowing us to deduce that

$$\frac{1}{Z_u} \cdot n^{-D_\gamma(r_{\min}, r)} \geq 1 - \Psi\big(\tau - \mu\big) \geq \frac{1}{Z_\ell} \cdot n^{D_\gamma(r_{\min}, r)},$$

so we need only show

$$\big|D_\gamma\left(s + \kappa_n, r\right) - D_\gamma\left(r_{\min}, r\right)\big| \leq \frac{s^{\frac{1-\gamma}{\gamma}} \log \zeta}{\log n}.$$

To prove this, we let

$$\tilde{r} \ := \min\big(s + \kappa_n, \ r_{\min}\big)$$

and note that by (3.16), we must have $\tilde{r} \in [s, \ r]$. Under this definition, we consider the function $f(x) = D_\gamma\left(\tilde{r} + x, r\right)$. A simple calculation shows that for $x \geq 0$, we have

$$f'(x) = \begin{cases} -\big(\tilde{r} + x\big)^{\frac{1-\gamma}{\gamma}} D_\gamma\left(\tilde{r} + x, r\right)^{\frac{\gamma-1}{\gamma}} & \text{if } \tilde{r} + x \leq r, \\ \big(\tilde{r} + x\big)^{\frac{1-\gamma}{\gamma}} D_\gamma\left(\tilde{r} + x, r\right)^{\frac{\gamma-1}{\gamma}} & \text{o.w.} \end{cases}$$

We observe that we only need to allow $0 \leq x \leq \max\big(s + \kappa_n, \ r_{\min}\big) - \tilde{r} =: \tilde{R} - \tilde{r}$, so in particular, we will always have $\tilde{r} + x \leq \tilde{R} \leq 2$. This, together with the lower bound $\tilde{r} \geq s$, yields

$$\sup_{0 \leq x \leq \tilde{R} - \tilde{r}} \big|f'(x)\big| \leq 2s^{\frac{1-\gamma}{\gamma}}.$$

Applying this result, we find

$$\begin{aligned}\left|D_\gamma\left(s+\kappa_n, r\right)-D_\gamma\left(r_{\min}, r\right)\right| &= \left|D_\gamma\left(\tilde{R}, r\right)-D_\gamma\left(\tilde{r}, r\right)\right| \\ &\leq 2 s^{\frac{1-\gamma}{\gamma}} \cdot\left(\tilde{R}-\tilde{r}\right) \\ &= 2 s^{\frac{1-\gamma}{\gamma}} \cdot \frac{\log \zeta}{\log n}.\end{aligned}$$

If we now consider $\beta_n = c\alpha_n$, we can recover the more refined statements in Lemmas 9 and 10, simply by noting that the same reasoning as above shows

$$\left|D_\gamma\left(s+\kappa_n, r\right)-D_\gamma\left(s+\kappa_n', r\right)\right| \leq 2 s^{\frac{1-\gamma}{\gamma}} \cdot \frac{\left|\log c\right|}{\log n},$$

concluding the argument.

Proof of Corollary 4

Although Corollary 4 can be proved from the statement of Theorem 5, we instead prove it more directly, as this allows us to reuse parts of the proof of Lemma 9, thereby saving some additional messy calculations.

First, we verify that there is indeed a unique solution κ_* to the fixed point equation (3.21). Define the function as $g(\kappa) := D_\gamma\left(s+\kappa, r\right)^{1/\gamma} - \kappa^{1/\gamma}$. Clearly the solutions to (3.21) are the roots of g. We would like to argue that any such root must occur in $[0,\ r-s)$ and that in fact g has a unique root in this interval. For the first claim, note that $g(r-s) = -(r-s)^{1/\gamma} < 0$. On the other hand, we have

$$g'(\kappa) = \begin{cases} -\frac{1}{\gamma}\left[(s+\kappa)^{-\frac{\gamma-1}{\gamma}} + \kappa^{-\frac{\gamma-1}{\gamma}}\right] & \text{if } 0 \leq \kappa < r-s, \\ \frac{1}{\gamma}\left[(s+\kappa)^{-\frac{\gamma-1}{\gamma}} - \kappa^{-\frac{\gamma-1}{\gamma}}\right] & \text{if } \kappa > r-s. \end{cases}$$

It is immediately clear that $g'(\kappa) < 0$ for $0 \leq \kappa < r-s$ and, since $s+\kappa > \kappa$, we may also deduce that $g'(\kappa) < 0$ for $\kappa > r-s$, so g is decreasing on its domain. Therefore, $g(\kappa) < g(r-s) < 0$ for all $\kappa > r-s$.[2] We conclude that any root of g must occur on $[0,\ r-s)$. To finish the argument, note that $g(0) > 0 > g(r-s)$, so that g does indeed have a root on $[0,\ r-s)$.

Turning now to the proof of the lower bound (3.22), let $\mathcal{I}$ be an arbitrary threshold-based multiple testing procedure. We may assume without loss of generality that

$$\mathrm{FDR}_n(\mathcal{I}) \leq \min\left\{n^{-\kappa_*}, \frac{1}{24}\right\} \leq c(s,\gamma) n^{-\kappa_*}, \tag{3.48}$$

[2]Note that we have suppressed the issue of non-differentiability of g at $\kappa = r-s$. We may do so because it is left- and right-differentiable at this point, and we argue separately for the intervals $[0,\ r-s)$ and $[r-s,\ \infty)$.

where the quantity $c(s,\gamma) \geq 1$ was defined in the statement of Theorem 5 (otherwise, the claimed lower bound (3.22) follows immediately).

Applying the second part of Lemma 9 and defining $\tilde{c} = 4c(s,\gamma)$, we conclude that

$$\begin{aligned}
\mathrm{FNR}_n\big(T_n\big) &\geq \frac{\mathrm{FNR}_n\big(\tau_{\min}\big(\tilde{c}n^{-\kappa_*}\big)\big)}{16} \\
&\geq \frac{\big(\tilde{c}\zeta\big)^{2s\frac{1-\gamma}{\gamma}}}{Z_\ell} \cdot n^{-D_\gamma(s+\kappa^*,r)} \\
&= \frac{\big(\tilde{c}\zeta\big)^{2s\frac{1-\gamma}{\gamma}}}{Z_\ell} \cdot n^{-\kappa_*} \\
&= c' n^{-\kappa_*}.
\end{aligned}$$

Proof details for Theorem 6

Achievability for the BH procedure

In this section, we prove that BH achieves the lower bound whenever $r_n > r_{\min}\big(c_{\mathrm{BH}}\alpha_n\big)$. Specifically, we prove the claim (3.26) stated in Theorem 6.

We first show how to derive the upper bound (3.26) from the probability bound (3.35) and then prove the probability bound itself. Note that since BH is a valid FDR control procedure, we necessarily have $\mathrm{FDR}_n(t_{\mathrm{BH}}) \leq \alpha_n$. To bound the FNR, first let $\mathcal{E} = \{t_{\mathrm{BH}} \leq \tau_{\min,\mathrm{BH}}\}$ and let $\mathrm{FNR}_n(\cdot \mid \mathcal{E})$ and $\mathrm{FNR}_n(\cdot \mid \mathcal{E}^c)$ denote the FNR_n conditional on the event and its complement, respectively. In this notation, the bound (3.35), together with Lemma 10, implies that

$$\begin{aligned}
\mathrm{FNR}_n(t_{\mathrm{BH}}) &\leq \mathbb{P}\big(\mathcal{E}\big) \cdot \mathrm{FNR}_n\big(\tau_{\min,\mathrm{BH}} \mid \mathcal{E}\big) + \mathbb{P}\big(\mathcal{E}^c\big) \\
&\leq \mathrm{FNR}_n\big(\tau_{\min,\mathrm{BH}}\big) + \mathbb{P}\big(\mathcal{E}^c\big) \\
&\leq \mathrm{FNR}_n\big(\tau_{\min,\mathrm{BH}}\big) + \exp\big(-\frac{n^{1-r_{\max}}}{24}\big) \\
&\leq \frac{\zeta_{\mathrm{BH}}^{2s_n\frac{1-\gamma}{\gamma}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r_n)} + \exp\big(-\frac{n^{1-r_{\max}}}{24}\big) \\
&\leq \frac{2\zeta_{\mathrm{BH}}^{2s_n\frac{1-\gamma}{\gamma}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r_n)},
\end{aligned}$$

where the final step uses the definition (3.23) of $n_{\min,u}$, and the fact that $\frac{1}{Z_u n} \leq \frac{\zeta_{\mathrm{BH}}^{2s_n\frac{1-\gamma}{\gamma}}}{Z_u} \cdot n^{-D_\gamma(s_n+\kappa_n,r_n)}$, which is easily verified by noting that $\zeta_{\mathrm{BH}}^{2s_n\frac{1-\gamma}{\gamma}} \geq 1$ and $D_\gamma\left(s_n+\kappa_n, r_n\right) \leq 1$.

We now prove the probability bound with an argument using p-values and survival functions that parallels that of Arias-Castro and Chen [4] but that sidesteps CDF asymptotics.

To carry out the analysis, we first study the relationship between the population survival function Ψ and the empirical survival function $\hat{\Psi}$, defined by

$$\hat{\Psi}\big(t\big) = \big(1 - \frac{1}{n^{s_n}}\big) \cdot \hat{\Psi}_0\big(t\big) + \frac{1}{n^{s_n}} \cdot \hat{\Psi}_1\big(t\big), \tag{3.49}$$

$$\text{where} \qquad \hat{\Psi}_0\big(t\big) = \frac{1}{n - n^{1-s_n}} \sum_{i \in \mathcal{H}_0} \mathbf{1}\big(X_i \geq t\big) \;\; \text{and} \;\; \hat{\Psi}_1\big(t\big) = \frac{1}{n^{1-s_n}} \sum_{i \notin \mathcal{H}_0} \mathbf{1}\big(X_i \geq t\big).$$

Now, sort the observations in *decreasing* order, so that $X_{(1)} \geq X_{(2)} \geq \cdots \geq X_{(n)}$, and define p-values

$$p_{(i)} = \Psi\left(X_{(i)}\right) \;\; \text{and} \;\; \hat{\Psi}\left(X_{(i)}\right) = \frac{i}{n}, \tag{3.50}$$

so that $p_{(1)} \leq p_{(2)} \leq \cdots \leq p_{(n)}$ are in *increasing* order. Then, we may characterize the indices rejected by BH as those satisfying $X_i \geq X_{(i_{\mathrm{BH}})}$, where

$$i_{\mathrm{BH}} = \max\left\{1 \leq i \leq n \colon \Psi\left(X_{(i)}\right) \leq \alpha_n \hat{\Psi}\left(X_{(i)}\right)\right\}. \tag{3.51}$$

Moving t_{BH} within $\left(X_{(i_{\mathrm{BH}}+1)},\; X_{(i_{\mathrm{BH}})}\right]$ if necessary, we may therefore assume $\Psi(t) > \alpha_n \hat{\Psi}(t)$ whenever $t < t_{\mathrm{BH}}$, and combining this knowledge with (3.49), we obtain the chain of inclusions

$$\begin{aligned} \mathcal{E}^c = \{t_{\mathrm{BH}} > \tau_{\min,\mathrm{BH}}\} &\subset \left\{\Psi\big(\tau_{\min,\mathrm{BH}}\big) > \alpha_n \hat{\Psi}\big(\tau_{\min,\mathrm{BH}}\big)\right\} \\ &\subset \left\{\Psi\big(\tau_{\min,\mathrm{BH}}\big) > \frac{\alpha_n}{n^{s_n}} \cdot \frac{W_n}{n^{1-s_n}}\right\} =: \widetilde{\mathcal{E}}^c, \end{aligned} \tag{3.52}$$

where $W_n = \sum_{i \notin \mathcal{H}_0} \mathbf{1}\big(X_i \geq \tau_{\min,\mathrm{BH}}\big) \sim \mathrm{Bin}\big(\Psi\big(\tau_{\min,\mathrm{BH}} - \mu_n\big),\; n^{1-s_n}\big)$.

We now argue that $\Psi\big(\tau_{\min,\mathrm{BH}}\big) \leq \frac{\alpha_n}{4n^{s_n}}$, so that $\mathbb{P}\left(\mathcal{E}^c\right) \leq \mathbb{P}\left(\widetilde{\mathcal{E}}^c\right) \leq \mathbb{P}\left(W_n \leq \frac{n^{1-s_n}}{4}\right)$. For this, observe that by the definition of $r_{\min}$ in (3.15) and the upper tail bound (3.5), we have

$$\begin{aligned} \log \Psi\big(\tau_{\min,\mathrm{BH}}\big) &\leq -r_{\min}\left(c_{\mathrm{BH}}\alpha_n\right) \log n + \log \frac{1}{Z_u} \\ &\leq -s_n \log n + \log(c_{\mathrm{BH}}\alpha_n) - \log \frac{1}{6Z_\ell} + \log \frac{1}{Z_u} \\ &= -s_n \log n + \log \alpha_n + \log \frac{6c_{\mathrm{BH}} Z_\ell}{Z_u} \\ &= \log \frac{\alpha_n}{6n^{s_n}} \;\; < \;\; \log \frac{\alpha_n}{4n^{s_n}}. \end{aligned}$$

We conclude

$$\mathbb{P}\left(t_{\mathrm{BH}} > \tau_{\min,\mathrm{BH}}\right) \;\; \leq \;\; \mathbb{P}\left(\widetilde{\mathcal{E}}^c\right) \;\; \leq \;\; 1 - \mathbb{P}\left(W_n > \frac{n^{1-s_n}}{4}\right) \;\; = \;\; \mathbb{P}\left(W_n \leq \frac{n^{1-s_n}}{4}\right).$$

Finally, by a Bernstein bound, we find

$$\begin{aligned}\mathbb{P}\left(W_n \leq \frac{n^{1-s_n}}{4}\right) &\leq \mathbb{P}\left(W_n \leq \frac{\mathbb{E}\left[W_n\right]}{2}\right) \\ &\leq \exp\big(-\frac{\mathbb{E}\left[W_n\right]}{12}\big) \\ &\leq \exp\big(-\frac{n^{1-s_n}}{24}\big) \\ &\leq \exp\big(-\frac{n^{1-r_{\max}}}{24}\big),\end{aligned}$$

where we have used the fact that $\tau_{\min,\mathrm{BH}} \leq \mu_n$ to conclude that $\Psi\left(\tau_{\min,\mathrm{BH}} - \mu_n\right) \geq \frac{1}{2}$ and therefore $\mathbb{E}[W_n] \geq \frac{n^{1-s_n}}{2}$. We have therefore established the required claim (3.35), concluding the proof of optimality of the BH procedure.

Achievability for the BC procedure

Our overall strategy for analyzing BC procedure resembles the one we used for the BH procedure. As with our analysis of the BH procedure, we define $\tau_{\min,\mathrm{BC}} := \tau_{\min}(c_{\mathrm{BC}}\alpha_n)$ and derive the bound (3.26) by controlling the algorithm's threshold as Since the proof of equation (3.27) from the bound (3.36) is essentially identical to the corresponding derivation for the BH procedure, we omit it. We now prove the bound (3.36) by an argument somewhat different than that used in analyzing the BH procedure. Define the integers

$$N_+\big(t\big) = \sum_{i=1}^{n} \mathbf{1}\big(X_i \geq t\big) \quad \text{and} \quad N_-\big(t\big) = \sum_{i=1}^{n} \mathbf{1}\big(X_i \leq -t\big).$$

Then, the definition of the BC procedure gives

$$t_{\mathrm{BC}} = \inf\big\{t \in \mathbb{R} \colon \frac{1 + N_-(t)}{1 \vee N_+(t)} \leq \alpha_n\big\}.$$

To prove (3.36), it therefore suffices to show that

$$\mathbb{P}\left(\frac{1 + N_-(\tau_{\min,\mathrm{BC}})}{1 \vee N_+(\tau_{\min,\mathrm{BC}})} > \alpha_n\right) \leq \alpha_n + \exp\big(-\frac{n^{1-r_{\max}}}{24}\big). \tag{3.53}$$

We prove the bound (3.53) in two parts:

$$\mathbb{P}\big(1 \vee N_+(\tau_{\min,\mathrm{BC}}) < \frac{n^{1-s_n}}{4}\big) \leq \exp\big(-\frac{n^{1-r_{\max}}}{24}\big), \tag{3.54a}$$

$$\mathbb{P}\left(1 + N_-(\tau_{\min,\mathrm{BC}}) > \alpha_n \cdot \frac{n^{1-s_n}}{4}\right) \leq \alpha_n. \tag{3.54b}$$

These bounds are a straightforward consequence of elementary Bernstein bounds, and together they imply the claim (3.53). We explain them below.

