
The Power of Adaptivity in Identifying Statistical Alternatives

Kevin Jamieson, Daniel Haas, Ben Recht
University of California, Berkeley
Berkeley, CA 94720
{kjamieson,dhaas,brecht}@eecs.berkeley.edu

Abstract

This paper studies the trade-off between two different kinds of pure exploration: breadth versus depth. We focus on the most biased coin problem, asking how many total coin flips are required to identify a “heavy” coin from an infinite bag containing both “heavy” coins with mean $\theta_1 \in (0, 1)$, and “light” coins with mean $\theta_0 \in (0, \theta_1)$, where heavy coins are drawn from the bag with proportion $\alpha \in (0, 1/2)$. When $\alpha, \theta_0, \theta_1$ are unknown, the key difficulty of this problem lies in distinguishing whether the two kinds of coins have very similar means, or whether heavy coins are just extremely rare. While existing solutions to this problem require some prior knowledge of the parameters $\theta_0, \theta_1, \alpha$, we propose an adaptive algorithm that requires no such knowledge yet still obtains near-optimal sample complexity guarantees. In contrast, we provide a lower bound showing that non-adaptive strategies require at least quadratically more samples. In characterizing this gap between adaptive and nonadaptive strategies, we make connections to anomaly detection and prove lower bounds on the sample complexity of differentiating between a single parametric distribution and a mixture of two such distributions.

1 Introduction

The trade-off between exploration and exploitation has been an ever-present trope in the online learning literature. In contrast, this paper studies the trade-off between two different kinds of pure exploration: breadth versus depth. Consider a bag that contains an infinite number of two kinds of biased coins: “heavy” coins with mean $\theta_1 \in (0, 1)$ and “light” coins with mean $\theta_0 \in (0, \theta_1)$. When a player picks a coin from the bag, with probability α the coin is “heavy” and with probability $(1 - \alpha)$ the coin is “light.” The player can flip any coin she picks from the bag as many times as she wants, and the goal is to identify a heavy coin using as few total flips as possible. When $\alpha, \theta_0, \theta_1$ are unknown, the key difficulty of this problem lies in distinguishing whether the two kinds of coins have very similar means, or whether heavy coins are just extremely rare. That is, how does one balance flipping an individual coin many times to better estimate its mean against considering many new coins to maximize the probability of observing a heavy one. Previous work has only proposed solutions that rely on some or full knowledge $\alpha, \theta_0, \theta_1$, limiting their applicability. In this work we propose the first algorithm that requires no knowledge of $\alpha, \theta_0, \theta_1$, is guaranteed to return a heavy coin with probability at least $1 - \delta$, and flips a total number of coins, in expectation, that nearly matches known lower bounds. Moreover, our fully adaptive algorithm supports more general sub-Gaussian sources in addition to just coins, and only ever has one “coin” outside the bag at a given time, a constraint of practical importance to some applications.

In addition, we connect the most biased coin problem to anomaly detection and prove novel lower bounds on the difficulty of detecting the presence of a mixture versus just a single component of a known family of distributions (e.g. $X \sim (1 - \alpha)g_{\theta_0} + \alpha g_{\theta_1}$ versus $X \sim g_{\theta}$ for some θ). We show that in detecting the presence of a mixture distribution, there is a stark difference of difficulty

between when the underlying distribution parameters are known (e.g. $\alpha, \theta_0, \theta_1$) and when they are not. The most biased coin problem can be viewed as an online, adaptive mixture detection problem where source distributions arrive one at a time that are either g_{θ_0} with probability $(1 - \alpha)$ or g_{θ_1} with probability α (e.g. null or anomalous) and the player adaptively chooses how many samples to take from each distribution (to increase the signal-to-noise ratio) with the goal of identifying an anomalous distribution f_{θ_1} using as few total number of samples as possible. One consequence of this work is drawing a contrast between the power of an adaptive versus non-adaptive (e.g. taking the same number of samples each time) approaches to this problem, specifically when $\alpha, \theta_0, \theta_1$ are unknown.

1.1 Motivation and Related Work for the Most Biased Coin Problem

The most biased coin problem characterizes the inherent difficulty of real-world problems including anomaly and intrusion detection and discovery of vacant frequencies in the radio spectrum. Our interest in the problem stemmed from automated hiring of crowd workers: data labeling for machine learning applications is often performed by humans, and recent work in the crowdsourcing literature accelerates labeling by organizing workers into pools of labelers and paying them to wait for incoming data [3, 11]. Workers hired on marketplaces such as Amazon’s Mechanical Turk [15] vary widely in skill, and identifying high-quality workers as quickly as possible is an important challenge. We can model each worker’s performance (e.g. accuracy or speed) as a random variable so that selecting a good worker is equivalent to identifying a worker with a high mean. Since we do not observe a worker’s expected performance directly, we must give them tasks from which we estimate it (like repeatedly flipping a biased coin).