The lower bound (3.54a) follows because $1 \vee N_+(\tau_{\min,\mathrm{BC}}) \geq N_+(\tau_{\min,\mathrm{BC}})$ and $N_+(\tau_{\min,\mathrm{BC}})$ is the sum of two binomial random variables, corresponding to nulls and signals, respectively, and the latter has a $\Psi\big(\tau_{\min,\mathrm{BC}} - \mu\big) \geq \frac{1}{2}$ probability of success. More precisely, we may write $N_+\left(\tau_{\min,\mathrm{BC}}\right) = N_+^{\mathrm{null}} + N_+^{\mathrm{signal}}$, with

$$N_+^{\mathrm{null}} \sim \mathrm{Bin}\left(\Psi\left(\tau_{\min,\mathrm{BC}}\right),\ n - n^{1-s_n}\right) \quad \text{and} \quad N_+^{\mathrm{signal}} \sim \mathrm{Bin}\left(\Psi\left(\tau_{\min,\mathrm{BC}} - \mu_n\right),\ n^{1-s_n}\right),$$

implying $N_+\left(\tau_{\min,\mathrm{BC}}\right) \geq N_+^{\mathrm{signal}}$, whence

$$\mathbb{E}\big[N_+(\tau_{\min,\mathrm{BC}})\big] \geq \mathbb{E}\big[N_+^{\mathrm{signal}}\big] = n^{1-s_n} \cdot \Psi\big(\tau_{\min,\mathrm{BC}} - \mu_n\big) \geq \frac{n^{1-s_n}}{2},$$

where we have used the fact that $\tau_{\min,\mathrm{BC}} \leq \mu_n$. With this bound in hand, a Bernstein bound yields

$$\begin{aligned}\mathbb{P}\big(N_+(\tau_{\min,\mathrm{BC}}) < \frac{n^{1-s_n}}{4}\big) &\leq \mathbb{P}\big(N_+(\tau_{\min,\mathrm{BC}}) \leq \frac{\mathbb{E}\left[N_+(\tau_{\min,\mathrm{BC}})\right]}{2}\big) \\ &\leq \exp\big(-\frac{n^{1-s_n}}{24}\big) \leq \exp\big(-\frac{n^{1-r_{\max}}}{24}\big),\end{aligned}$$

as required to prove equation (3.54a). The proof of equation (3.54b) follows a similar pattern. Here, we note that $N_-(\tau_{\min,\mathrm{BC}})$ is a sum of two binomial random variables, with a total of n trials, such that—using the definition (3.15) of $r_{\min}$ and the upper bound on the tail (3.5)—each one has probability of success upper bounded by $1 - \Psi\big(-\tau_{\min,\mathrm{BC}}\big) \leq \frac{6Z_\ell}{Z_u} \cdot c_{\mathrm{BC}}\alpha_n n^{-s_n} = \frac{1}{8} \cdot \alpha_n n^{-s_n}$. Formally, we may write $N_-(\tau_{\min,\mathrm{BC}}) = N_-^{\mathrm{null}} + N_-^{\mathrm{signal}}$, with

$$N_-^{\mathrm{null}} \sim \mathrm{Bin}\left(1 - \Psi\left(-\tau_{\min,\mathrm{BC}}\right),\ n - n^{1-s_n}\right) \quad \text{and} \quad N_-^{\mathrm{signal}} \sim \mathrm{Bin}\left(1 - \Psi\left(-\tau_{\min,\mathrm{BC}} - \mu\right),\ n^{1-s_n}\right).$$

Since $1 - \Psi\left(-\tau_{\min,\mathrm{BC}} - \mu\right) \leq 1 - \Psi\left(-\tau_{\min,\mathrm{BC}}\right)$, we deduce

$$\mathbb{E}\big[N_-(\tau_{\min,\mathrm{BC}})\big] \leq \big[1 - \Psi\left(-\tau_{\min,\mathrm{BC}}\right)\big] \cdot n \leq \frac{\alpha_n}{2} \cdot \frac{n^{1-s_n}}{4}.$$

On the other hand, using the lower bound in (3.5), we find $1 - \Psi\big(-\tau_{\min,\mathrm{BC}}\big) \geq 6c_{\mathrm{BC}}\alpha_n n^{-s_n}$. Using the additional fact that $n - n^{1-s_n} \geq \frac{n}{2}$ by (3.14a), we may conclude that

$$\begin{aligned}\mathbb{E}\big[N_-(\tau_{\min,\mathrm{BC}})\big] &\geq \mathbb{E}\big[N_-^{\mathrm{null}}\big] \\ &= \big(n - n^{1-s_n}\big) \cdot \big[1 - \Psi\big(-\tau_{\min,\mathrm{BC}}\big)\big] \\ &\geq \frac{n}{2} \cdot 6c_{\mathrm{BC}}\alpha_n n^{-s_n} \\ &\geq 3c_{\mathrm{BC}}\alpha_n n^{1-s_n}.\end{aligned}$$

By a Bernstein bound, it follows that

$$\begin{aligned}\mathbb{P}\left(N_-(\tau_{\min,\mathrm{BC}}) \geq \alpha_n \cdot \frac{n^{1-s_n}}{4}\right) &\leq \mathbb{P}\left(N_-(\tau_{\min,\mathrm{BC}}) \geq 2\mathbb{E}\big[N_-(\tau_{\min,\mathrm{BC}})\big]\right) \\ &\leq \exp\left(-\frac{\mathbb{E}\big[N_-(\tau_{\min,\mathrm{BC}})\big]}{4}\right) \\ &\leq \exp\big(-\frac{3c_{\mathrm{BC}}}{4} \cdot \alpha_n n^{1-s_n}\big) \\ &\leq \alpha_n,\end{aligned}$$

where we have invoked the decay condition (3.25) for the last step.

Proof of Corollary 5

The corollary is a nearly immediate consequence of Theorem 6. We will prove it for both algorithms simultaneously. Observe that

$$r_{\min}(\kappa_n(c_{\mathrm{A}}\alpha_*)) = s + \kappa_* + \frac{\log \frac{1}{6c_* c_{\mathrm{A}} Z_\ell}}{\log n}. \tag{3.55}$$

Suppose for now that the decay condition (3.25) holds for α_* and some choice of $n_{\min,\mathrm{BC}}$. Then, using (3.55) and the fact that $r > s + \kappa_*$, we may choose $n'_{\min} \geq n_{\min,\mathrm{BC}}$ large enough so that $r > r_{\min}(\kappa_n(c_{\mathrm{A}}\alpha_*))$ for all $n \geq n'_{\min}$ and $\mathrm{A} \in \{\mathrm{BH}, \mathrm{BC}\}$. From Theorem 6, we conclude that there exists a constant c' such that both algorithms satisfy

$$n \geq n'_{\min} \Longrightarrow \mathcal{R}_n \leq c' n^{-\kappa_*}.$$

By replacing c' by $\tilde{c} = \max\{c', \ (n'_{\min})^{\kappa_*}\}$ (and recalling $\mathcal{R}_n \leq 1$ always), we obtain $\mathcal{R}_n \leq \tilde{c} n^{-\kappa_*}$ for all $n \geq 1$, obtaining the claimed result.

In order to check the decay condition (3.25), note that, as $\kappa_* \leq r - s \leq 1 - s$, we have for sufficiently large n that

$$\frac{\alpha_n}{\log \frac{1}{\alpha_n}} = \frac{n^{-\kappa_*}}{\kappa_* \log n} \geq \frac{4}{3c_{\mathrm{BC}}} \cdot n^{-(1-s)},$$

which completes the proof.

Chapter 4

A General Framework for Minimax Theory for Multiple Testing

4.1 Overview

In the previous chapter, we gave a power analysis for multiple testing under specific modeling assumptions – an analogue of the longstanding tradition of power analysis in single hypothesis testing. We considered false discovery rate (FDR) as the Type-I error concept, and obtained a matching Type-II concept by considering the false non-discovery rate (FNR) — that is, the fraction of tests in which the null was incorrectly *not* rejected.

We now attempt to go beyond the limitations of the results in Chapter 3. Like most power results in the multiple testing literature, the findings of Chapter 3 were limited to models where test statistics are constrained to be independent of one another, as well as Gaussian-like, and in which alternatives are assumed to be location shifts of the null. Models of this type provide the advantage of analytical tractability, while still permitting the expression of central features of many multiple testing problems. Indeed, the results in question mostly apply to variants of the rare-weak (RW) model in which problem difficulty is parameterized by the *rarity* of signals and their *weakness* — a model initially introduced for multiple testing under familywise error rate (FWER) as the Type-I error concept [27, 29, 58],

Unfortunately, the analysis strategies of Chapter 3 do not appear to extend beyond the analytically tractable realm of independent, Gaussian-like test statistics. This limitation appears most dramatic for the most precise, non-asymptotic results, as these depend critically on the ability to control tail probabilities and apply concentration inequalities [78].

In this chapter, we introduce a modeling and analysis strategy less dependent on independence and analytical tractability. Building on the proof strategy of Chapter 3 [78], we seek to compare multiple testing procedures to derandomized versions that always make a *fixed* number of discoveries. Given the fixed proxy k^* for the number of discoveries, we show how to further eliminate the randomness associated with the location of the data to obtain deterministic proxies ℓ^* lower bounding the number of *false* discoveries with constant prob-

ability. Using the two proxies k^* and ℓ^* together yields constant-probability lower bounds on the proportion of false discoveries and false non-discoveries, and these translate directly into corresponding lower bounds on the FDR and FNR that hold for any procedure applied to the given model.

Our core contribution in this chapter is a meta-theorem that establishes a tradeoff between FDR and FNR in terms of these proxies for *any* member of the large model class we consider (see Section 4.2). The result is meta-level in that the evaluation of the proxies may itself pose a substantial technical challenge in many settings. In contrast to minimax theory for estimation, where tools like Le Cam's method, Fano's method, the Yang-Barron method, and Assouad's lemma play the role of meta-theorems that can be instantiated—with some effort—to answer a given minimax rate question [90], such general tools are currently absent from the multiple testing literature. Ours appears to be the first of its kind.

To illustrate how application of our meta-theorem works in practice, we apply it to several specific models. These include the previously-studied independent generalized Gaussians model [4, 78], as well as Gaussian models with dependence, Gaussian models with scale-transformation, and an exponentiation model for p-values.

Related work. Our work is most closely related to the work described in Chapter 3, which gives non-asymptotic tradeoffs between FDR and FNR [78]. Those results apply to a model where the test statistics are assumed to have tails on the order of $\exp\big(-|x|^\gamma\big)$ for some $\gamma \geq 1$ — in other words, generalized Gaussian-like tails. The tradeoff results, as well as previous asymptotic lower bounds [4], apply only to the independent case. Under exact Gaussianity, Ji and Jin [56] and Ji and Zhao [57] provide lower bounds that continue to hold under dependence, at the cost of replacing FDR with the so-called *modified* FDR (mFDR), in which the expectation is moved inside the ratio in the definition of FDR (cf. (3.3)). The mFDR and FDR measures ought to behave similarly for large numbers of tests, but they are distinct metrics, and the analysis strategies that work for mFDR generally do not apply to FDR itself.

The present work is partly inspired by classic results by Jin and Donoho on phase transitions in multiple testing based on Tukey's higher-criticism [27, 29, 58]. Their results apply to the Gaussian sequence model with location shifts, in which the number of signals is a polynomially small fraction n^{-s} of the total number of tests and in which signals are weak (achieved by scaling the shift as $\mu = \sqrt{2r \log n}$ for $0 < r < 1$). Their work establishes the regime of s and r in which asymptotic consistency is possible, under the standard Type-I and Type-II error measures for testing the global null.

Finally, the present work is related at least in spirit to the general toolkit of minimax theory for estimation. Unlike the multiple testing literature, in which minimax theory depends generally on *ad hoc* model-specific calculations, estimation theory benefits from a well-established set of tools that frequently succeed in establishing lower bounds for specific estimation problems of interest [90]. This chapter aims to provide a first such tool for multiple testing.

4.2 Background and problem formulation

In this section, we provide background and a precise formulation of the problem under study. Here, we define only the structure of the multiple testing model and procedures we consider, since the relevant background on the multiple testing problem and its error metrics appears in Section 3.2.

Model structure

In this chapter, we consider a flexible class of models that includes as special cases the location and scale families that have been studied in past work. For a testing problem with n hypotheses, we assume that the vector $X = (X_1, \ldots, X_n)$ of test statistics is generated from some underlying random vector $W = (W_1, \ldots, W_n)$ in the following way. The vector W may be drawn from an arbitrary (not necessarily product) distribution $\mathbb{P}_n$, while the test statistics are related via

$$X_i = \begin{cases} W_i & \text{if } i \in \mathcal{H}_0 \\ f(W_i) & \text{otherwise.} \end{cases} \tag{4.1}$$

The main restriction on this model is that $f \colon \mathbb{R} \to \mathbb{R}$ must be a non-decreasing function such that $f(w) \geq w$ for all w in the support of a marginal of $\mathbb{P}_n$. For convenience, we shall also assume that the marginals of $\mathbb{P}_n$ are atom-free. In some instances, we consider the restricted case where all the W_i are iid—that is, $\mathbb{P}_n = \mathbb{P}_0^{\otimes n}$. We refer to this as the iid model. Conceptually, a model is therefore a tuple $(\mathbb{P}_n,\ f)$ satisfying these constraints, and where appropriate we shall denote models by such tuples and sometimes name them.

Prototypical examples of this general set-up include the following:

Location model: The variables W_i are drawn from a generalized Gaussian distribution (i.e., with density proportional to $\exp(-|w|^\gamma)$ for some $\gamma \in [1, 2]$), and the function $f(w) = \mu + w$ for some $\mu > 0$.

Scale models: The variables W_i are the absolute values of standard normal variates, and the function $f(w) = \sigma w$ for some $\sigma \geq 1$.

Lehmann alternative model: The variables W_i are uniform on the unit interval, and the function $f(w) = 1 - \big(1 - w\big)^{1/\gamma}$ for some parameter $\gamma \in (0, 1)$. This set-up models the situation in which the W_i represent p-values and the signals have p-values that are stochastically closer to zero than those of the nulls. Since we have chosen to model the transformation as *non-decreasing*, we represent the unit interval backwards, which leads to the form written rather than $w^{1/\gamma}$.

All three of these examples have been studied in past work [4, 30, 58, 71, 78]

Top-K procedures

Many popular procedures, including the Benjamini-Hochberg (BH) and several variants thereof [8, 33, 80], are based on thresholding the order statistics

$$\min_{i=1,\ldots,n} X_i =: X_{(1)} \leq X_{(2)} \leq X_{(3)} \leq \cdots \leq X_{(n)} := \max_{i=1,\ldots,n} X_i. \tag{4.2}$$

More precisely, a *top-K* procedure is a method that rejects the top K order statistics, where $K = K_n(X_1^n)$ is a non-negative integer that can depend on the observed statistics. The testing rule $\mathcal{I}\colon \mathbb{R}^n \to 2^{[n]}$ defined by any top-K procedure has the form

$$\mathcal{I}\big(X_1^n\big) = \Big\{ i \in [n] \colon X_i \geq X_{(K_n(X_1^n))} \Big\}, \tag{4.3}$$

where $K_n\colon \mathbb{R}^n \to \mathbb{N}$ is some (possibly randomized) mapping. Alternatively, such procedures can be described in terms of choosing a threshold $\tau = \tau_n(X_1^n)$, and rejecting all nulls i for which $X_i \geq \tau$.

4.3 Main results

We now turn to the statements of our main results. We begin in Section 4.3 by defining the deterministic proxies that play a central role in our analysis; see Section 4.3 for the intuition that underlies these definitions. In Section 4.3, we state a general lower bound (Theorem 7) on the pairs of FDR and FNR that are achievable. In the remaining sections, we illustrate the consequences of this general bound to various specific models.

A general bound based on deterministic proxies for FDR and FNR

We say that a FDR-FNR pair $(\alpha, \beta) \in [0,1] \times [0,1]$ that is *achievable* if there exists a top-K procedure K such that

$$\mathrm{FDR}\big(K\big) \leq \alpha \quad \text{and} \quad \mathrm{FNR}\big(K\big) \leq \beta. \tag{4.4}$$

Any top-K procedure satisfying condition (4.4) is said to be $(\alpha,\ \beta)$-*controlled.* Our high-level goal is to provide bounds on the region of achievable (α, β) pairs.

Defining the deterministic proxies

In order to characterize the space of achievable (α, β) pairs, we construct two sets of deterministic proxies. One set of proxies is useful in the regime $\alpha \geq \beta$, while the other is useful in the opposite setting—namely, $\beta \geq \alpha$. The proxies used in the former regime are denoted by FDP_-^* and FNP_-^*, whereas those in the latter regime are denoted by FNP_+^* and FDP_+^*. The

reasoning underlying our choice of notation should be clear once we detail their construction below. The numerical values of these proxies depend on α, β, the model $\mathbb{M} = \big(\mathbb{P}_n, f\big)$ under consideration, and a parameter $\epsilon \in (0,1)$ that controls the strength of the bounds. We make these dependencies explicit when needed, suppressing them otherwise.