The most biased coin problem was first proposed by Chandrasekaran and Karp [7]. In that work, it was shown that if $\alpha, \theta_0, \theta_1$ were known then there exists an algorithm based on the sequential probability ratio test (SPRT) that is optimal in that it minimizes the expected number of total flips to find a “heavy” coin whose posterior probability of being heavy is at least $1 - \delta$, and the expected sample complexity of this algorithm was upper-bounded by

$$\frac{16}{(\theta_1 - \theta_0)^2} \left(\frac{1 - \alpha}{\alpha} + \log \left(\frac{(1 - \alpha)(1 - \delta)}{\alpha\delta} \right) \right). \quad (1)$$

However, the practicality of the proposed algorithm is severely limited as it relies critically on knowing α, θ_0 , and θ_1 exactly. In addition, the algorithm returns to coins it has previously flipped and thus requires more than one coin to be outside the bag at a time, ruling out some applications. Malloy et al. [14] addressed some of the shortcomings of [8] (a preprint of [7]) by considering both an alternative SPRT procedure and a sequential thresholding procedure. Both of these proposed algorithms only ever have one coin out of the bag at a time. However, the former requires knowledge of all relevant parameters $\alpha, \theta_0, \theta_1$, and the latter requires knowledge of α, θ_0 . Moreover, these results are only presented for the asymptotic case where $\delta \rightarrow 0$.

The most biased coin problem can be viewed through the lens of multi-armed bandits. In the best-arm identification problem, the player has access to K distributions (arms) such that if arm $i \in [K]$ is sampled (pulled), an iid random variable with mean μ_i is observed; the objective is to identify the arm associated with the highest mean with probability at least $1 - \delta$ using as few pulls as possible (see [13] for a short survey). In the *infinite* armed bandit problem, the player is not confined to K arms but an infinite reservoir of arms such that a draw from this reservoir results in an arm with a mean μ drawn from some distribution; the objective is to identify the highest mean possible after n total pulls for any $n > 0$ with probability $1 - \delta$ (see [6]). The most biased coin problem is an instance of this latter game with the arm reservoir distribution of means μ defined as $\mathbb{P}(\mu \geq \theta_1 - \epsilon) = \alpha \mathbf{1}_{\epsilon > 0} + (1 - \alpha) \mathbf{1}_{\epsilon \geq \theta_1 - \theta_0}$ for all ϵ . Previous work has focused on an alternative arm distribution reservoir that satisfies $E\epsilon^\beta \leq \mathbb{P}(\mu \geq \mu_* - \epsilon) \leq E'\epsilon^\beta$ for some $\mu_* \in [0, 1]$ where E, E' are constants and β is known [4, 20, 5, 6]. Because neither arm distribution reservoir can be written in terms of the other, neither work subsumes the other. Note that one can always apply an algorithm designed for the infinite armed bandit problem to any finite K -armed bandit problem by defining the arm reservoir as placing a uniform distribution over the K arms. This is appealing when K is very large and one wishes to guarantee nontrivial performance when the number of pulls is much less than K^1 . The most biased problem is a special case of the K -armed reservoir distribution where one arm has mean θ_1 and $K - 1$ arms have mean θ_0 with $\alpha = \frac{1}{K}$.

¹All algorithms for K -armed bandit problem known to these authors begins by sampling each arm once so that until the number of pulls exceeds K , performance is no better than random selection.

Given that [7] and [14] are provably optimal algorithms for the most biased coin problem given knowledge of $\alpha, \theta_0, \theta_1$, it is natural to consider a procedure that first estimates these unknown parameters first and then uses these estimates in the algorithms of [7] or [14]. Indeed, in the β -parameterized arm reservoir setting discussed above, this is exactly what Carpentier and Valko [6] propose to do, suggesting a particular estimator for β given a lower bound $\hat{\beta} \leq \beta$. They show that this estimator is sufficient to obtain the same sample complexity result up to log factors as when β was known. Sadly, through upper and lower bounds we show that for the most biased coin problem this *estimate-then-explore* approach requires quadratically more flips than our proposed algorithm that adapts to these unknown parameters. Specifically, we show that when $\theta_1 - \theta_0$ is sufficiently small one cannot use a static estimation step to determine whether $\alpha = 0$ or $\alpha > 0$ unless a number of samples *quadratic* in the optimal sample complexity are taken.