Our first step is to define deterministic approximations of the number of total discoveries. Letting $m = \mathrm{card}(\mathcal{H}_1)$ denote the number of signals and given any $\epsilon > \max\{\alpha, \beta\}$, we define

$$k^*_-(\beta, \epsilon) := \left(1 - \frac{\beta}{\epsilon}\right) m \quad \text{and} \quad k^*_+(\alpha, \epsilon) := \left(1 - \frac{\alpha}{\epsilon}\right)^{-1} m. \tag{4.5}$$

Roughly speaking, the integer k^*_- functions as a lower approximation to the number of *total discoveries*, whereas the quantity k^*_+ provides an upper approximation to the same quantity. Note that these lower and upper bounds converge as $\alpha, \beta \to 0$; in the limit $\alpha = \beta = 0$, we have $k^*_-(0, \epsilon) = k^*_+(0, \epsilon) = m$, since in this case, the total number of discoveries must be equal to the number of signals m.

For each of these approximations of the discovery number, we construct a corresponding false discovery proxy. Recalling the underlying random vector $W = (W_{\mathcal{H}_0}, W_{\mathcal{H}_1})$ that underlies our generic model, these quantities involve the order statistics

$$W_{\mathcal{H}_0,(1)} \leq W_{\mathcal{H}_0,(2)} \leq \cdots \leq W_{\mathcal{H}_0,(|\mathcal{H}_0|)},$$

with the order statistics for $W_{\mathcal{H}_1}$ defined analogously. For any $\epsilon > \max\{\alpha, \beta\}$ and adopting the shorthand $k^*_- = k^*_-(\beta, \epsilon)$ and $k^*_+ = k^*_+(\alpha, \epsilon)$, we then define

False discovery proxies:

$$\ell^*_-(\beta, \epsilon, \mathbb{M}) = \arg \max_{\ell \in [1, k^*_-]} \left\{ \mathbb{P}\Big[W_{\mathcal{H}_0,(\ell)} > f\big(W_{\mathcal{H}_1,(k^*_- - \ell + 1)}\big) \Big] \geq 1 - \epsilon \right\}, \qquad \text{and} \tag{4.6a}$$

$$\ell^*_+(\alpha, \epsilon, \mathbb{M}) = \arg \max_{\ell \in [k^*_+ - m, k^*_+]} \left\{ \mathbb{P}\Big[W_{\mathcal{H}_0,(\ell)} > f\big(W_{\mathcal{H}_1,(k^*_+ - \ell + 1)}\big) \Big] \geq 1 - \epsilon \right\}. \tag{4.6b}$$

Roughly, the quantities ℓ^*_- and ℓ^*_+ represent, respectively, lower and upper approximations to the number of false discoveries.

Finally, by taking appropriate ratios, we define:

Proxies to FDR and FNR:

$$\mathrm{FDP}^*_-(\beta, \epsilon, \mathbb{M}) = \frac{\ell^*_-(\beta, \epsilon, \mathbb{M})}{m} \quad \text{and} \quad \mathrm{FNP}^*_-(\beta, \epsilon, \mathbb{M}) = \frac{m - k^*_-(\beta, \epsilon) + \ell^*_-(\beta, \epsilon, \mathbb{M})}{m}, \tag{4.7a}$$

$$\mathrm{FDP}^*_+(\alpha, \epsilon, \mathbb{M}) = \frac{\ell^*_+(\alpha, \epsilon, \mathbb{M})}{m} \quad \text{and} \quad \mathrm{FNP}^*_+(\alpha, \epsilon, \mathbb{M}) = \frac{m - k^*_+(\alpha, \epsilon) + \ell^*_+(\alpha, \epsilon, \mathbb{M})}{m}. \tag{4.7b}$$

To be clear, in defining FDP^*_+ (respectively FDP^*_-), it might be more natural to use k^*_+ (respectively k^*_-) in the denominator, but as noted above, when (α, β) are small, both of these quantities are close to m.

The underlying intuition

Let us now describe the intuition that underlies the definitions (4.6), ignoring the difference between the + and − versions to simplify matters. First, suppose that we accept that k is a good approximation to the total number of discoveries, and that ℓ is a good approximation to the number of false discoveries. In this case, ℓ/k is a good approximation to the FDR, and since $k-\ell$ of the discoveries must be false, then $\frac{m-(k-\ell)}{m}$ should be a good approximation to the FNR. As we have argued above, when β and α are small, then k is actually relatively close to m, so that ℓ/m should also be a good approximation to the FDR.

It remains to justify why ℓ, as defined in equation (4.6), is a reasonable proxy to the number of false discoveries. Consider a procedure that rejects exactly k hypotheses, of which ℓ are nulls. It must then be case that ℓ^{th} largest null value exceeds the value of the $\big(k-\ell+1\big)^{\text{th}}$ largest signal value, or else only $\ell-1$ nulls would be in the top k test statistics. Using the definition of the model, we can re-express this relation in symbols:

$$W_{\mathcal{H}_{0,(\ell)}} > f\left(W_{\mathcal{H}_{1},(k-\ell+1)}\right). \tag{4.8}$$

The definitions (4.6) are motivated by this fact.

A general lower bound

Our main result is that our choice of proxies yield constant-factor lower bounds on the attainable FDR and FNR of any top-K procedure.

Theorem 7. *Given a model $\mathbb{M}$, consider any $(\alpha,\ \beta)$-controlled top-K procedure such that $2\max\{\alpha,\beta\} < \frac{1}{3}$. Then for any scalar $\epsilon \in \big(2\max\{\alpha,\beta\}, \frac{1}{3}\big)$, there exists a constant $c_0(\epsilon) \geq 1$ such that*

$$\alpha \geq c_0^{-1}\mathrm{FDP}_-^*(\beta,\epsilon) \quad \textit{and} \quad \max\big\{\alpha,\ \beta\big\} \geq c_0^{-1}\mathrm{FNP}_-^*(\beta,\epsilon), \qquad \textit{as well as} \tag{4.9a}$$

$$\beta \geq c_0^{-1}\mathrm{FNP}_+^*(\alpha,\epsilon) \quad \textit{and} \quad \max\big\{\alpha,\ \beta\big\} \geq c_0^{-1}\mathrm{FDP}_+^*(\alpha,\epsilon). \tag{4.9b}$$

The slightly unorthodox form of (4.9a) and (4.9b) calls for some discussion. The presence of the maximum reflects the fact that generally only one set of proxies will be suitable for lower bounding α and β simultaneously. If $\alpha > \beta$, the bound in (4.9b) is the meaningful one, while (4.9a) gives the desired bound when $\beta > \alpha$. When $\alpha = \beta$, either equation will do.

The two regimes arise because ℓ_+^* and k_+^* well-approximate the false discovery number and total discovery number only when $\alpha > \beta$, while ℓ_-^* and k_-^* well-approximate these quantities only when $\alpha < \beta$. Intuitively, the dichotomy arises because k_+^* may be *larger* than the actual number of discoveries by an amount as large as order of $\alpha+\beta$, so that FDP_+^* can only be upper-bounded by a quantity of this order, or, equivalently (disregarding constants), a quantity on the order of $\max\{\alpha,\beta\}$. A similar but inverted phenomenon occurs for the − proxies.

Application 1: Independent Gaussians model

In this section, we investigate models in which the vector W has iid Gaussian entries, and the signal structure is specified by either a location shift or a scale factor.

Gaussian location model

We begin by analyzing the Gaussian location model, in which the function f takes the form

$$f(w) = w + \mu \qquad \text{for some } \mu > 0. \tag{4.10}$$

By applying Theorem 7 to this particular model, we obtain the following:

Corollary 6. *Consider the iid Gaussian location model with $m = n^{1-s}$ signals and $\mu = \sqrt{2r \log n}$ with parameters (s, r) satisfying the inequality $0 < s < r < 1$. Suppose that there exists a constant $c > 0$ such that for all $n \geq 1$, there is an $(\alpha_n,\ \beta_n)$-controlled top-K procedure with $\alpha_n = cn^{-\kappa_\alpha}$ and $\beta_n = cn^{-\kappa_\beta}$. Then we must have*

$$\sqrt{s + \kappa_\alpha} + \sqrt{\kappa_\beta} \leq \sqrt{r}, \quad \text{and} \quad \min\left\{\kappa_\alpha,\ \kappa_\beta\right\} \leq \underbrace{\frac{\left(r - s\right)^2}{4r}}_{=:\,\kappa^*}. \tag{4.11}$$

The result encapsulated in (4.11) is essentially the main lower bound of Rabinovich et al. [78], derived by other means and applicable to an exactly Gaussian rather than a Gaussian-like model. Thus, Corollary 6 can be seen as an extension of those results, illustrating how the methods developed in this work can expand the scope of multiple testing lower bounds. Moreover, since the lower bound of Rabinovich et al. [78] is known to be sharp, it seems likely the bound proven here is likewise sharp for the Gaussian location model.

Gaussian scale model

We now turn attention to the Gaussian scale model. It is specified by the transformation

$$f(w) = \sigma w, \qquad \text{for some } \sigma \geq 1.$$

By applying Theorem 7 to this model, we obtain a rather different lower bound on pairs (α, β). At a high-level, the main take-away is that the FDR and FNR can only decay as inverse polynomial functions of n when the signal strength σ is extremely strong—in particular, the scalar σ has to grow polynomially in n.

Corollary 7. *Consider the iid Gaussian scale model with m signals and signal strength $\sigma \geq 1$ where $s_n = \frac{\log \frac{n}{m}}{\log n}$ lies in the interval $[\rho,\ 1 - \rho]$ for some $\rho \in (0, 0.5)$. Suppose that there exists an $(\alpha,\ \beta)$-controlled procedure such that $\max\left\{\alpha,\ \beta\right\} \leq \frac{1}{3c_0}$. Then there exists some $\eta_n \in (0, 1)$ such that*

$$\sigma \geq \frac{1}{\sqrt{2\pi} c_0} \cdot \left(1 - \eta_n\right) \left(\frac{1}{m} + \beta\right)^{-1} \sqrt{2 s_n \log n + 2 \log\left(\alpha + \frac{1}{m}\right)^{-1}}. \tag{4.12}$$

As in the statement of Corollary 6 on the location model, by assuming certain scalings of the number of signals, FDR and FNR, we can give an asymptotic statement. In particular, suppose that the number of signals scales as $m \propto n^{1-s}$ for a fixed s, whereas the FDR and FNR scale as $\alpha_n \propto n^{-\kappa_\alpha}$, and $\beta_n \propto n^{-\kappa_\beta}$ for some scalars $\kappa_\alpha, \kappa_\beta$ such that $\max\{\kappa_\alpha,\ \kappa_\beta\} \leq 1 - s$. Then there is a universal constant $c > 0$ such that

$$\sigma \geq c\, n^{\kappa_\beta} \sqrt{2\big(s + \kappa_\alpha\big) \log n}. \tag{4.13}$$

Consequently, we see that whenever $\kappa_\beta > 0$, the signal strength σ must grow polynomially in n. This is quite a dramatic contrast from the location model, where the analogous quantity μ need only grow proportionally to $\sqrt{\log n}$.

Application 2: Gaussian location models with dependence

Given the presence of dependence in many target applications of multiple testing (e.g., [6]), it makes sense to ask how dependence changes the performance of multiple testing procedures. In this section, we provide answer for two models of Gaussian dependence, which lie at opposite extremes of dependence structures. In both cases, we consider only location shifts.

Spiked dependence model

We begin by considering a doubly spiked covariance model, with one spike within the nulls and a separate spike within the signals. This specification corresponds to coupling all the nulls (and, separately, all the signals) through a single random variable that captures all shared randomness.

In the spiked dependence model, the model distribution $\mathbb{P}_n$ is a multivariate Gaussian $\mathcal{N}\big(0,\ \Sigma\big)$, with covariance matrix of the block-partitioned form

$$\Sigma_{ij} = \begin{cases} 1 & \text{if } i = j, \\ \rho_0 & \text{if } i \neq j \ \text{ and } \ i, j \in \mathcal{H}_0, \\ \rho_1 & \text{if } i \neq j \ \text{ and } \ i, j \notin \mathcal{H}_0, \\ \pm\rho_{\mathrm{c}} & \text{if } i \neq j \ \text{ and } i \in \mathcal{H}_0,\ j \notin \mathcal{H}_0 \ \text{ or vice versa.} \end{cases} \tag{4.14}$$

for some parameters $0 \leq \rho_0, \rho_1, \rho_{\mathrm{c}} < 1$.

Corollary 8. *Consider the spiked dependence model with $m = n^{1-s}$ and $\mu = \sqrt{2r \log n}$ for some pair (s, r) satisfying the inequalities $0 < s < r < 1$. Suppose that for each n, there exists an $(\alpha_n,\ \beta_n)$-controlled top-k procedure with $\alpha_n = cn^{-\kappa_\alpha}$ and $\beta_n = cn^{-\kappa_\beta}$. Then we must have*

$$\sqrt{1 - \rho_0}\sqrt{s + \kappa_\alpha} + \sqrt{1 - \rho_1}\sqrt{\kappa_\beta} \leq \sqrt{r}. \tag{4.15}$$

Note that the bound (4.15) is a generalization of the bound (4.11) for iid Gaussians, to which it reduces when $\rho_0 = \rho_1 = \rho_c = 0$. Relative to this iid case, the bound (4.15) allows for larger values of the pair $(\kappa_\alpha, \kappa_\beta)$—which translates into faster decay rates of FDR and FNR—when either ρ_1 or ρ_0 is non-zero. While this might be counterintuitive at first sight, note that our spiked dependence makes all nulls "more like" each other when $\rho_0 > 0$, and all signals "more like" each other when $\rho_1 > 0$. This similarity in either the nulls or signals means that it becomes easier to control the FDR and FNR. What may still be surprising is that ρ_c does not play any role in the rates.

Grouped dependence model

We now turn to the opposite extreme of dependence models. In the grouped dependence model, we match each signal with a different set of A nulls that are strongly coupled to that signal, but independent of all other signals and all nulls in different groups. More formally, first fix a value $1 \leq A \leq \min\{m, \frac{n}{m}\}$. We then write $\mathcal{H}_1 = \{i_1 < \cdots < i_m\}$ and for each $1 \leq g \leq m$, we define a set of A nulls $\mathcal{H}_0^{(g)}$ corresponding to the g^{th} signal. Finally, we define the independent nulls as $\mathcal{H}_0^{(0)} = \mathcal{H}_0 \setminus \cup_{g=1}^m \mathcal{H}_0^{(g)}$.

Rather than providing an explicit form of the covariance matrix for this model, it is more informative to specify the underlying generative model, given by

$$W_{\mathcal{H}_1} \sim \mathcal{N}\big(0,\ I_m\big), \quad \text{with} \quad W_i \mid W_{\mathcal{H}_1} \sim \begin{cases} \mathcal{N}\big(0,\ 1\big) & \text{if } i \in \mathcal{H}_0^{(0)}, \\ W_{i_g} & \text{if } i \in \mathcal{H}_0^{(g)}. \end{cases} \tag{4.16}$$

By applying Theorem 7 to this model, we obtain the folllowing:

Corollary 9. *Consider the grouped dependence model with $m = n^{1-s}$, $\mu = \sqrt{2r \log n}$, and $A = \frac{n-m}{n^{1-t}}$ where the parameters (s, r, t) satisfy the inequalities $0 < s < r < 1$ and $0 \leq t < s$. Suppose that for some constant $c > 0$ and for each positive integer n, there is an $(\alpha_n,\ \beta_n)$-controlled procedure with $\alpha_n = cn^{-\kappa_\alpha}$ and $\beta_n = cn^{-\kappa_\beta}$. Then*

$$\sqrt{s + \kappa_\alpha} + \sqrt{\kappa_\beta} \leq \sqrt{r}. \tag{4.17}$$

Note that this result is actually rather surprising: the bound (4.17) is identical to our earlier bound (4.11) from the iid case. If the lower bound of (4.17) is sharp, this coincidence reflects a deep fact about the difficulty of multiple testing in the grouped Gaussians model. We do not at this point know, however, whether Corollary 9 is sharp, and the sharpness of this result (like the others established in this chapter) remains an important question for future work.

Application 3: Lehmann alternatives

One model that has at times been used in the multiple testing literature is the so-called Lehmann alternative model for p-values. In this model, the statistics are now p-values; nulls

are then assumed to come from a uniform distribution, while alternatives follow a CDF given by

$$F(p) = p^{\gamma} \qquad \text{for some } \gamma \in (0, 1). \tag{4.18}$$

In order to formulate this problem within our framework, let W_i be iid uniform random variables on the unit interval $[0,\ 1]$, and define the transformation

$$f(w) = 1 - \big(1 - w\big)^{1/\gamma}.$$

Note that here $1 - w$ plays the role of the p-value, so that the w values for signals cluster more around 1 than the nulls.