Our contributions to the most biased coin problem include a novel algorithm that never has more than one coin outside the bag at a time, has no knowledge of the distribution parameters, supports distributions on $[0, 1]$ rather than just “coins,” and comes within log factors of the known information-theoretic lower bound and Equation 1 which is achieved by an algorithm that knows the parameters. See Table 1 for an overview of the upper and lower bounds proved in this work for this problem. We believe that our algorithm is the first solution to the most biased coin problem that does not require prior knowledge of the problem parameters and that the same approach can be reworked to solve more general instances of the infinite-armed bandit problem, including the β -parameterized and K -armed reservoir cases described of above. Finally, if an algorithm is desired for arbitrary arm reservoir distributions, this work rules out an estimate-then-explore approach.

1.2 Problem Statement

Let $\theta \in \Theta$ index a family of single-parameter probability density functions g_θ and fix $\theta_0, \theta_1 \in \Theta$, $\alpha \in [0, 1/2]$. For any $\theta \in \Theta$ assume that g_θ is known to the procedure. Note that in the most biased coin problem, $g_\theta = \text{Bernoulli}(\theta)$, but in general it is arbitrary (e.g. $\mathcal{N}(\theta, 1)$). Consider a sequence of iid Bernoulli random variables $\xi_i \in \{0, 1\}$ for $i = 1, 2, \dots$ where each $\mathbb{P}(\xi_i = 1) = 1 - \mathbb{P}(\xi_i = 0) = \alpha$. Let $X_{i,j}$ for $j = 1, 2, \dots$ be a sequence of random variables drawn from g_{θ_1} if $\xi_i = 1$ and g_{θ_0} otherwise, and let $\{\{X_{i,j}\}_{j=1}^{M_i}\}_{i=1}^N$ represent the sampling history generated by a procedure for some $N \in \mathbb{N}$ and $(M_1, \dots, M_N) \in \mathbb{N}^N$. Any valid procedure behaves accordingly:

Algorithm 1 The most biased coin problem definition. Only the last distribution drawn may be sampled or declared heavy, enforcing the rule that only one coin may be outside the bag at a time.

Initialize an empty history ($N = 1, M = (0, 0, \dots)$).

Repeat until heavy distribution declared:

Choose one of

1. draw a sample from distribution N , $M_N \leftarrow M_N + 1$
 2. draw a sample from the $(N + 1)$ st distribution, $M_{N+1} = 1, N \leftarrow N + 1$
 3. declare distribution N as heavy
-

Definition 1 We say a strategy for the most biased coin problem is δ -**probably correct** if for all $(\alpha, \theta_0, \theta_1)$ it identifies a “heavy” g_{θ_1} distribution with probability at least $1 - \delta$.

Definition 2 (Strategies for the most biased coin problem) An **estimate-then-explore strategy** is a strategy that, for any fixed $m \in \mathbb{N}$, begins by sampling each successive coin exactly m times for a number of coins that is at least the minimum necessary for any test to determine that $\alpha \neq 0$ with probability at least $1 - \delta$, then optionally continues sampling with an arbitrary strategy that declares a heavy coin. An **adaptive strategy** is any strategy that is not an estimate-then-explore strategy.

We study the *estimate-then-explore strategy* because there exist optimal algorithms [7, 14] for the most biased coin problem if $\alpha, \theta_0, \theta_1$ are known, so it is natural to consider estimating these quantities then using one of these algorithms. Note that the algorithm of [6] for the β -parameterized infinite armed bandit problem discussed above can be considered an *estimate-then-explore strategy* since it first estimates β by sampling a fixed number of samples from a set of arms, and then uses this estimate to draw a fixed number of arms and applies a UCB-style algorithm to these arms. A contribution of this work is showing that such a strategy is infeasible for the most biased coin problem.

For all strategies that are δ -probably correct and follow the interface of Algorithm 1, our goal is to provide lower and upper bounds on the quantity $\mathbb{E}[T] := \mathbb{E}[\sum_{i=1}^N M_i]$ for any $(\alpha, \theta_0, \theta_1)$ if N denotes the final number of coins considered.

2 From Identifying Coins to Detecting Mixture Distributions

Addressing the most biased coin problem, [14] analyzes perhaps the most natural strategy: fix an $m \in \mathbb{N}$ and flip each successive coin exactly m times. The relevant questions are how large does m have to be in order to guarantee correctness with probability $1 - \delta$, and for a given m how long must one wait to declare a “heavy” coin? The authors partially answer these questions and we improve upon them (see Section 3.2.1) which leads us to our study of the difficulty of detecting the presence of a mixture distribution. As an example of the kind of lower bounds shown in this work, if we observe a sequence of random variables X_1, \dots, X_n , consider the following hypothesis test:

$$\begin{aligned} \mathbf{H}_0 &: \forall i \ X_1, \dots, X_n \sim \mathcal{N}(\theta, \sigma^2) \quad \text{for some } \theta \in \mathbb{R}, \\ \mathbf{H}_1 &: \forall i \ X_1, \dots, X_n \sim (1 - \alpha)\mathcal{N}(\theta_0, \sigma^2) + \alpha\mathcal{N}(\theta_1, \sigma^2) \end{aligned} \quad (\text{P1})$$

which will henceforth be referred to as Problem P1 or just (P1). We can show that if $\theta_0, \theta_1, \alpha$ are *known* and $\theta = \theta_0$, then it is *sufficient* to observe just $\max\{1/\alpha, \frac{\sigma^2}{\alpha^2(\theta_1 - \theta_0)^2} \log(1/\delta)\}$ samples to determine the correct hypothesis with probability at least $1 - \delta$. However, if $\theta_0, \theta_1, \alpha$ are *unknown* then it is *necessary* to observe at least $\max\{1/\alpha, (\frac{\sigma^2}{\alpha(\theta_1 - \theta_0)^2})^2 \log(1/\delta)\}$ samples in expectation whenever $\frac{(\theta_1 - \theta_0)^2}{\sigma^2} \leq 1$ and $\max\{1/\alpha, \frac{\sigma^2}{\alpha^2(\theta_1 - \theta_0)^2} \log(1/\delta)\}$ otherwise (see Appendix C).

Recognizing $\frac{(\theta_1 - \theta_0)^2}{\sigma^2}$ as the KL divergence between two Gaussians of \mathbf{H}_1 , we observe startling consequences for anomaly detection when the parameters of the underlying distributions are unknown: if the anomalous distribution is well separated from the null distribution, then detecting an anomalous component is only about as hard as observing just one anomalous sample (i.e. $1/\alpha$) multiplied by the inverse KL divergence between the null and anomalous distributions. However, when the two distributions are *not* well separated then the necessary sample complexity explodes to this latter quantity *squared*. In Section 4 we will investigate adaptive methods for dramatically decreasing this sample complexity.

Our lower bounds are based on the detection of the presence of a mixture of two distributions of an exponential family versus just a single distribution of the same family. There has been extensive work in the estimation of mixture distributions [12, 10] but this literature often assumes that the mixture coefficient α is bounded away from 0 and 1 to ensure a sufficient number of samples from each distribution. In contrast, we highlight the regime when α is arbitrarily small, as is the case in statistical anomaly detection [9, 19, 2]. Property testing, e.g. unimodality, [1] is relevant but can lack interpretability or strength in favor of generality. Considering the exponential family allowing us to make interpretable statements about the relevant problem parameters in different regimes.

Preliminaries Let P and Q be two probability distributions with densities p and q , respectively. For simplicity, assume p and q have the same support. Define the *KL Divergence* between P and Q as $KL(P, Q) = \int \log\left(\frac{p(x)}{q(x)}\right) dp(x)$. Define the χ^2 *Divergence* between P and Q as $\chi^2(P, Q) = \int \left(\frac{p(x)}{q(x)} - 1\right)^2 dq(x) = \int \frac{(p(x) - q(x))^2}{q(x)} dx$. Note that by Jensen’s inequality

$$KL(P, Q) = \mathbb{E}_p\left[\log\left(\frac{p}{q}\right)\right] \leq \log\left(\mathbb{E}_p\left[\frac{p}{q}\right]\right) = \log(\chi^2(P, Q) + 1) \leq \chi^2(P, Q). \quad (2)$$

Examples: If $P = \mathcal{N}(\theta_1, \sigma^2)$ and $Q = \mathcal{N}(\theta_0, \sigma^2)$ then $KL(P, Q) = \frac{(\theta_1 - \theta_0)^2}{2\sigma^2}$ and $\chi^2(P, Q) = e^{\frac{(\theta_1 - \theta_0)^2}{\sigma^2}} - 1$. If $P = \text{Bernoulli}(\theta_1)$ and $Q = \text{Bernoulli}(\theta_0)$ then $KL(P, Q) = \theta_1 \log\left(\frac{\theta_1}{\theta_0}\right) + (1 - \theta_1) \log\left(\frac{1 - \theta_1}{1 - \theta_0}\right) \leq \frac{(\theta_1 - \theta_0)^2/2}{\theta_0(1 - \theta_0) - [(\theta_1 - \theta_0)(2\theta_0 - 1)]_+}$ and $\chi^2(P, Q) = \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}$. All proofs appear in the appendix.