Corollary 10. *Under the Lehmann alternative model with parameter $\gamma \in (0, 1)$, fix some triple $(\alpha, \beta, \epsilon)$ such that $\alpha \leq \frac{\epsilon}{3}$ and $t := \frac{3\beta}{\epsilon} + \frac{1}{m} + \sqrt{\frac{3c_0\beta}{m\epsilon}} < 1$, where c_0 is the constant from Theorem 7. Further, let $\pi_1 = \frac{m}{n}$. Then for any (α, β)-controlled procedure, we must have*

$$\frac{1}{\gamma} \geq \frac{1-t}{t} \cdot \log\left(\frac{\epsilon}{3\pi_1\alpha}\left[1 + 4\log\frac{3}{\epsilon}\right]^{-1}\right). \tag{4.19}$$

The bound of Corollary 10 requires some interpretation. Intuitively, t is on the order of β, while the argument of the logarithm is on the order of $\frac{1}{\pi_1\alpha}$, so the high-level takeaway is that

$$\frac{1}{\gamma} \gtrsim \frac{1}{\beta}\log\frac{1}{\pi_1\alpha} \tag{4.20}$$

Since the signal – which is to say the difference between nulls and alternatives – becomes greater as γ becomes smaller, $\frac{1}{\gamma}$ is a measure of signal strength, and thus (4.20) is similar to our previous bounds in that it lower bounds the required signal strength in terms of the problem parameters. In this case, the dependence on the FNR β is inverse polynomial, while the dependence on both the FDR α and the sparsity π_1 is logarithmic.

4.4 Proofs

We now turn to the proofs of our results.

Technical tools

Before giving proofs of our main results, we develop two technical tools that we apply repeatedly in our arguments.

Derandomization under concentration

Our proxies depend on the model only through the false discovery number proxies ℓ^*_- and ℓ^*_+. Unfortunately, the dependence is of a rather complicated form, since the definitions (4.6) involve the probabilities of rather complex events specified in terms of the order statistics of nulls and signals. Ideally, we would like to simplify these definitions so as to obtain modified versions that are more tractable. In this section, we show that, provided the model's order statistics admit a suitable concentration bound, we can reduce the probabilistic comparison to a deterministic comparison of expected order statistics. In particular, we make use of the following family of assumptions, which are parameterized by $T \in \{\mathcal{H}_0, \mathcal{H}_1\}$, and an integer k.

Concentration assumption (T, k)**:** There exists a function $\Delta_{T,k} \colon (0,\ 1) \to [0,\ \infty)$ such that

$$\mathbb{P}\left[\left|X_{T,(k)} - \mathbb{E}\big[X_{T,(k)}\big]\right| \geq \Delta_{T,k}\big(\epsilon\big)\right] \leq \epsilon. \tag{4.21}$$

Depending on the nature of the function $\Delta_{T,k}$, condition (4.21) might be a more or less stringent (and a more or less useful) assumption. In general, when we apply this bound, we shall be able to prove it holds for a reasonable choice of $\Delta_{T,k}$.

Our analysis invokes two particular cases of the concentration assumption:

Case I: The concentration assumption (4.21) holds for $(\mathcal{H}_0, \ell^*)$ and $(\mathcal{H}_1, k^* - \ell^* + 1)$.

Case II: The concentration assumption (4.21) holds for $(\mathcal{H}_0, \ell^* + 1)$ and $(\mathcal{H}_1, k^* - \ell^*)$.

Here as always, the integer k^* is one of k^*_- and k^*_+, and the integer ℓ^* is set correspondingly. The following lemma allows us to reduce from probabilities of events to differences in expected order statistics:

Lemma 13. *Suppose that* $\ell^* \in \{\ell^*_-\big(\beta, \epsilon\big), \ell^*_+\big(\alpha, \epsilon\big)\}$ *and define* k^* *accordingly. Then under Cases I and II, we have the following bounds, respectively:*

$$(\mathbf{I}): \quad \mathbb{E}\left[f\big(W_{\mathcal{H}_1,(k^*-\ell^*)}\big)\right] + \Delta_{\mathcal{H}_1,(k^*-\ell^*)}\left(\frac{\epsilon}{3}\right) > \mathbb{E}\left[W_{\mathcal{H}_0,(\ell^*+1)}\right] - \Delta_{\mathcal{H}_0,(\ell^*+1)}\left(\frac{\epsilon}{3}\right). \tag{4.22a}$$

$$(\mathbf{II}): \quad \mathbb{E}\left[f\big(W_{\mathcal{H}_1,(k^*-\ell^*+1)}\big)\right] - \Delta_{\mathcal{H}_1,(k^*-\ell^*+1)}\left(\frac{\epsilon}{3}\right) < \mathbb{E}\left[W_{\mathcal{H}_0,(\ell^*)}\right] + \Delta_{\mathcal{H}_0,(\ell^*)}\left(\frac{\epsilon}{3}\right). \tag{4.22b}$$

See Section 4.6 for the proof of this claim.

Remarks: The main value of Lemma 13 lies in inequality (4.22a). Indeed, this inequality places a lower bound on an expected signal order statistic in terms of an expected null order

statistic (plus some deviations). Since signals are shifted rightward relative to nulls, a lower bound of this kind gives a lower bound on the signal strength in terms of ℓ^*. Meanwhile, Theorem 7 provides upper bounds on ℓ^* (and $m - k^* + \ell^*$) in terms of the realized FDR and FNR. Together, these bounds yield a lower bound on signal strength in terms of FDR and FNR, which can be interpreted as a lower bound on FDR and FNR in terms of the signal strength.

Transfering results between models

It is useful to be able to transfer results from simple models to more complex models that are in some sense "close" to them. In this section, we specify a notion of closeness that makes sense for our problem and prove a technical result that allows us to transfer lower bounds between close models.

Our definition of closeness for models has some unusual features that bear explanation. First, it only applies to models that share a single transformation function f. This limitation is imposed for convenience and is not fundamental. The definition is also asymmetric, with some base model $\mathbb{M}$ given, and the closeness of another model $\mathbb{M}'$ assessed relative to $\mathbb{M}$. The asymmetry arises from the fact that we wish to define proximity of models based on a single fixed order statistic distribution, rather than uniformly over all order statistics of the models, and the single fixed reference point we choose arises from the discovery and false discovery number proxies $k^*_\pm$ and $\ell^*_\pm$, which depend on the model.

Definition 1. *We say that two models* $\mathbb{M} = \big(\mathbb{P}_n,\ f\big)$ *and* $\mathbb{M}' = \big(\mathbb{P}'_n,\ f\big)$ *are* $(\Delta_0, \Delta_1, \ell^*_+, \delta)$*-close if*

$$\max\left\{\mathbb{P}\big(\big|W_{\mathcal{H}_{0,(\ell^*)}} - W'_{\mathcal{H}_{0,(\ell^*)}}\big| \geq \Delta_0\big),\ \mathbb{P}\big(\big|f\big(W_{\mathcal{H}_{1,(k^*-\ell^*)}}\big) - f\big(W'_{\mathcal{H}_{1,(k^*-\ell^*)}}\big)\big| \geq \Delta_1\big)\right\} \leq \delta,$$

where $\ell^* = \ell^*_+(\alpha, \delta, \mathbb{M})$ *and* $k^* = k^*_+(\alpha, \delta)$*. Similarly, we say that they are* $(\Delta_0, \Delta_1, \ell^*_-, \delta)$*-close if the same condition holds with* $\ell^* = \ell^*_-(\beta, \delta, \mathbb{M})$ *and* $k^* = k^*_-(\beta, \delta)$*.*

Based on this definition, we can transfer lower bounds from the base model in Definition 1 to the other model using the following technical lemma.

Lemma 14. *For a given pair* (α, β) *with* $2\max\{\alpha, \beta\} < \frac{1}{3}$*, consider some* $\epsilon \in \big(2\max\{\alpha, \beta\}, \frac{1}{3}\big)$*.*

(a) If $\big(\mathbb{P}'_n,\ f\big)$ *is* $(\Delta_0, \Delta_1, \ell^*_+, \epsilon/3)$*-close to* $\big(\mathbb{P}_n,\ f\big)$*, then*

$$\ell^*_+\big(\alpha,\ \epsilon,\ \big(\mathbb{P}'_n,\ f - \Delta_0 - \Delta_1\big)\big) \geq \ell^*_+(\alpha, \epsilon, (\mathbb{P}_n, f)). \tag{4.23a}$$

(b) If $\big(\mathbb{P}'_n,\ f\big)$ *is* $(\Delta_0, \Delta_1, \ell^*_-, \epsilon/3)$*-close to* $\big(\mathbb{P}_n,\ f\big)$*, then*

$$\ell^*_-\big(\beta,\ \epsilon,\ \big(\mathbb{P}'_n,\ f - \Delta_0 - \Delta_1\big)\big) \geq \ell^*_-(\beta, \epsilon, (\mathbb{P}_n, f)). \tag{4.23b}$$

See Section 4.6 for the proof of this claim.

We use Lemma 14 primarily to remove dependence. In that context, a particularly useful specialization of it is the following decoupling lemma, which allows us to remove dependence between nulls and signals provided that we can verify the concentration condition (4.21).

Lemma 15. *For a given model* $\mathbb{M} = \big(\mathbb{P}_n,\ f\big)$*, let* $\mathbb{M}'$ *denote the same model but with nulls and signals sampled independently from their marginals under* $\mathbb{P}_n$*. Suppose that* $2\max\{\alpha, \beta\} < \frac{1}{3}$*, and that* $\mathbb{M}'$ *satisfies Case I of the concentration assumption* (4.21)*. Then for any* $\epsilon \in \big(2\max\{\alpha, \beta\}, \frac{1}{3}\big)$*, with the integers* $k^* = k^*_+(\alpha, \epsilon)$*,* $\ell^{*\prime} := \ell^*_+(\alpha, \mathbb{M}', \frac{\epsilon}{3})$*, the scalar* $\Delta = 2\big[\Delta_{\mathcal{H}_{0,(\ell^{*\prime})}}\big(\epsilon/6\big) + \Delta_{\mathcal{H}_{1,(k^*-\ell^{*\prime}+1)}}\big(\epsilon/6\big)\big]$*, and the model* $\mathbb{M}'' = \big(\mathbb{P}_n,\ w \mapsto f(w) - 2\Delta\big)$*, we have*

$$\ell^*_+(\alpha, \mathbb{M}'', \epsilon) \geq \ell^*_+(\alpha, \mathbb{M}', \frac{\epsilon}{3}). \tag{4.24a}$$

Similarly, with $k^* = k^*_-(\beta, \epsilon)$*,* $\ell^{*\prime} := \ell^*_-(\beta, \mathbb{M}', \frac{\epsilon}{3})$*, and* $(\Delta, \mathbb{M}'')$ *re-defined accordingly, we have*

$$\ell^*_-(\beta, \mathbb{M}'', \epsilon) \geq \ell^*_-(\beta, \mathbb{M}', \frac{\epsilon}{3}). \tag{4.24b}$$

See Section 4.6 for the proof of this claim.

Proof of Theorem 7

In this section, we prove Theorem 7. Two main ideas underlie the proof. First, we show that any top-K procedure that is (α, β)-controlled must have at least a constant probability of making approximately the right number of discoveries. Here "right" means equal to the true number of signals. In order to formalize this idea, for a given top-K procedure K, define the event

$$\mathcal{E}_{\text{band}} := \left\{ K \in \left[k^*_-(\beta),\ k^*_+(\alpha) \right] \right\}, \tag{4.25}$$

where k^*_- and k^*_+ are the discovery proxies (4.5). The width of this band is determined by α and β and by the constant ϵ, which will play the role of a parameter in the analysis throughout our proofs.

Lemma 16. *For any* $(\alpha,\ \beta)$*-controlled top-*K *procedure, we have* $\mathbb{P}[\mathcal{E}_{\text{band}}] \geq 1 - 2\epsilon$.

We defer the proof of this lemma to Section 4.4.

The main second ingredient is precise version of the argument that led to the inequality (4.8). Essentially, we need to know that the event defined by (4.8) really is the same as

the event defined by the number of false discoveries in the top k being lower-bounded by ℓ. We define

$$L(k) = \left| \{ i : X_i \geq X_{(k)} \} \cap \mathcal{H}_0 \right|, \tag{4.26}$$

corresponding to the number of false discoveries in the top k. In terms of this notation, we have the following:

Lemma 17. *We have*

$$\{L(k) \geq \ell\} = \left\{ W_{\mathcal{H}_0,(\ell)} > f\left(W_{\mathcal{H}_1,(k-\ell+1)}\right) \right\} \qquad \text{for each } k = 1, 2, \ldots, n. \tag{4.27}$$

See Section 4.4 for the proof of this claim.

Equipped with these lemmas, we now turn to the proof of the theorem. Define the events

$$\mathcal{E}_{\text{proxy},-} = \left\{ L(k^*_-(\beta)) \geq \ell^*_-(\beta) \right\} \quad \text{and} \quad \mathcal{E}_{\text{proxy},+} = \left\{ L(k^*_+(\alpha)) \geq \ell^*_+(\alpha) \right\}.$$

By applying Lemma 17 twice, once with the choice $k = k^*_-(\beta)$ and then with the choice $k = k^*_+(\alpha)$, and using the definitions of ℓ^*_- and ℓ^*_+ (see equation (4.6)), we have

$$\min \left\{ \mathbb{P}[\mathcal{E}_{\text{proxy},-}], \ \mathbb{P}[\mathcal{E}_{\text{proxy},+}] \right\} \geq 1 - \epsilon. \tag{4.28}$$

Next, combining Lemma 16 and the bound (4.28) yields

$$\mathbb{P}\Bigg(\overbrace{\mathcal{E}_{\text{band}} \cap \mathcal{E}_{\text{proxy},-}}^{\mathcal{E}_-} \Bigg) \geq 1 - 3\epsilon \ \text{ and } \ \mathbb{P}\Bigg(\overbrace{\mathcal{E}_{\text{band}} \cap \mathcal{E}_{\text{proxy},+}}^{\mathcal{E}_+} \Bigg) \geq 1 - 3\epsilon. \tag{4.29}$$

Argument for $-$ proxies: We now use the bound (4.29) to proof the theorem's claims for the negative-subscript proxies. Note that $L(K) = K \cdot \text{FDP}(K)$, so that on conditioned on $\mathcal{E}_-$, we have

$$\begin{aligned} \ell^*_- \leq K \cdot \text{FDP}(K) &\leq k^*_+ \cdot \text{FDP}(K) \\ \implies \text{FDP}(K) &\geq \frac{\ell^*_-}{k^*_+} = \left(1 - \frac{\alpha}{\epsilon}\right) \cdot \text{FDP}^*_-. \end{aligned}$$

We now take expectations to find that

$$\text{FDR}(K) \geq \mathbb{P}\big[\mathcal{E}_-\big] \cdot \mathbb{E}\big[\text{FDP}(K) \ \mid \ \mathcal{E}_-\big] \overset{(i)}{\geq} \Big(\big(1 - 3\epsilon\big) \left(1 - \frac{\alpha}{\epsilon}\right) \ \text{FDP}^*_-,$$

where step (i) uses the lower bound (4.29).

Recalling that $\mathrm{FDR}(K) \leq \alpha$ by assumption and rearranging the inequality, we find that

$$\mathrm{FDP}_-^* \leq \frac{1}{\big(1 - \frac{\alpha}{\epsilon}\big) \cdot \big(1 - 3\epsilon\big)} \cdot \alpha \leq \frac{2}{1 - 3\epsilon} \cdot \alpha,$$

where we have also used the assumed inequality $\frac{\alpha}{\epsilon} \leq \frac{1}{2}$. This establishes the first inequality in line (4.9a).

We now prove the second inequality in line (4.9a). Observe that the number of non-discovered signals in the top k_-^* statistics can be lower bounded as

$$m \cdot \mathrm{FNP}(k_-^*) = m - \big(k_-^* - L(k_-^*)\big) \;\geq\; m - \big(k_-^* - \ell_-^*\big) \;=\; m \cdot \mathrm{FNP}_-^*.$$

Next note that conditioned on the event $\mathcal{E}_-$, we have

$$\begin{aligned} m \cdot \mathrm{FNP}\big(k_-^*\big) &\leq m \cdot \mathrm{FNP}\big(K\big) + \big(K - k_-^*\big) \\ &\leq m \cdot \mathrm{FNP}\big(K\big) + \big(k_+^* - k_-^*\big) \\ &\leq m \cdot \Big[\mathrm{FNP}\big(K\big) + 2\epsilon^{-1}\big(\alpha + \beta\big)\Big], \end{aligned}$$

where we have used the fact that

$$k_+^* - k_-^* \leq \frac{\epsilon^{-1}\big(\alpha + \beta\big)}{1 - \epsilon^{-1}\alpha} \leq 2\epsilon^{-1}\big(\alpha + \beta\big).$$

Once again taking conditional expectations, dividing through by $\mathbb{P}[\mathcal{E}_-]$, and using the bound on $\mathrm{FNR}(K)$, we find

$$m \cdot \mathrm{FNP}\big(k_-^*\big) \leq m \cdot \left[\frac{\beta}{1 - 3\epsilon} + 2\epsilon^{-1}\big(\alpha + \beta\big)\right]$$

Putting it all together, we conclude that

$$\begin{aligned} \mathrm{FNP}_-^* &\leq \frac{\beta}{1 - 3\epsilon} + 2\epsilon^{-1}\big(\alpha + \beta\big) \\ &\leq \left(\frac{1}{1 - 3\epsilon} + 2\epsilon^{-1}\right) \cdot \big(\alpha + \beta\big) \\ &\leq \left(\frac{2}{1 - 3\epsilon} + 4\epsilon^{-1}\right) \cdot \max\big\{\alpha,\ \beta\big\}. \end{aligned}$$

Argument for + proxies: The argument for the bounds in line (4.9b) based on positive-subscripted proxies is similar, so that we merely sketch it. By reasoning similar to that used for FNR above, we can show that

$$\ell_+^* \leq m \cdot \left[\frac{2}{1 - 3\epsilon} \cdot \alpha + 2\epsilon^{-1}\big(\alpha + \beta\big)\right].$$

Rearranging as before yields the inequality $\text{FDP}_+^* \leq \left(\frac{4}{1-3\epsilon}+4\epsilon^{-1}\right)\cdot\max\left\{\alpha,\ \beta\right\}$. Conditioned on $\mathcal{E}_+$, we have

$$\begin{aligned} m - k_+^* + \ell_+^* \leq m \cdot \text{FNP}(k_+^*) \leq m \cdot \text{FNP}(K) \\ \implies \text{FNP}(K) \geq \text{FNP}_+^*. \end{aligned}$$

Taking conditional expectations and dividing by the probability $\mathbb{P}[\mathcal{E}_+]$, we conclude that

$$\text{FNP}_+^* \geq \frac{1}{1-3\epsilon} \cdot \beta,$$

which completes the proof of the first inequality. The proof of the second inequality is analogous to the negative-subscripted case.