3 Lower bounds

We present lower bounds on the sample complexity of δ -probably correct strategies for the most biased coin problem that follow the interface of Algorithm 1. Lower bounds are stated for any

adaptive strategy in Section 3.1, non-adaptive strategies that may have knowledge of the parameters but sample each distribution the same number of times in Section 3.2.1, and *estimate-then-explore* strategies that do not have prior knowledge of the parameters in Section 3.2.2. Our lower bounds, with the exception of the adaptive strategy, are based on the difficulty of detecting the presence of a mixture distribution, and this reduction is explained in Section 3.2.

3.1 Adaptive strategies

The following theorem, reproduced from [14], describes the sample complexity of any δ -probably correct algorithm for the most biased coin identification problem. Note that this lower bound holds for any procedure even if it returns to previously seen distributions to draw additional samples and even if it knows $\alpha, \theta_0, \theta_1$.

Theorem 1 [14, Theorem 2] Fix $\delta \in (0, 1)$. Let T be the total number of samples taken of any procedure that is δ -probably correct in identifying a heavy distribution. Then

$$\mathbb{E}[T] \geq c_1 \max \left\{ \frac{1 - \delta}{\alpha}, \frac{(1 - \delta)}{\alpha KL(g_{\theta_0} | g_{\theta_1})} \right\}$$

whenever $\alpha \leq c_2 \delta$ where $c_1, c_2 \in (0, 1)$ are absolute constants.

The above theorem is directly applicable to the special case where g_θ is a Bernoulli distribution, implying a lower bound of $\max \left\{ \frac{1 - \delta}{\alpha}, \frac{2 \min\{\theta_0(1 - \theta_0), \theta_1(1 - \theta_1)\}}{\alpha(\theta_1 - \theta_0)^2} \right\}$ for the most biased coin problem. The upper bounds of our proposed procedures for the most biased coin problem presented later will be compared to this benchmark.

3.2 The detection of a mixture distribution and the most biased coin problem

First observe that identifying a specific distribution $i \leq N$ as heavy (i.e. $\xi_i = 1$) or determining that α is strictly greater than 0, is at least as hard as detecting that *any* of the distributions up to distribution N is heavy. Thus, a lower bound on the total expected number of samples of all considered distributions for this strictly easier detection problem is also a lower bound for the estimate-then-explore strategy for the most biased coin identification problem.

The estimate-then-explore strategy fixes an $m \in \mathbb{N}$ prior to starting the game and then samples each distribution exactly m times, i.e. $M_i = m$ for all $i \leq N$ for some N . To simplify notation let f_θ denote the distribution of the sufficient statistics of these m samples. In general f_θ is a product distribution, but when g_θ is a Bernoulli distribution, as in the biased coin problem, we can take f_θ to be a Binomial distribution with parameters (m, θ) . Now our problem is more succinctly described as:

$$\begin{aligned} \mathbf{H}_0 : \forall i \ X_i \sim f_\theta \quad \text{for some } \theta \in \tilde{\Theta} \subseteq \Theta, \\ \mathbf{H}_1 : \forall i \ \xi_i \sim \text{Bernoulli}(\alpha), \quad \forall i \ X_i \sim \begin{cases} f_{\theta_0} & \text{if } \xi_i = 0 \\ f_{\theta_1} & \text{if } \xi_i = 1 \end{cases} \end{aligned} \quad (\text{P2})$$

If θ_0 and θ_1 are close to each other, or if α is very small, it can be very difficult to decide between \mathbf{H}_0 and \mathbf{H}_1 even if $\alpha, \theta_0, \theta_1$ are known a priori. Note that when the parameters are *known*, one can take $\tilde{\Theta} = \{\theta_0\}$. However, when the parameters are *unknown*, one takes $\tilde{\Theta} = \Theta$ to prove a lower bound on the sample complexity of the estimate-then-explore algorithm, which is tasked with deciding whether or not samples are coming from a mixture of distributions or just a single distribution within the family. That is, lower bounds on the sample complexity when the parameters are known and unknown follow by analyzing a simple binary and composite hypothesis test, respectively. In what follows, for any event A , let $\mathbb{P}_i(A)$ and $\mathbb{E}_i[A]$ denote probability and expectation of A under hypothesis \mathbf{H}_i for $i \in \{0, 1\}$ (the specific value of θ in \mathbf{H}_0 will be clear from context). The next claim is instrumental in our ability to prove lower bounds on the difficulty of the hypothesis tests.