Proof of Lemma 16

It suffices to establish the inequalities $\mathbb{P}[K \geq k_+^*] \leq \epsilon$ and $\mathbb{P}[K \leq k_-^*] \leq \epsilon$. Beginning with the first inequality, note that $\text{FDP}\big(K\big) \geq \frac{K-m}{K} = 1 - \frac{m}{K}$ and that the lower bound is an increasing function of K. Thus, we have the inclusions

$$\begin{aligned} \left\{K \geq k_+^*\right\} &\subset \left\{\text{FDP}\big(K\big) \geq 1 - \frac{m}{k_+^*}\right\} \\ &= \left\{\text{FDP}\big(K\big) \geq \frac{1}{\epsilon}\alpha\right\} \\ &\subset \left\{\text{FDP}\big(K\big) \geq \frac{1}{\epsilon}\text{FDR}\big(K\big)\right\}. \end{aligned}$$

Given this set inclusion, we have

$$\mathbb{P}\big[K \geq k_+^*\big] \leq \mathbb{P}\left[\text{FDP}(K) \geq \frac{\text{FDR}(K)}{\epsilon}\right] \ \leq \ \epsilon,$$

where the final line follows by Markov's inequality.

As for the second inequality, note that $\text{FNP}\big(K\big) \geq \frac{m-K}{m} = 1 - \frac{K}{m}$ and that this lower bound is a decreasing function of K. Thus, we have the inclusions

$$\begin{aligned} \left\{K \leq k_-^*(\beta)\right\} &\subset \left\{\text{FNP}\big(K\big) \geq 1 - \frac{k_-^*}{m}\right\} \\ &= \left\{\text{FNP}\big(K\big) \geq \frac{1}{\epsilon}\beta\right\} \\ &\subset \left\{\text{FNP}\big(K\big) \geq \frac{1}{\epsilon}\text{FNR}\big(K\big)\right\}. \end{aligned}$$

As before, applying Markov's inequality yields the claim.

Proof of Lemma 17

Suppose that $W_{\mathcal{H}_0,(\ell)} > f\big(W_{\mathcal{H}_1,(k-\ell+1)}\big)$, or equivalently, $X_{\mathcal{H}_0,(\ell)} > X_{\mathcal{H}_1,(k-\ell+1)}$. Define the set $I = \big\{i \colon X_i \geq X_{\mathcal{H}_0,(\ell)}\big\}$, and note that if $\big|I\big| \leq k$, then necessarily $X_{\mathcal{H}_0,(\ell)}$ is one of the top k statistics, so that $L(k) \geq \ell$. But, by the hypothesis and the definition of order statistics,

$$\big|I \cap S\big| \leq k - \ell \;\; \text{and} \;\; \big|I \cap \mathcal{H}_0\big| = \ell$$

Thus $\big|I\big| = \big|I \cap S\big| + \big|I \cap \mathcal{H}_0\big| \leq k$, as required.

We now turn to the converse implication. Concretely, fixing some $k \in [n] := \{1, 2, \ldots, n\}$ for which $L(k) \geq \ell$, we prove that $W_{\mathcal{H}_0,(\ell)} > f\big(W_{\mathcal{H}_1,(k-\ell+1)}\big)$. Let $i_k \in [n]$ be the k^{th}-largest rank statistic—that is, the index corresponding to the order statistic $X_{(k)}$. and we break our analysis into two cases, depending on whether $i_k \in \mathcal{H}_1$ or $i \in \mathcal{H}_0$.

Case 1, $i_k \in \mathcal{H}_1$**:** In this case, since there are at most $k - \ell$ signals in the top k statistics, we must have $X_{i_k} \geq X_{\mathcal{H}_1,(k-\ell)}$. On the other hand, for any $j \in \mathcal{H}_0$ such that X_j falls in the top k, we must have $X_j > X_{i_k}$. Since there are at least ℓ such indices, we conclude

$$W_{\mathcal{H}_0,(\ell)} = X_{\mathcal{H}_0,(\ell)} > X_{i_k} \geq X_{\mathcal{H}_1,(k-\ell)} = f\big(W_{\mathcal{H}_1,(k-\ell)}\big) > f\big(W_{\mathcal{H}_1,(k-\ell+1)}\big).$$

Rearranging yields the claim.

Case 2, $i_k \in \mathcal{H}_0$**:** Since the number of signals in the top k is $< k - \ell + 1$, it must be that $X_{\mathcal{H}_1,(k-\ell+1)} < X_{i_k} \leq X_{\mathcal{H}_0,(\ell)}$. Rearranging again gives the claim.

Proof of Corollary 6

Let $\chi_{k,n}$ denote the expected value of the k^{th}-largest value in a sample of n independent standard Gaussians. Recalling the definition (4.21) of the concentration function, classical results on Gaussian order statistics ensure that we apply the concentration assumption with $\Delta_{k,n}(\epsilon) := \sqrt{2 \log \frac{2}{\epsilon}}$. Although this specification is not the sharpest possible, it suffices for our purposes.

Our proof of Corollary 6 is based primarily on comparing Gaussian order statistics to μ. In particular, we wish to establish that the inequality

$$\chi_{\ell^*+1,n-m} + \chi_{m-k^*+\ell^*,m} > \mu - 2\sqrt{2 \log \frac{6}{\epsilon}} \tag{4.30}$$

holds with the choices $(\ell^*,\ k^*) = (\ell^*_-(\beta),\ k^*_-(\beta))$, or $(\ell^*,\ k^*) = (\ell^*_+(\alpha),\ k^*_+(\alpha))$. The proof is identical for these two cases, so we simply use the shorter (ℓ^*, k^*) notation throughout.

Taking inequality (4.30) as given for the moment, we first use it to prove Corollary 6. In order to do so, we require the following:

Lemma 18. *We have*

$$\frac{\sqrt{2\log\frac{n}{k}}}{\xi_{k,n}} = 1 \pm o(1) \quad \textit{for } k = 1, \ldots, 2, \ldots, n, \tag{4.31}$$

where the $o(1)$ *decay holds as* $n \to \infty$ *and/or* $k \to \infty$.

We now proceed with the proof, suppressing n subscripts throughout so as to avoid clutter. Combining inequality (4.30), Lemma 18, and the fact that $m \ll n$, we find that

$$\mu > \chi_{\ell^*+1,n-m} + \chi_{k^*-\ell^*,m} - c_1 \geq \big(1 - o(1)\big)\left[\sqrt{\log\frac{n}{\ell^*+1}} + \sqrt{\log\frac{m}{m-k^*+\ell^*}}\right],$$

where $c_1 > 0$ is a constant that may depend on ϵ. [1]

We now invoke Theorem 7 with $(\ell^*, k^*) = (\ell^*_-, k^*_-)$ or $(\ell^*, k^*) = (\ell^*_+, k^*_+)$, according to whether we are in the regime $\beta \geq \alpha$ or vice versa. Applying the theorem, rearranging, and substituting the value of μ yields

$$\sqrt{\log\frac{n}{c_0\alpha m+1}} + \sqrt{\log\frac{m}{c_0\beta m}} \leq \big(1 - o(1)\big)^{-1}\sqrt{2r\log n} \;= \big(1 + o(1)\big)\sqrt{2r\log n}. \tag{4.32}$$

We claim that it suffices to prove that $\min\big\{\alpha,\ \beta\big\} \cdot m \to \infty$. Indeed, under this scaling, for large enough (n, m), we would ahve

$$\sqrt{2\log\frac{n}{2c_0\alpha m}} + \sqrt{2\log\frac{m}{2c_0\beta m}} \leq (1 + o(1)) \cdot \sqrt{2r\log n}.$$

Substituting the assumed scalings $\alpha = cn^{-\kappa_\alpha}$ and $\beta = cn^{-\kappa_\beta}$ then yields

$$\sqrt{2\big(s+\kappa_\alpha\big)\log n + \log\frac{1}{2c_0}} + \sqrt{2\kappa_\beta\log n + \log\frac{1}{2c_0}} \leq \big(1 + o(1)\big) \cdot \sqrt{2r\log n},$$

and letting $n \to \infty$ yields the claimed inequality $\sqrt{s+\kappa_\alpha} + \sqrt{\kappa_\beta} \leq \sqrt{r}$.

It remains to prove that $\min\{\alpha,\ \beta\} \cdot m \to \infty$, and we split our analysis into two cases.

Case 1: Suppose first that $\alpha \leq \beta$ and assume by way of contradiction that there exists a constant c_2 such that $\alpha m \leq c_2$ for all n. Combined with the inequality (4.32), we find that

$$\sqrt{2\log\frac{n}{c_0c_2+1}} + \sqrt{2\log\frac{m}{c_0\beta m}} \leq \big(1 + o(1)\big) \cdot \sqrt{2r\log n}.$$

Since $r < 1$, this inequality cannot hold once n is large enough, which establishes the desired contradiction.

[1] To clarify a subtle point that we have elided: Lemma 18 requires that $k < \frac{n}{\log n}$, so we need to check $\ell^* + 1 \leq \frac{n}{\log n}$ and $m - k^* + \ell^* \leq \frac{n}{\log n}$. Since $\max\big\{\ell^*,\ m - k^* + \ell^*\big\} \leq \max\big\{\alpha,\ \beta\big\} \cdot m$ by Theorem 7, this condition is in fact easily verified under the given scalings for α and β.

Case 2: Turning to the other case, suppose that $\beta \leq \alpha$, and assume by way of contradiction that there exists a constant c_2 such that $\beta m \leq c_2$ for all n. In this case, again by inequality (4.32), we have

$$\sqrt{2\log\frac{n}{c_0\alpha m+1}}+\sqrt{2\log\frac{m}{c_0c_2}} \leq \big(1+o(1)\big)\cdot\sqrt{2r\log n}.$$

On the other hand, for a suitable choice of c_3, we have

$$\begin{aligned}\sqrt{2\log\frac{n}{c_0\alpha m+1}}+\sqrt{2\log\frac{m}{c_0c_2}} &\geq \sqrt{2\log\frac{n}{4c_0m}}+\sqrt{2\log\frac{m}{c_0c_2}} \\ &\geq \sqrt{2s\log n+\log\frac{1}{c_3}}+\sqrt{2\big(1-s\big)\log n+\log\frac{1}{c_3}}.\end{aligned}$$

Putting together the pieces, we have shown that

$$\sqrt{2s\log n+\log\frac{1}{c_3}}+\sqrt{2\big(1-s\big)\log n+\log\frac{1}{c_3}} \leq \big(1+o(1)\big)\cdot\sqrt{2r\log n}.$$

Since $\sqrt{s}+\sqrt{1-s}>1>\sqrt{r}$, this inequality cannot hold once n is sufficiently large, which establishes the desired contradiction in this case.

Proof of inequality (4.30)

Applying Lemma 13 with $\Delta_{\mathcal{H}_0,k}=\Delta_{\mathcal{H}_1,k}=\sqrt{2\log\frac{2}{\epsilon}}$, we find that

$$\mu+\chi_{k^*-\ell^*,m}+\sqrt{2\log\frac{2}{\epsilon}} > \chi_{\ell^*+1,n-m}-\sqrt{2\log\frac{2}{\epsilon}}.$$

Rearranging yields

$$\mu > \chi_{\ell^*+1,n-m}-\chi_{k^*-\ell^*,m}-2\sqrt{2\log\frac{2}{\epsilon}}.$$

Since the Gaussian distribution is symmetric around zero, we can replace $-\chi_{k^*-\ell^*,m}$ by $\chi_{m-k^*+\ell^*,m}$, which yields the desired inequality.

Proof of Corollary 7

By analogy to the notation in Section 4.4, let $\chi_{k,n}$ denote the expected value of the k^{th}-largest value in a sample of n variables, each of which is the absolute value of a standard Gaussian. Other notational conventions are also preserved. In particular, we let (k^*,ℓ^*) stand in for either $(k^*_-,\ \ell^*_-)$ or $(k^*_+,\ \ell^*_+)$, depending on whether $\beta \geq \alpha$ or vice versa. We also suppress n subscripts.

By applying Lemma 13 in this case, we find that

$$\sigma \cdot \chi_{k^*-\ell^*,m} + \sqrt{2\log\frac{6}{\epsilon}} \geq \chi_{\ell^*+1,n-m} - \sqrt{2\log\frac{6}{\epsilon}},$$

and rearranging yields

$$\sigma \geq \chi^{-1}_{k^*-\ell^*,m}\left[\chi_{\ell^*+1,n-m} + 2\sqrt{2\log\frac{6}{\epsilon}}\right]. \tag{4.33}$$

We now need to evaluate the χ values, and we make use of the following result due to Gordon et al. [43]:

Lemma 19. *For all $k \geq n/2$, we have $\chi_{k,n} \leq \sqrt{2\pi} \cdot \frac{n-k+1}{n+1}$. Moreover, we have*

$$\frac{\sqrt{2\log\frac{n}{k}}}{\chi_{k,n}} = 1 \pm o(1) \qquad \text{for all } k = 1, 2, \ldots, \lfloor \tfrac{n}{\log n} \rfloor, \tag{4.34}$$

where the $o(1)$ scaling holds as n and possibly k go to infinity.

Suppose for now that $\ell^* + 1 \leq \frac{n-m}{\log n - m}$ and that $k^* - \ell^* \geq \frac{m}{2}$. Recall from Theorem 7, applied with the appropriate choice of $+$ or $-$ proxies, that

$$\ell^* \leq c_0 \alpha m \;\text{ and }\; m - k^* + \ell^* \leq c_0 \beta m. \tag{4.35}$$

Consequently, we have the lower bound

$$\begin{aligned}
\chi_{\ell^*+1,n-m} &\geq \big(1 - o(1)\big) \cdot \sqrt{2\log\frac{n-m}{\ell^*+1}} \\
&\geq \big(1 - o(1)\big) \cdot \sqrt{2\log\frac{\frac{n}{m} - 1}{c_0\alpha + 1/m}} \\
&\geq \big(1 - o(1)\big) \cdot \sqrt{2s\log n + 2\log\left(\alpha + \frac{1}{m}\right)^{-1}}
\end{aligned} \tag{4.36a}$$

On the other hand, Lemma 19 also implies that

$$\begin{aligned}
\chi_{k^*-\ell^*,m} &\leq \sqrt{2\pi} \cdot \frac{m - k^* + \ell^* + 1}{m+1} \\
&\leq \sqrt{2\pi} \cdot \left(\frac{m - k^* - \ell^*}{m} + \frac{1}{m}\right) \\
&\leq \sqrt{2\pi} c_0 \cdot \big(\beta + \frac{1}{m}\big).
\end{aligned} \tag{4.36b}$$

Combining the bounds (4.36a) and (4.36b) with inequality (4.33), we find that

$$\sigma \geq \frac{1}{\sqrt{2\pi}c_0} \cdot \big(1 - o(1)\big)\left(\beta + \frac{1}{m}\right)^{-1}\left[\sqrt{2s\log n + 2\log\left(\alpha + \frac{1}{m}\right)^{-1}} + 2\sqrt{2\log\frac{6}{\epsilon}}\right].$$

Since $s \geq \rho > 0$, we have that $\sqrt{2\beta\log n + 2\log\left(\alpha + \frac{1}{m}\right)^{-1}} \gg 2\sqrt{2\log\frac{6}{\epsilon}}$, so that

$$\sigma \geq \frac{1}{\sqrt{2\pi}c_0} \cdot \big(1 - o(1)\big)\left(\beta + \frac{1}{m}\right)^{-1}\sqrt{2s\log n + 2\log\left(\alpha + \frac{1}{m}\right)^{-1}},$$

as claimed.

Under the specified scalings, we have $\alpha + \frac{1}{m} \approx n^{\kappa_\alpha}$ and $\beta + \frac{1}{m} \approx n^{\kappa_\beta}$, which directly implies the comparison

$$\sigma \gtrsim n^{\kappa_\beta}\sqrt{2\big(s + \kappa_\alpha\big)\log n}.$$

We now need to verify that $\ell^* + 1 \leq \frac{n-m}{\log\binom{n-m}{}}$ and $k^* - \ell^* \geq \frac{m}{2}$. From the inequalities (4.35), we deduce that

$$\ell^* + 1 \leq c_0\alpha m + 1 \;\text{ and }\; k^* - \ell^* \geq \big(1 - c_0\beta\big)m.$$

Note that, by the assumption that $s \geq \rho$, we have $m \leq n^{1-\rho}$, we have $\frac{n-m}{\log\binom{n-m}{}} \geq \frac{n}{2\log n}$ for large enough n. On the other hand, we also have $\alpha m + 1 \leq n^{1-\rho} + 1 \leq \frac{n}{2\log n}$ (once n large enough—say for all n such that $n^\rho \geq 4\log n$, for instance). For the second case, recall the assumption $c_0\beta \leq \frac{1}{3} < \frac{1}{2}$, from which the claim follows.

Proof of Corollary 8

At a high level, this proof involves reducing to an independent model with altered variances and using Lipschitz concentration to verify the closeness condition of Lemma 14.

We carry out the reduction in two steps: in Step 1, we reduce to a model with dependence *only* between nulls and signals, and then in Step 2, we reduce to an independent model. The models in Steps 1 and 2 are Gaussian models with covariance matrices Σ' and Σ'', respectively, of the form

$$\Sigma'_{ij} = \begin{cases} 1 - \rho_0 + \rho_{\mathrm{c}} & \text{if } i = j,\ i \in \mathcal{H}_0, \\ 1 - \rho_1 + \rho_{\mathrm{c}} & \text{if } i = j,\ i \notin \mathcal{H}_0, \\ \rho_{\mathrm{c}} & \text{if } i \in \mathcal{H}_0,\ j \notin \mathcal{H}_0, \\ 0 & \text{o.w.} \end{cases} \quad \text{and} \quad \Sigma''_{ij} = \begin{cases} 1 - \rho_0 & \text{if } i = j,\ i \in \mathcal{H}_0, \\ 1 - \rho_1 & \text{if } i = j,\ i \notin \mathcal{H}_0, \\ 0 & \text{o.w.} \end{cases}$$

We let W' and W'' corresponding the corresponding Gaussian random vectors in $\mathbb{R}^n$. The shifts associated with these models are set to be constant-scale perturbations of μ, so that overall, we have the two models

$$\mathbb{M}' = \left(\mathbb{P}'_n,\ \mu + c_1\sqrt{2\log\frac{c_2}{\epsilon}}\right) \quad \text{and} \quad \mathbb{M}'' = \left(\mathbb{P}''_n,\ \mu + 2c_1\sqrt{2\log\frac{c_2}{\epsilon}}\right),$$

where $\mathbb{P}'_n$ and $\mathbb{P}''_n$ are the Gaussian distributions associated with Σ' and Σ''. We also introduce the convenient shorthand notation

$$\ell^{*\prime} = \ell^*\left(\mathbb{M}',\ \frac{\epsilon}{3}\right) \quad \text{and} \quad \ell^{*\prime\prime} = \ell^*\left(\mathbb{M}'',\ \frac{\epsilon}{9}\right).$$

The main idea of the proof is to represent the W variables as functions of higher-dimensional Gaussians. This representation is helpful in decoupling the test statistics from each other. Basically, the constant covariance within the nulls and signals, and across the two, allows us to represent each test statistic as independent of all the others after conditioning on three standard Gaussians that contain all the shared randomness: one each for the within-nulls, within-signals, and between-nulls-and-signals randomness. More precisely, we can write

$$W_i = \begin{cases} \sqrt{1-\rho_0}\cdot U_i + \sqrt{\rho_0-\rho_{\mathrm{c}}}\cdot V_0 + \sqrt{\rho_{\mathrm{c}}}\cdot V_{\mathrm{c}} & \text{if } i \in \mathcal{H}_0, \\ \sqrt{1-\rho_1}\cdot U_i + \sqrt{\rho_1-\rho_{\mathrm{c}}}\cdot V_1 + \sqrt{\rho_{\mathrm{c}}}\cdot V_{\mathrm{c}} & \text{if } i \notin \mathcal{H}_0 \end{cases},$$

$$W'_i = \begin{cases} \sqrt{1-\rho_0}\cdot U'_i + \sqrt{\rho_{\mathrm{c}}}\cdot V'_{\mathrm{c}} & \text{if } i \in \mathcal{H}_0, \\ \sqrt{1-\rho_1}\cdot U'_i + \sqrt{\rho_{\mathrm{c}}}\cdot V'_{\mathrm{c}} & \text{if } i \notin \mathcal{H}_0 \end{cases},$$

$$W''_i = \begin{cases} \sqrt{1-\rho_0}\cdot U''_i & \text{if } i \in \mathcal{H}_0, \\ \sqrt{1-\rho_1}\cdot U''_i & \text{if } i \notin \mathcal{H}_0 \end{cases},$$

The link functions that connect the U and V variables to the order statistics of the W variables are given, in the three cases, by

$$\begin{aligned} a_{0,\ell}\big(u,\ v_{0:1},\ v_c\big) &= \big(\sqrt{1-\rho_0}\cdot u + \sqrt{\rho_0-\rho_{\mathrm{c}}}\cdot v_0 + \sqrt{\rho_{\mathrm{c}}}\cdot v_c\big)_{(\ell)}, \\ b_{0,\ell}\big(u,\ v_c\big) &= \big(\sqrt{1-\rho_0}\cdot u + \sqrt{\rho_{\mathrm{c}}}\cdot v_c\big)_{(\ell)}, \\ c_{0,\ell}\big(u\big) &= \sqrt{1-\rho_0}\cdot u_{(\ell)}, \end{aligned}$$

and similarly for the signals, for which we denote the link functions by a_1, b_1, and c_1.

Our first aim is to prove $\ell^* \geq \ell^{*\prime}$ using Lemma 14. We begin by observing that

$$\Big|a_{0,\ell}\big(u,\ v_{0:1},\ v_{\mathrm{c}}\big) - b_{0,\ell}\big(u,\ v_{\mathrm{c}}\big)\Big| \leq \big|v_0\big|,$$

from which it follows that

$$\begin{aligned}\left\{\left|W_{\mathcal{H}_0,(\ell^{*\prime})}-W'_{\mathcal{H}_0,(\ell^{*\prime})}\right|\leq\Delta\right\}&=\left\{\left|a_{0,\ell^{*\prime}}\big(U,\ V_{0:1},\ V_c\big)-b_{0,\ell^{*\prime}}\big(U',\ V'_c\big)\right|\leq\Delta\right\}\\&\supset\left\{|V_0|\leq\frac{\Delta}{2}\right\}\cap\left\{\left|b_{0,\ell^{*\prime}}\big(U,\ V_c\big)-b_{0,\ell^{*\prime}}\big(U',\ V'_c\big)\right|\leq\frac{\Delta}{2}\right\}\\&\supset\left\{|V_0|\leq\frac{\Delta}{2}\right\}\cap\left\{\left|b_{0,\ell^{*\prime}}\big(U,\ V_c\big)-\mathbb{E}\big[b_{0,\ell^{*\prime}}\big(U,\ V_c\big)\big]\right|\leq\frac{\Delta}{4}\right\}\\&\qquad\cap\left\{\left|b_{0,\ell^{*\prime}}\big(U',\ V'_c\big)-\mathbb{E}\big[b_{0,\ell^{*\prime}}\big(U',\ V'_c\big)\big]\right|\leq\frac{\Delta}{4}\right\}.\end{aligned}$$

A similar analysis yields that the event $\left\{\left|W_{\mathcal{H}_1,(k^{*\prime}-\ell^{*\prime})}-W'_{\mathcal{H}_1,(k^{*\prime}-\ell^{*\prime})}\right|\leq\Delta\right\}$ contains the event

$$\begin{aligned}&\left\{|V_1|\leq\frac{\Delta}{2}\right\}\cap\left\{\left|b_{1,k^{*\prime}-\ell^{*\prime}}\big(U,\ V_c\big)-\mathbb{E}\big[b_{1,k^{*\prime}-\ell^{*\prime}}\big(U,\ V_c\big)\big]\right|\leq\frac{\Delta}{4}\right\}\\&\qquad\cap\left\{\left|b_{1,k^{*\prime}-\ell^{*\prime}}\big(U',\ V'_c\big)-\mathbb{E}\big[b_{1,k^{*\prime}-\ell^{*\prime}}\big(U',\ V'_c\big)\big]\right|\leq\frac{\Delta}{4}\right\}\end{aligned}$$

By Lipschitz concentration, we may choose $\Delta=c'_1\sqrt{2\log\frac{c'_2}{\epsilon}}$ such that

$$\max\left\{\mathbb{P}\Big(\big|W_{\mathcal{H}_0,(\ell^{*\prime})}-W'_{\mathcal{H}_0,(\ell^{*\prime})}\big|>\Delta\Big),\ \mathbb{P}\Big(\big|W_{\mathcal{H}_1,(k^{*\prime}-\ell^{*\prime})}-W'_{\mathcal{H}_1,(k^{*\prime}-\ell^{*\prime})}\big|>\Delta\Big)\right\}\leq\frac{\epsilon}{3}.$$

In other words, $\mathbb{M}$ and $\mathbb{M}'$ are both $\big(\Delta,\ \Delta,\ \ell^*_-,\ \frac{\epsilon}{3}\big)$ and $\big(\Delta,\ \Delta,\ \ell^*_+,\ \frac{\epsilon}{3}\big)$ close. Consequently, if we choose $c_1\geq 2c'_1$ and $c_2\geq c'_2$, then applying Lemma 14 guarantees that $\ell^{*\prime}\geq\ell^*$.

Next observe that $\big|b_{0,\ell}(u,\ v_c)-c_{0,\ell}(u)\big|\leq|v_c|$. We may therefore apply a variant of the previous argument with $\Delta'=c''_1\sqrt{2\log\frac{c''_2}{\epsilon}}$ to show

$$\max\left\{\mathbb{P}\Big(\big|W'_{\mathcal{H}_0,(\ell^{*\prime\prime})}-W''_{\mathcal{H}_0,(\ell^{*\prime\prime})}\big|>\Delta'\Big),\ \mathbb{P}\Big(\big|W'_{\mathcal{H}_1,((k^{*\prime\prime}-\ell^{*\prime\prime})}-W''_{\mathcal{H}_1,(k^{*\prime\prime}-\ell^{*\prime\prime})}\big|>\Delta'\Big)\right\}\leq\frac{\epsilon}{9}.$$

We then find by Lemma 14 that $\ell^{*\prime\prime}\geq\ell^{*\prime}$ provided that $c_1\geq 2c''_1$ and $c_2\geq c''_2$.

Combining the two pieces of our argument, we are guaranteed to have $\ell^*\geq\ell^{*\prime\prime}$ as long $c_1\geq 2\cdot\max\big\{c'_1,\ c''_1\big\}$ and $c_2\geq\max\big\{c'_2,\ c''_2\big\}$. Since we have now reduced to the independent case with μ changed by a constant, applying suitably rescaled version of Theorem 6 yields the conclusion of the corollary.

Proof of Corollary 9

By our previous arguments for order statistics of Gaussians, we know that for the grouped Gaussian model, Case I of the concentration assumption (4.21) holds with

$$\Delta_{\mathcal{H}_0,\ell^*} = \Delta_{\mathcal{H}_1,k^*-\ell^*+1} = c_1\sqrt{2\log\frac{c_2}{\epsilon}}$$

If we choose the constants c_1, c_2 sufficiently large, we can ensure that Case II of the concentration condition (4.21) holds at the same scale in a modified form $\mathbb{M}'$ of the model in which the dependence between nulls and signals is broken and the shift is altered to $\mu - c_1'\sqrt{2\log\frac{c_2'}{\epsilon}}$. We may therefore apply Lemma 15 to obtain that $\ell^* \geq \ell^{*\prime}$, where $\ell^{*\prime}$ is computed in $\mathbb{M}'$.

Since Case II of the concentration condition (4.21) holds in the new model, we may apply Lemma 13. Specifically, if we set $T_0 = \mathcal{H}_0^{(0)}$ and $T_1 = \mathcal{H}_0 \setminus T_0$, then

$$\begin{aligned}
\mu - c_1'\sqrt{2\log\frac{c_2'}{\epsilon}} &\geq \mathbb{E}\big[W'_{\mathcal{H}_0,(\ell^{*\prime}+1)}\big] + \mathbb{E}\big[W'_{\mathcal{H}_1,(m-k^*+\ell^{*\prime}+1)}\big] - 2c_1'\sqrt{2\log\frac{c_2'}{\epsilon}} \\
&\geq \mathbb{E}\big[W'_{T_0,(\ell^{*\prime}+1)}\big] + \mathbb{E}\big[W'_{\mathcal{H}_1,(m-k^*+\ell^{*\prime}+1)}\big] - 2c_1'\sqrt{2\log\frac{c_2'}{\epsilon}} \\
&= \chi_{\ell^{*\prime}+1,|T_0|} + \chi_{m-k^*+\ell^{*\prime}+1,m} - 2c_1'\sqrt{2\log\frac{c_2'}{\epsilon}}.
\end{aligned}$$

Since $t < s$, we have $|T_0| \geq n - (\frac{1}{\gamma}+1)m \geq n - (n^t+1)n^{1-s} \geq n - 2n^{1-t} = \big(1-o(1)\big)n$, and an application of Lemma 18 yields the claim.

Proof of Corollary 10

In order to simplify the proof, it is convenient to pass to an equivalent model. Consider the new random vector $V = (V_1, \ldots, V_n)$ with components $V_i := \log\frac{1}{1-W_i}$. Note that V_i is distributed as a standard exponential and that if we define the transformation function $g(v) = v/\gamma = Av$ with $A\colon = 1/\gamma$, then

$$\log\frac{1}{1-f(W_i)} = g(V_i).$$

With this set-up, the test statistics in the new model are related to the test statistics in the original model by the transformation $x \mapsto \log\frac{1}{1-x}$. Since this transformation is monotonic, any top-K procedure for one can be translated into a top-K procedure for the other, with no change in performance. Likewise, the proxy values are the same for all α and β. In summary, the two models are equivalent for our purposes.

As in previous proofs, we use $(k^*,\ \ell^*)$ as a stand-in for $(k^*_-,\ \ell^*_-)$ or $(k^*_+,\ \ell^*_+)$, and we suppress n subscripts. We claim that it is sufficient to show that

$$\mathbb{P}\left[V_{\mathcal{H}_0,(\ell^*+1)} \le v_-\right] \le \frac{\epsilon}{3}, \tag{4.37a}$$

$$\mathbb{P}\left[g(V_{\mathcal{H}_1,(k^*-\ell^*)}) \ge v_+\right] \le \frac{\epsilon}{3}. \tag{4.37b}$$

where

$$v_+ = \frac{1}{\gamma}\frac{t}{1-t} \ \text{ and } \ v_- = \log\left(\frac{1}{c_0\pi_1\alpha}\left(1+4\log\frac{3}{\epsilon}\right)^{-1}\right).$$

Taking inequalities (4.37a) and (4.37b) as given for the moment, by the definition of ℓ^*, we have

$$\mathbb{P}\left[g(V_{\mathcal{H}_1,(k^*-\ell^*)}) > V_{\mathcal{H}_0,(\ell^*+1)}\right] \ge \epsilon.$$

On the other hand, combining the two bounds above, we see that

$$\mathbb{P}\left(V_{\mathcal{H}_0,(\ell^*+1)} \le v_-\right) + \mathbb{P}\left(g(V_{\mathcal{H}_1,(k^*-\ell^*)}) \ge v_+\right) \le \frac{2\epsilon}{3} < \epsilon.$$

It follows that

$$\mathbb{P}\left[v_+ > g(V_{\mathcal{H}_1,(k^*-\ell^*)}) > V_{\mathcal{H}_0,(\ell^*+1)} > v_-\right] > 0,$$

so $v_- \le v_+$. Rearranging yields

$$\frac{1}{\gamma} \ge \frac{1-t}{t}\cdot\log\left(\frac{1}{c_0\pi_1\alpha}\left(1+4\log\frac{3}{\epsilon}\right)^{-1}\right),$$

as claimed.

The only remaining detail is to prove inequalities (4.37a) and (4.37b).

Proof of inequality (4.37a)

Applying Lemma 4.3 from Boucheron and Thomas [14] yields

$$\mathbb{P}\left[V_{\mathcal{H}_0,(\ell^*+1)} \le \log\frac{n}{\ell^*} - z\right] \le \exp\left(-\frac{\ell^*\left(e^z-1\right)}{4}\right) \qquad \text{for each } z>0.$$

In particular, choosing $z = \log\left(1+\frac{4\log\frac{3}{\epsilon}}{\ell^*}\right) \le \log\left(1+4\log\frac{3}{\epsilon}\right)$, we deduce that

$$\mathbb{P}\left[V_{\mathcal{H}_0,(\ell^*+1)} \le \log\left(\frac{n}{\ell^*}\left(1+4\log\frac{3}{\epsilon}\right)^{-1}\right)\right] \le \epsilon/3.$$

We complete the proof by noting that $\frac{n}{\ell^*} \ge \frac{n}{c_0\alpha m} = \frac{1}{c_0\pi_1\alpha}$.