Claim 1 Any procedure that is δ -probably correct also satisfies $\mathbb{P}_0(N < \infty) \leq \delta$ whenever $\alpha = 0$.

3.2.1 Sample complexity when parameters are known

Theorem 2 Fix $\delta \in (0, 1)$. Consider the hypothesis test of Problem P2 for any fixed $\theta \in \tilde{\Theta} \subseteq \Theta$. Let N be the random number of distributions considered before stopping and declaring a

hypothesis. If a procedure satisfies $\mathbb{P}_0(N < \infty) \leq \delta$ and $\mathbb{P}_1(\cup_{i=1}^N \{\xi_i = 1\}) \geq 1 - \delta$, then $\mathbb{E}_1[N] \geq \max \left\{ \frac{1-\delta}{\alpha}, \frac{\log(1/\delta)}{KL(\mathbb{P}_1|\mathbb{P}_0)} \right\} \geq \max \left\{ \frac{1-\delta}{\alpha}, \frac{\log(1/\delta)}{\chi^2(\mathbb{P}_1|\mathbb{P}_0)} \right\}$. In particular, if $\tilde{\Theta} = \{\theta_0\}$ then

$$\mathbb{E}_1[N] \geq \max \left\{ \frac{1-\delta}{\alpha}, \frac{\log(1/\delta)}{\alpha^2 \chi^2(f_{\theta_1}|f_{\theta_0})} \right\}.$$

The next corollary relates Theorem 2 to the most biased coin problem and is related to Malloy et al. [14, Theorem 4] that considers the limit as $\alpha \rightarrow 0$ and assumes m is sufficiently large (specifically, large enough for the Chernoff-Stein lemma to apply). In contrast, our result holds for all finite δ, α, m .

Corollary 1 *Fix $\delta \in (0, 1)$. For any $m \in \mathbb{N}$ consider a δ -probably correct strategy that flips each coin exactly m times. If N_m is the number of coins considered before declaring a coin as heavy then*

$$\min_{m \in \mathbb{N}} \mathbb{E}[mN_m] \geq \frac{(1-\delta) \log \left(\frac{\log(1/\delta)}{\alpha} \right) \theta_0(1-\theta_0)}{\alpha (\theta_1 - \theta_0)^2}.$$

One can show the existence of such a strategy with a nearly matching upperbound when $\alpha, \theta_0, \theta_1$ are known (see Appendix B.1). Note that this is at least $\log(1/\alpha)$ larger than the sample complexity of (1) that can be achieved by an adaptive algorithm when the parameters are known.

3.2.2 Sample complexity when parameters are unknown

If α, θ_0 , and θ_1 are unknown, we cannot test f_{θ_0} against the mixture $(1-\alpha)f_{\theta_0} + \alpha f_{\theta_1}$. Instead, we have the general composite test of *any* individual distribution against *any* mixture, which is at least as hard as the hypothesis test of Problem P2 with $\tilde{\Theta} = \{\theta\}$ for some particular worst-case setting of θ . Without any specific form of f_{θ} , it is difficult to pick a worst case θ that will produce a tight bound. Consequently, in this section we consider single parameter exponential families (defined formally below) to provide us with a class of distributions in which we can reason about different possible values for θ . Since exponential families include Bernoulli, Gaussian, exponential, and many other distributions, the following theorem is general enough to be useful in a wide variety of settings. The constant C referred to in the next theorem is an absolute constant under certain conditions that we outline in the following remark and corollary, its explicit form is given in the proof.

Theorem 3 *Suppose f_{θ} for $\theta \in \Theta \subset \mathbb{R}$ is a single parameter exponential family so that $f_{\theta}(x) = h(x) \exp(\eta(\theta)x - b(\eta(\theta)))$ for some scalar functions h, b, η where η is strictly increasing. If $\tilde{\Theta} = \{\theta_*\}$ where $\theta_* = \eta^{-1}((1-\alpha)\eta(\theta_0) + \alpha\eta(\theta_1))$ and N is the stopping time of any procedure that satisfies $\mathbb{P}_0(N < \infty) \leq \delta$ and $\mathbb{P}_1(\cup_{i=1}^N \{\xi_i = 1\}) \geq 1 - \delta$, then*

$$\mathbb{E}_1[N] \geq \max \left\{ \frac{1-\delta}{\alpha}, \frac{\log(\frac{1}{\delta})}{C(\frac{1}{2}\alpha(1-\alpha)(\eta(\theta_1) - \eta(\theta_0))^2)^2} \right\}.$$

where C is a constant that may depend on $\alpha, \theta_0, \theta_1$.