Proof of inequality (4.37b)

The proof is based on the fact that $V_i = \log \frac{1}{U_i}$ where U_i is a uniform random variable. Let $U_{\mathcal{H}_1,(j)}$ denote the j^{th}-*smallest* value in the sample, which follows a beta distribution with parameters j and $m - j + 1$. We thus have

$$\begin{aligned}\mathbb{E}\big[U_{\mathcal{H}_1,(k^*-\ell^*)}\big] &= \frac{k^* - \ell^*}{m+1}, \qquad \text{and} \\ \operatorname{Var}\big[U_{\mathcal{H}_1,(k^*-\ell^*)}\big] &\leq \mathbb{E}\big[U_{\mathcal{H}_1,(k^*-\ell^*)}\big] \cdot \left(1 - \mathbb{E}\big[U_{\mathcal{H}_1,(k^*-\ell^*)}\big]\right) \cdot \frac{1}{m}.\end{aligned}$$

Applying Chebyshev's inequality yields

$$\mathbb{P}\left(U_{\mathcal{H}_1,(k^*-\ell^*)} \leq \frac{k^* - \ell^*}{m+1} - \sqrt{\frac{3}{\epsilon} \cdot \frac{k^* - \ell^*}{m+1} \cdot \frac{m - k^* + \ell^*}{m+1} \cdot \frac{1}{m}}\right) \leq \epsilon/3.$$

Using the fact that $\frac{m-k+\ell}{m+1} \leq c_0\beta$, we thus have

$$\mathbb{P}\left[U_{\mathcal{H}_1,(k^*-\ell^*)} \leq 1 - c_0\beta - 1/m - \sqrt{c_{\epsilon/3} c_0 \beta / m}\right] = \mathbb{P}\left[U_{\mathcal{H}_1,(k^*-\ell^*)} \leq 1 - t\right] \leq \epsilon/3. \tag{4.38}$$

We now note that

$$\begin{aligned}V_{\mathcal{H}_1,(k^*-\ell^*)} = \log \frac{1}{U_{\mathcal{H}_1,(k^*-\ell^*)}} &= \log \frac{1}{1 - (1 - U_{\mathcal{H}_1,(k^*-\ell^*)})} \\ &= \log\left(1 + \frac{1 - U_{\mathcal{H}_1,(k^*-\ell^*)}}{1 - (1 - U_{\mathcal{H}_1,(k^*-\ell^*)})}\right) \\ &\leq \frac{1 - U_{\mathcal{H}_1,(k^*-\ell^*)}}{1 - (1 - U_{\mathcal{H}_1,(k^*-\ell^*)})}.\end{aligned}$$

Applying the bound (4.38) yields $\mathbb{P}\left[V_{\mathcal{H}_1,(k^*-\ell^*)} \geq \frac{t}{1-t}\right] \leq \epsilon/3$, as claimed. This inequality completes the proof since g is monotonically increasing and $g\big(\frac{t}{1-t}\big) = \frac{1}{\gamma} \cdot \frac{t}{1-t}$.

4.5 Discussion

4.6 Additional proofs

In this section, we collect the proofs of various technical lemmas used in the chapter.

Proof of Lemma 13

The main idea is to pass from probability statements to expectation statements. We prove the forward direction, as the converse admits a similar proof.

Consider the "good" event

$$\mathcal{E} = \left\{W_{\mathcal{H}_0,(\ell^*+1)} \leq f\big(W_{\mathcal{H}_1,(k^*-\ell^*)}\big)\right\},$$

as well as the two "bad" events

$$\mathcal{E}_0 = \left\{W_{\mathcal{H}_0,(\ell^*+1)} < \mathbb{E}\big[W_{\mathcal{H}_0,(\ell^*+1)}\big] - \Delta_{\mathcal{H}_0,\ell^*+1}\big(\frac{\epsilon}{3}\big)\right\}, \qquad \text{and}$$
$$\mathcal{E}_1 = \left\{W_{\mathcal{H}_1,(k^*-\ell^*)} > \mathbb{E}\big[W_{\mathcal{H}_1,(k^*-\ell^*)}\big] + \Delta_{\mathcal{H}_1,k^*-\ell^*}\big(\frac{\epsilon}{3}\big)\right\}.$$

By the maximality of ℓ^*, we have $\mathbb{P}(\mathcal{E}) \geq \epsilon$, while the definition of the concentration functions ensures $\max\big\{\mathbb{P}(\mathcal{E}_0),\ \mathbb{P}(\mathcal{E}_1)\big\} \leq \frac{\epsilon}{3}$. Thus, if we define the event $\mathcal{E}_* = \mathcal{E} \setminus \big(\mathcal{E}_0 \cup \mathcal{E}_1\big)$, we are guaranteed that $\mathbb{P}[\mathcal{E}_*] \geq \frac{\epsilon}{3}$.

Conditioned on $\mathcal{E}_*$, we have

$$\mathbb{E}\big[W_{\mathcal{H}_0,(\ell^*+1)}\big] - \Delta_{\mathcal{H}_0,\ell^*+1}\big(\frac{\epsilon}{3}\big) \leq W_{\mathcal{H}_0,(\ell^*+1)} \leq W_{\mathcal{H}_1,(k^*-\ell^*)} \leq \mathbb{E}\big[W_{\mathcal{H}_1,(k^*-\ell^*)}\big] + \Delta_{\mathcal{H}_1,k^*-\ell^*}\big(\frac{\epsilon}{3}\big).$$

Comparing the left-hand side and right-hand side of this string of inequalities yields the desired conclusion.

Proof of Lemma 14

Again, the proof is mostly a matter of playing with events. First note that the "good" event

$$\mathcal{E} = \left\{W_{\mathcal{H}_0,(\ell^*)} > f(W_{\mathcal{H}_1,(k^*-\ell^*+1)})\right\}$$

satisfies $\mathbb{P}(\mathcal{E}) \geq 1 - \frac{\epsilon}{3}$. The entire point of the proof is to show that the corresponding event for the primed model, namely

$$\mathcal{E}' = \left\{W'_{\mathcal{H}_0,(\ell^*)} > f(W_{\mathcal{H}_1,(k^*-\ell^*+1)}) - \Delta_0 - \Delta_1\right\}$$

satisfies $\mathbb{P}(\mathcal{E}') \geq 1 - \epsilon$. If this claim is proven, the conclusion of the lemma will follow from the maximality of $\ell^{*\prime}$ (cf. equation (4.6)).

In order to establish the latter claim, we observe that $\mathcal{E}' \supset \mathcal{E} \setminus \big(\mathcal{E}_0 \cup \mathcal{E}_1\big)$, where the bad events are defined by

$$\mathcal{E}_0 = \left\{W'_{\mathcal{H}_0,(\ell^*)} < W_{\mathcal{H}_0,(\ell^*)} - \Delta_0\right\},$$
$$\mathcal{E}_1 = \left\{W'_{\mathcal{H}_1,(k^*-\ell^*+1)} < W_{\mathcal{H}_1,(k^*-\ell^*+1)} - \Delta_1\right\}.$$

Given these definitions, the inclusion is clear. Likewise, it's immediate from the assumption of closeness that $\max\big\{\mathbb{P}(\mathcal{E}_0), \mathbb{P}(\mathcal{E}_1)\big\} \leq \frac{\epsilon}{3}$, so that

$$\mathbb{P}\Big(\mathcal{E} \setminus \big(\mathcal{E}_0 \cup \mathcal{E}_1\big)\Big) \geq 1 - \epsilon.$$

Proof of Lemma 15

The proof is similar to that of Lemmas 13 and 14. In this case, the "good" event is given by

$$\mathcal{E}'' = \left\{ W_{\mathcal{H}_0,(\ell^{*\prime})} > f(W_{\mathcal{H}_1,(k^*-\ell^{*\prime}+1)}) - 4\Delta \right\}.$$

If we can show that $\mathbb{P}\big(\mathcal{E}''\big) \geq 1 - \epsilon$, the conclusion of the lemma will follow from the maximality of $\ell^{*\prime}$ and $\ell^{*\prime\prime}$ (see definition (4.6)). Let

$$\mathcal{E}' = \left\{ W_{\mathcal{H}_0,(\ell^{*\prime})} > f(W_{\mathcal{H}_1,(k^*-\ell^{*\prime}+1)}) \right\}$$

denote the corresponding event for the primed model $\mathbb{M}'$. Note that by the definition of $\ell^{*\prime}$, we have $\mathbb{P}\big(\mathcal{E}'\big) \geq 1 - \frac{\epsilon}{3}$.

In order to control $\mathbb{P}[\mathcal{E}'']$, consider as usual the "bad" events

$$\begin{aligned}
\mathcal{E}_0 &= \left\{ \left| W_{\mathcal{H}_0,(\ell^{*\prime})} - W'_{\mathcal{H}_1,(\ell^{*\prime})} \right| \geq 2\Delta_{\mathcal{H}_0,\ell^{*\prime}}\big(\epsilon/6\big) \right\}, \\
\mathcal{E}_1 &= \left\{ \left| W_{\mathcal{H}_1,(k^*-\ell^{*\prime}+1)} - W_{\mathcal{H}_1,(k^*-\ell^{*\prime}+1)} \right| \geq 2\Delta_{\mathcal{H}_1,k^*-\ell^{*\prime}+1} \right\}.
\end{aligned}$$

By two applications of Case I of the concentration condition (4.21), we find that

$$\max\left\{ \mathbb{P}\big(\mathcal{E}_0\big),\ \mathbb{P}\big(\mathcal{E}_1\big) \right\} \leq \frac{\epsilon}{3}.$$

Given the set inclusion $\mathcal{E}'' \supset \mathcal{E}' \setminus \big(\mathcal{E}_0 \cup \mathcal{E}_1\big)$, we conclude that

$$\mathbb{P}\big(\mathcal{E}''\big) \geq 1 - \frac{\epsilon}{3} - \frac{2\epsilon}{3} = 1 - \epsilon,$$

as claimed.

Chapter 5

Conclusion

In the preceding chapters, we examined two statistical techniques with wide applications and a deep theoretical literature: posterior simulation via Markov chain Monte Carlo (MCMC) and multiple hypothesis testing under the false discovery rate (FDR) metric. In both cases, we targeted questions relevant to practice and provided theoretical analyses shedding light on them.

Within the domain of MCMC, we examined how for the purposes of estimating posterior moments, convergence can occur much faster than predicted by standard ways of analyzing mixing. In Chapter 2, we introduced the f-mixing time of a function, and showed that it can be characterized by the interaction between the function and the eigenspaces of the transition operator. Using these tools, we proved that the empirical averages of a function f concentrate around their equilibrium values at a rate characterized by the f-mixing time; in so doing, we eliminated the worst-case dependence on the spectral gap of the chain, characteristic of previous results. Our methodology yields sharper confidence intervals, and better rates for sequential hypothesis tests, and we have provided evidence that our theory's predictions are accurate in some real instances of MCMC and therefore of potential practical interest.

Our investigation also suggests several further questions, notably concerning the continuous and non-reversible cases. Both arise frequently in statistical applications—for example, when sampling continuous parameters or when performing Gibbs sampling with systematic scan. As uniform Hoeffding bounds do exist for the continuous case and, more recently, have been established for the non-reversible case, we believe many of our conclusions should carry over to these settings.

Furthermore, it would be desirable to have methods for estimating or bounding the f-mixing time based on samples. Likewise, while we have shown what can be done with spectral methods, the classical theory provides a much larger arsenal of techniques, some of which may generalize to yield sharper f-mixing time bounds. We leave these and other problems to future work.

Within the domain of multiple hypothesis testing, we investigated the nonasymptotic trade-off between controlling false positives (via the false discovery rate, or FDR) and con-

trolling false negatives (via the false non-discovery rate, or FNR). In Chapter 3, we explored this issue in the context of the sparse generalized Gaussian model, deriving both the optimal tradeoff between FDR and FNR and the first nonasymptotic lower bounds on the sum of the FDR and FNR. We complemented these lower bounds by establishing the nonasymptotic minimaxity of the Benjamini-Hochberg (BH) multiple testing algorithm. For the more recent Barber-Candès (BC) algorithm we were able to show minimaxity in certain regimes and near-minimaxity across the board. Our theoretical predictions are validated using simple simulations, and the analytical results include recent asymptotic results [4] as special cases. Our work introduces a simple proof strategy based on a reduction to deterministic and data-oblivious procedures.

In Chapter 4, we elaborated this derandomization-based proof strategy into a framework that could handle a much larger class of multiple testing models. We showed that the existence of a multiple testing algorithm with a given level of Type I and Type II errors implies constraints on the relationship between the distributions of the order statistics of the null and signal test statistics. Using this fact, we derived a meta-theorem providing an avenue to proving concrete lower bounds for a wide range of multiple testing scenarios. We then instantiated this meta-theorem for the Gaussian location-shift model with and without dependence, the Gaussian scale model, and the Lehmann alternatives model of p-values.

A number of future directions for minimax theory for multiple testing present themselves. The results of Chapter 4 might be seen as a first step toward the kind of general minimax theory available in estimation. There, techniques like Le Cam's method and the Fano method have provided general avenues to proving lower bounds for a wide array of statistical problems. Yet nothing of the kind exists in the multiple testing world. We believe it would be highly desirable to develop general principles that can be used to prove lower bounds in multiple testing – either by exporting idea from estimation theory or otherwise.

Progress might be made by expanding the applications of the meta-theorem of Chapter 4. The modern multiple testing literature contains a wide variety of testing scenarios that we have not yet explored. For example, much of that literature focuses on developing valid FDR control procedures that can gain power or precision by explicitly using prior knowledge and structure in various ways, including using null-proportion adaptivity [88, 89], grouping of hypotheses [6, 53], prior or penalty weights [9, 39], or other forms of structure [66, 80].

Similarly, the issue of dependence—either positive or arbitrary—between test statistics has been an area of focus [10, 12, 80]. (Dependence has already been explored for the higher criticism statistic applied to the detection problem [48, 58, 49].) Still others have studied the nonexchangeability of hypotheses, either in the context of multiple scales of signal strength, or in the context of online FDR procedures [32, 55]. We have gone part of the way to addressing these kinds of scenarios, analyzing simple forms of dependence, but the problem of establishing lower bounds under more flexible models of dependence remains open. All the more so in the non-exchangeable case, which may require a fundamentally different approach.

It is also far from clear that known procedures are optimal under assumptions of structure, dependence, or various kinds of non-exchangeability. Unlike Chapter 3, which provided matching lower and upper bounds, the more general setting of Chapter 4 so far includes only

lower bounds. Providing upper bounds in a more general setting is thus necessary – and may even yield improved algorithms. Chen and Arias-Castro [16] have made progress in this direction by providing *upper* bounds for existing procedures for the online FDR problem [55], but much still remains unknown.

Bibliography

[1] F. Abramovich, Y. Benjamini, D. Donoho, and I.M. Johnstone. Adapting to unknown sparsity by controlling the false discovery rate. *Annals of Statistics*, 34(2):584–653, 2006.

[2] David Aldous and Persi Diaconis. Shuffling cards and stopping times. *American Mathematical Monthly*, 93(5):333–348, 1986.

[3] Joshua D. Angrist and Alan B. Krueger. Does compulsory school attendance affect schooling and earnings? *The Quarterly Journal of Economics*, 106(4):979–1014, 1991. ISSN 00335533, 15314650. URL `http://www.jstor.org/stable/2937954`.

[4] E. Arias-Castro and S. Chen. Distribution-free multiple testing. *Electronic Journal of Statistics*, 11(1):1983–2001, 2017.

[5] ATLAS Collaboration. Observation of a new particle in the search for the standard model higgs boson with the atlas detector at the lhc. *Physics Letters B*, 716(1):1 – 29, 2012. ISSN 0370-2693. doi: https://doi.org/10.1016/j.physletb.2012.08.020. URL `http://www.sciencedirect.com/science/article/pii/S037026931200857X`.

[6] Rina Barber and Aaditya Ramdas. The p-filter: Multilayer false discovery rate control for grouped hypotheses. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79, 11 2016. doi: 10.1111/rssb.12218.

[7] Thomas R Belin and Donald B Rubin. The analysis of repeated-measures data on schizophrenic reaction times using mixture models. *Statistics in Medicine*, 14(8):747–768, 1995.

[8] Y. Benjamini and Y. Hochberg. Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 57(1):289–300, 1995.

[9] Y. Benjamini and Y. Hochberg. Multiple hypotheses testing with weights. *Scandinavian Journal of Statistics*, 24(3):407–418, 1997.

[10] Y. Benjamini and D. Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of Statistics*, 29(4):1165–1188, 2001.

[11] Scott M Berry, Bradley P Carlin, J Jack Lee, and Peter Muller. *Bayesian Adaptive Methods for Clinical Trials*. CRC press, 2010.

[12] G. Blanchard and E. Roquain. Two simple sufficient conditions for FDR control. *Electronic Journal of Statistics*, 2:963–992, 2008.

[13] M. Bogdan, A. Chakrabarti, F. Frommlet, and J.K. Ghosh. Asymptotic Bayes-optimality under sparsity of some multiple testing procedures. *Annals of Statistics*, 39(3):1551–1579, 2011.

[14] Stéphane Boucheron and Maud Thomas. Concentration inequalities for order statistics. *Electronic Communications in Probability*, 17:12 pp., 2012.

[15] David Card and Alan B Krueger. Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania. *American Economic Review*, 84(4):772–793, September 1994. URL `https://ideas.repec.org/a/aea/aecrev/v84y1994i4p772-93.html`.

[16] Shiyun Chen and Ery Arias-Castro. Sequential multiple testing. *arXiv preprint arXiv:1705.10190*, 2017.

[17] Raj Chetty, John N. Friedman, and Jonah E. Rockoff. Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9): 2593–2632, September 2014. doi: 10.1257/aer.104.9.2593. URL `http://www.aeaweb.org/articles?id=10.1257/aer.104.9.2593`.

[18] Z. Chi. On the performance of FDR control: Constraints and a partial solution. *Annals of Statistics*, 35(4):1409–1431, 2007.

[19] Hee Min Choi and James P. Hobert. The polya-gamma gibbs sampler for bayesian logistic regression is uniformly ergodic. *Electronic Journal of Statistics*, 7:2054–2064, 2013.

[20] Kai-Min Chung, Henry Lam, Zhenming Liu, and Michael Mitzenmacher. Chernoff-hoeffding bounds for markov chains: Generalized and simplified. In *29th International Symposium on Theoretical Aspects of Computer Science, STACS 2012*, pages 124–135, 2012.

[21] Andrea Cipriani, Toshi A Furukawa, Georgia Salanti, Anna Chaimani, Lauren Z Atkinson, Yusuke Ogawa, Stefan Leucht, Henricus G Ruhe, Erick H Turner, Julian P T Higgins, Matthias Egger, Nozomi Takeshima, Yu Hayasaka, Hissei Imai, Kiyomi Shinohara, Aran Tajika, John P A Ioannidis, and John R Geddes. Comparative efficacy and acceptability of 21 antidepressant drugs for the acute treatment of adults with major depressive disorder: a systematic review and network meta-analysis. *The Lancet*, 391(10128):1357 – 1366, 2018. ISSN 0140-6736. doi: https://doi.org/10.1016/S0140-6736(17)32802-7. URL `http://www.sciencedirect.com/science/article/pii/S0140673617328027`.

[22] Mark Conger and Divakar Viswanath. Riffle shuffles of decks with repeated cards. *The Annals of Probability*, 34(2):804–819, 2006.

[23] Alexander D'Amour, Peng Ding, Avi Feller, Lihua Lei, and Jasjeet Sekhon. Overlap in observational studies with high-dimensional covariates. *arXiv:1711.02582*, 2017.

[24] Persi Diaconis and James Allen Fill. Strong stationary times via a new form of duality. *The Annals of Probability*, 18(4):1483–1522, 10 1990.

[25] Persi Diaconis and Bob Hough. Random walk on unipotent matrix groups. *arXiv preprint arXiv:1512.06304*, 2015.

[26] Rebecca Diamond, Tim McQuade, and Franklin Qian. The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. *American Economic Review*, 109(9):3365–94, September 2019. doi: 10.1257/aer.20181289. URL `http://www.aeaweb.org/articles?id=10.1257/aer.20181289`.

[27] D. Donoho and J. Jin. Higher criticism for detecting sparse heterogeneous mixtures. *Annals of Statistics*, 32(3):962–994, 2004.

[28] D. Donoho and J. Jin. Asymptotic minimaxity of false discovery rate thresholding for sparse exponential data. *Annals of Statistics*, 34(6):2980–3018, 2006.

[29] D. Donoho and J. Jin. Higher criticism for large-scale inference, especially for rare and weak effects. *Statistical Science*, 30(1):1–25, 2015.

[30] Bradley Efron. *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press, 2012.

[31] James M Flegal, Murali Haran, and Galin L Jones. Markov chain monte carlo: Can we trust the third significant figure? *Statistical Science*, 23(2):250–260, 2008.

[32] D.P. Foster and R.A. Stine. α-investing: A procedure for sequential control of expected false discoveries. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(2):429–444, 2008.

[33] R. Foygel Barber and E.J. Candès. Controlling the false discovery rate via knockoffs. *Annals of Statistics*, 43(5):2055–2085, 2015.

[34] Andrew Gelman and John Carlin. Beyond power calculations: Assessing type s (sign) and type m (magnitude) errors. *Perspectives on Psychological Science*, 9(6): 641–651, 2014. doi: 10.1177/1745691614551642. URL `https://doi.org/10.1177/1745691614551642`. PMID: 26186114.

[35] Andrew Gelman and Gary King. Estimating incumbency advantage without bias. *American Journal of Political Science*, 34(4):1142–1164, 1990. ISSN 00925853, 15405907. URL `http://www.jstor.org/stable/2111475`.

[36] Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. *Statist. Sci.*, 7(4):457–472, 11 1992. doi: 10.1214/ss/1177011136. URL `https://doi.org/10.1214/ss/1177011136`.

[37] Andrew Gelman, John B. Carlin, Hal S. Stern, David B. Dunson, Aki Vehtari, and Donald B. Rubin. *Bayesian Data Analysis, Third Edition*. Chapman and Hall/CRC, 2013.

[38] C.R. Genovese and L. Wasserman. Operating characteristics and extensions of the false discovery rate procedure. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):499–517, 2002.

[39] C.R. Genovese, K. Roeder, and L. Wasserman. False discovery control with p-value weighting. *Biometrika*, 93(3):509–524, 2006.

[40] Zoubin Ghahramani and Thomas L Griffiths. Infinite latent feature models and the indian buffet process. In *Advances in Neural Information Processing Systems 18: Annual Conference on Neural Information Processing Systems, NIPS 2005*, pages 475–482, 2005.

[41] David Gillman. A chernoff bound for random walks on expander graphs. *SIAM Journal on Computing*, 27(4):1203–1220, 1998.

[42] Peter W Glynn and Eunji Lim. Asymptotic validity of batch means steady-state confidence intervals. In *Advancing the Frontiers of Simulation*, pages 87–104. Springer, 2009.

[43] Y. Gordon, A.E. Litvak, S. Mendelson, and A. Pajor. Gaussian averages of interpolated bodies and applications to approximate reconstruction. *Journal of Approximation Theory*, 149(1):59 – 73, 2007. ISSN 0021-9045.

[44] Thomas L. Griffiths and Mark Steyvers. Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(suppl 1):5228–5235, 2004.

[45] Benjamin M Gyori and Daniel Paulin. Non-asymptotic confidence intervals for mcmc in practice. *arXiv preprint arXiv:1212.2016*, 2012.

[46] Benjamin M. Gyori and Daniel Paulin. Hypothesis testing for markov chain monte carlo. *Statistics and Computing*, 26(6):1–12, July 2015. ISSN 0960-3174.

[47] Jens Hainmueller and Dominik Hangartner. Who gets a swiss passport? a natural experiment in immigrant discrimination. *American Political Science Review*, 107(1): 159–187, 2013. doi: 10.1017/S0003055412000494.

[48] P. Hall and J. Jin. Properties of higher criticism under strong dependence. *Annals of Statistics*, 36(1):381–402, 2008.

[49] P. Hall and J. Jin. Innovated higher criticism for detecting sparse signals in correlated noise. *Annals of Statistics*, 38(3):1686–1732, 2010.

[50] Masahito Hayashi and Shun Watanabe. Information geometry approach to parameter estimation in markov chains. *The Annals of Statistics*, 44(4):1495–1535, 08 2016.

[51] Daniel Hsu, Aryeh Kontorovich, David A. Levin, Yuval Peres, and Csaba Szepesvári. Mixing time estimation in reversible markov chains from a single sample path. *arXiv preprint arXiv:1708.07367*, 2017.

[52] Daniel J Hsu, Aryeh Kontorovich, and Csaba Szepesvári. Mixing time estimation in reversible markov chains from a single sample path. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems, NIPS 2015*, pages 1459–1467, 2015.

[53] J. Hu, H. Zhao, and H. Zhou. False discovery rate control with groups. *Journal of the American Statistical Association*, 105(491):1215–1227, 2010.

[54] Sonia Jain, Radford M Neal, et al. Splitting and merging components of a nonconjugate dirichlet process mixture model. *Bayesian Analysis*, 2(3):445–472, 2007.

[55] Adel Javanmard and Andrea Montanari. Online rules for control of false discovery rate and false discovery exceedance. *Ann. Statist.*, 46(2):526–554, 04 2018. doi: 10.1214/17-AOS1559. URL `https://doi.org/10.1214/17-AOS1559`.

[56] P. Ji and J. Jin. UPS delivers optimal phase diagram in high-dimensional variable selection. *Annals of Statistics*, 40(1):73–103, 2012.

[57] P. Ji and Z. Zhao. Rate optimal multiple testing procedure in high-dimensional regression. *arXiv preprint arXiv:1404.2961*, 2014.

[58] J. Jin and T. Ke. Rare and weak effects in large-scale inference: Methods and phase diagrams. *arXiv preprint arXiv:1410.4578*, 2014.

[59] Galin L Jones and James P Hobert. Honest exploration of intractable probability distributions via markov chain monte carlo. *Statistical Science*, 16(4):312–334, 2001.

[60] Aldéric Joulin, Yann Ollivier, et al. Curvature, concentration and error estimates for markov chain monte carlo. *The Annals of Probability*, 38(6):2418–2442, 2010.

[61] Aryeh Kontorovich, Roi Weiss, et al. Uniform chernoff and dvoretzky-kiefer-wolfowitz-type inequalities for markov chains and related processes. *Journal of Applied Probability*, 51(4):1100–1113, 2014.

[62] Carlos A. Léon and François Perron. Optimal hoeffding bounds for discrete reversible markov chains. *The Annals of Applied Probability*, 14(2):958–970, 05 2004. doi: 10.1214/105051604000000170.

[63] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2008.

[64] Pascal Lezaud. Chernoff-type bound for finite markov chains. *The Annals of Applied Probability*, 8(3):849–867, 1998.

[65] Pascal Lezaud. Chernoff and berry–esséen inequalities for markov processes. *ESAIM: Probability and Statistics*, 5:183–201, 2001.

[66] A. Li and R. Foygel Barber. Multiple testing with the structure adaptive Benjamini-Hochberg algorithm. *arXiv preprint arXiv:1606.07926*, 2016.

[67] Ang Li and Rina Foygel Barber. Multiple testing with the structure-adaptive benjamini–hochberg algorithm. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 81(1):45–74, 2019.

[68] Sean P Meyn and Richard L Tweedie. *Markov Chains and Stochastic Stability*. Springer Science & Business Media, 2012.

[69] David M. Mimno, Matthew D. Hoffman, and David M. Blei. Sparse stochastic inference for latent dirichlet allocation. In *Proceedings of the 29th International Conference on Machine Learning, ICML 2012*, 2012.

[70] Radford M Neal. Markov chain sampling methods for dirichlet process mixture models. *Journal of Computational and Graphical Statistics*, 9(2):249–265, 2000.

[71] Pierre Neuvial and Etienne Roquain. On false discovery rate thresholding for classification under sparsity. *Annals of Statistics*, 40(5):2572–2600, 2012.

[72] Brian A. Nosek, Charles R. Ebersole, Alexander C. DeHaven, and David T. Mellor. The preregistration revolution. *Proceedings of the National Academy of Sciences*, 115 (11):2600–2606, 2018. ISSN 0027-8424. doi: 10.1073/pnas.1708274114. URL `https://www.pnas.org/content/115/11/2600`.

[73] Yann Ollivier. Ricci curvature of markov chains on metric spaces. *Journal of Functional Analysis*, 256(3):810–864, 2009.

[74] Open Science Collaboration. Estimating the reproducibility of psychological science. *Science*, 349(6251), 2015. ISSN 0036-8075. doi: 10.1126/science.aac4716. URL `https://science.sciencemag.org/content/349/6251/aac4716`.

[75] Daniel Paulin. Concentration inequalities for markov chains by marton couplings and spectral methods. *arXiv preprint arXiv:1212.2015*, 2012.

[76] Benjamin M. Peter, Emilia Huerta-Sanchez, and Rasmus Nielsen. Distinguishing between selective sweeps from standing variation and from a de novo mutation. *PLOS Genetics*, 8(10):1–14, 10 2012. doi: 10.1371/journal.pgen.1003011. URL `https://doi.org/10.1371/journal.pgen.1003011`.

[77] Maxim Rabinovich, Aaditya Ramdas, Michael I. Jordan, and Martin J. Wainwright. Function-specific mixing times and concentration away from equilibrium, 2019. URL `https://doi.org/10.1214/19-BA1151`. Advance publication.

[78] Maxim Rabinovich, Aaditya Ramdas, Michael I. Jordan, and Martin J. Wainwright. Optimal rates and tradeoffs in multiple hypothesis testing. *Statistica Sinica*, 30:1–22, 2020.

[79] Fernando Racimo, Jeremy J. Berg, and Joseph K. Pickrell. Detecting polygenic adaptation in admixture graphs. *Genetics*, 208(4):1565–1584, 2018. ISSN 0016-6731. doi: 10.1534/genetics.117.300489. URL `https://www.genetics.org/content/208/4/1565`.

[80] A. Ramdas, R. Foygel Barber, M.J. Wainwright, and M.I. Jordan. A unified treatment of multiple testing with prior knowledge. *arXiv preprint arXiv:1703.06222*, 2017.

[81] Aaditya Ramdas, Tijana Zrnic, Martin Wainwright, and Michael Jordan. Saffron: an adaptive algorithm for online control of the false discovery rate. In *International Conference on Machine Learning*, pages 4283–4291, 2018.

[82] Christian P. Robert and George Casella. *Monte Carlo Statistical Methods (Springer Texts in Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2005. ISBN 0387212396.

[83] Jorge Carlos Román and James P. Hobert. Geometric ergodicity of Gibbs samplers for Bayesian general linear mixed models with proper priors. *Linear Algebra and its Applications*, 473:54 – 77, 2015. ISSN 0024-3795. Special issue on Statistics.

[84] Paul-Marie Samson et al. Concentration of measure inequalities for Markov chains and ϕ-mixing processes. *The Annals of Probability*, 28(1):416–461, 2000.

[85] Sebastian Schneeweiss, Jeremy Rassen, Robert Glynn, Jerry Avorn, Helen Mogun, and M Brookhart. High-dimensional propensity score adjustment in studies of treatment effects using health care claims data. *Epidemiology (Cambridge, Mass.)*, 20:512–22, 08 2009. doi: 10.1097/EDE.0b013e3181a663cc.

[86] Alistair Sinclair. Improved bounds for mixing rates of markov chains and multicommodity flow. *Combinatorics, Probability, and Computing*, 1(4):351–370, 1992.

[87] Aaron J. Stern, Peter R. Wilton, and Rasmus Nielsen. An approximate full-likelihood method for inferring selection and allele frequency trajectories from dna sequence data.

PLOS Genetics, 15(9):1–32, 09 2019. doi: 10.1371/journal.pgen.1008384. URL https://doi.org/10.1371/journal.pgen.1008384.

[88] J. Storey. A direct approach to false discovery rates. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 64(3):479–498, 2002.

[89] J. Storey, J. Taylor, and D. Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: A unified approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(1):187–205, 2004.

[90] Alexandre B. Tsybakov. *Introduction to Nonparametric Estimation*. Springer Publishing Company, Incorporated, 1st edition, 2008.

[91] J.W. Tukey. *The Problem of Multiple Comparisons: Introduction and Parts A, B, and C*. Princeton University, 1953.

[92] Aki Vehtari, Andrew Gelman, Daniel Simpson, Bob Carpenter, and Paul-Christian Bürkner. Rank-normalization, folding, and localization: An improved $\hat{R}$ for assessing convergence of mcmc. *arXiv:1903.08008*, 2019.

[93] Stefan Wager, Wenfei Du, Jonathan Taylor, and Robert J. Tibshirani. High-dimensional regression adjustments in randomized experiments. *Proceedings of the National Academy of Sciences*, 113(45):12673–12678, 2016. ISSN 0027-8424. doi: 10.1073/pnas.1614732113. URL https://www.pnas.org/content/113/45/12673.

[94] Ronald L. Wasserstein and Nicole A. Lazar. The asa statement on p-values: Context, process, and purpose. *The American Statistician*, 70(2):129–133, 2016. doi: 10.1080/00031305.2016.1154108. URL https://doi.org/10.1080/00031305.2016.1154108.

[95] Shun Watanabe and Masahito Hayashi. Finite-length analysis on tail probability for markov chain and application to simple hypothesis testing. *The Annals of Applied Probability*, 27(2):811–845, 04 2017. doi: 10.1214/16-AAP1216.