The following remark and corollary apply Theorem 3 to the special cases of Gaussian mixture model detection and the most biased coin problem, respectively.

Remark 1 *When $\alpha, \theta_0, \theta_1$ are unknown, any procedure has no knowledge of $\tilde{\Theta}$ in Problem P2 and consequently it cannot rule out $\theta = \theta_*$ for \mathbf{H}_0 where θ_* is defined in Theorem 3. If $f_{\theta} = \mathcal{N}(\theta, \sigma^2)$ for known σ , then whenever $\frac{(\theta_1 - \theta_0)^2}{\sigma^2} \leq 1$ the constant C in Theorem 3 is an absolute constant and consequently, $\mathbb{E}_1[N] = \Omega\left(\left(\frac{\sigma^2}{\alpha(\theta_1 - \theta_0)^2}\right)^2 \log(1/\delta)\right)$. Conversely, when $\alpha, \theta_0, \theta_1$ are known, then we simply need to determine whether samples came from $\mathcal{N}(\theta_0, \sigma^2)$ or $(1-\alpha)\mathcal{N}(\theta_0, \sigma^2) + \alpha\mathcal{N}(\theta_1, \sigma^2)$, and we show that it is sufficient to take just $O\left(\frac{\sigma^2}{\alpha^2(\theta_1 - \theta_0)^2} \log(1/\delta)\right)$ samples (see Appendix C).*

Corollary 2 *Fix $\delta \in [0, 1]$ and assume θ_0, θ_1 are bounded sufficiently far from $\{0, 1\}$ such that $2(\theta_1 - \theta_0) \leq \min\{\theta_0(1 - \theta_0), \theta_1(1 - \theta_1)\}$. For any m let N_m be the number of coins a δ -probably correct estimate-then-explore strategy that flips each coin m times in the exploration step. Then*

$$m\mathbb{E}[N_m] \geq \frac{c' \min\{\frac{1}{m}, \theta_*(1 - \theta_*)\}}{\left(\alpha(1-\alpha)\frac{(\theta_1 - \theta_0)^2}{\theta_*(1-\theta_*)}\right)^2} \log(\frac{1}{\delta}) \quad \text{whenever} \quad m \leq \frac{\theta_*(1 - \theta_*)}{(\theta_1 - \theta_0)^2}.$$

where c' is an absolute constant and $\theta_* = \eta^{-1}((1-\alpha)\eta(\theta_0) + \alpha\eta(\theta_1)) \in [\theta_0, \theta_1]$.

Remark 2 If $\alpha, \theta_0, \theta_1$ are unknown, any estimate-then-explore strategy (or the strategy described in Corollary 1) would be unable to choose an m that depended on these parameters, so we can treat it as a constant. Thus, for the case when θ_0 and θ_1 are bounded away from $\{0, 1\}$ (e.g. $\theta_0, \theta_1 \in [1/8, 7/8]$), the above corollary states that for any fixed m , whenever $\theta_1 - \theta_0$ is sufficiently small the number of samples necessary for these strategies to identify a heavy coin scales like $(\frac{1}{\alpha(\theta_1 - \theta_0)^2})^2 \log(1/\delta)$. This is striking example of the difference when parameters are known versus when they are not and effectively rules out an estimate-then-explore strategy for practical purposes.

Setting	Upper Bound	Lower Bound
Fixed, known $\alpha, \theta_0, \theta_1$	$\frac{\log(1/(\delta\alpha))}{\alpha\epsilon^2}$, Thm. 7	$\frac{\log(\log(1/\delta)/\alpha)}{\alpha\epsilon^2}$ Cor. 1
Adaptive, known $\alpha, \theta_0, \theta_1$	$\frac{1}{\epsilon^2} (\frac{1}{\alpha} + \log(\frac{1}{\delta}))$ [7, 14], Thm. 4	$\frac{1}{\alpha\epsilon^2}$ [14]
Est+Expl, unknown $\alpha, \theta_0, \theta_1$	Unconsidered [†]	$(\frac{1}{\alpha\epsilon^2})^2 \log(\frac{1}{\delta})$ Cor. 2
Adaptive, unknown $\alpha, \theta_0, \theta_1$	$\frac{c \log(\frac{1}{\alpha\epsilon^2}) \log(\log(\frac{1}{\alpha\epsilon^2})/\delta)}{\alpha\epsilon^2}$ Thm. 5	$\frac{1}{\alpha\epsilon^2}$ [14]

Table 1: Upper and lower bounds on the expected sample complexity of different δ -probably correct strategies. Fixed refers to the strategy of Corollary 1. For this table, we assume $\min\{\theta_0(1 - \theta_0), \theta_1(1 - \theta_1)\}$ is lower bounded by a constant (e.g. $\theta_0, \theta_1 \in [1/8, 7/8]$) and $\epsilon = \theta_1 - \theta_0$ is sufficiently small. Also note that the upperbounds apply to distributions supported on $[0, 1]$, not just coins. All results without bracketed citations were unknown prior to this work. [†] Due to our discouraging lower bound for any estimate-then-explore strategy, it is inadvisable to propose an algorithm.

4 Near optimal adaptive algorithm

In this section we propose an algorithm that has no prior knowledge of the parameters $\alpha, \theta_0, \theta_1$ yet yields an upper bound that matches the lower bound of Theorem 1 up to logarithmic factors. We assume that samples from heavy or light distributions are supported on $[0, 1]$, and that drawn samples are independent and unbiased estimators of the mean, i.e., $\mathbb{E}[X_{i,j}] = \mu_i$ for $\mu_i \in \{\theta_0, \theta_1\}$. All results can be easily extended to sub-Gaussian distributions. Consider Algorithm 2, an SPRT-like procedure [17] for finding a heavy distribution given δ and lower bounds on α and $\epsilon = \theta_1 - \theta_0$. It improves upon prior work by supporting arbitrary distributions on $[0, 1]$ and requires only bounds α, ϵ .

Algorithm 2 Adaptive strategy for heavy distribution identification with inputs $\alpha_0, \epsilon_0, \delta$

Given $\delta \in (0, 1/4), \alpha_0 \in (0, 1/2), \epsilon_0 \in (0, 1)$.

Initialize $n = \lceil 2 \log(9)/\alpha_0 \rceil, m = \lceil 64\epsilon_0^{-2} \log(14n/\delta) \rceil, A = -8\epsilon_0^{-1} \log(21),$

$B = 8\epsilon_0^{-1} \log(14n/\delta), k_1 = 5, k_2 = \lceil 8\epsilon_0^{-2} \log(2k_1/\min\{\delta/8, m^{-1}\epsilon_0^{-2}\}) \rceil.$

Draw k_1 distributions and sample them each k_2 times.

Estimate $\hat{\theta}_0 = \min_{i=1, \dots, k_1} \hat{\mu}_{i, k_2}, \hat{\gamma} = \hat{\theta}_0 + \epsilon_0/2.$

Repeat for $i = 1, \dots, n$:

Draw distribution i .

Repeat for $j = 1, \dots, m$:

Sample distribution i and observe $X_{i,j}$.

If $\sum_{k=1}^j (X_{i,k} - \hat{\gamma}) > B$:

Declare distribution i to be heavy and **Output** distribution i .

Else if $\sum_{k=1}^j (X_{i,k} - \hat{\gamma}) < A$:

break.

Output null.

Theorem 4 If Algorithm 2 is run with $\delta \in (0, 1/4), \alpha_0 \in (0, 1/2), \epsilon_0 \in (0, 1)$, then the expected number of total samples taken by the algorithm is no more than

$$\frac{c' \alpha \log(1/\alpha_0) + c'' \log(\frac{1}{\delta})}{\alpha_0 \epsilon_0^2} \quad (3)$$

for some absolute constants c, c' , and all of the following hold: 1) with probability at least $1 - \delta$, a light distribution is not returned, 2) if $\epsilon_0 \leq \theta_1 - \theta_0$ and $\alpha_0 \leq \alpha$, then with probability $\frac{4}{5}$ a heavy distribution is returned, and 3) the procedure takes no more than $\frac{c \log(1/(\alpha_0 \delta))}{\alpha_0 \epsilon_0^2}$ total samples.

The second claim of the theorem holds only with constant probability (versus with probability $1 - \delta$) since the probability of observing a heavy distribution among the $n = \lceil 2 \log(4)/\alpha_0 \rceil$ distributions only occurs with constant probability. One can show that if the outer loop of algorithm is allowed to run indefinitely (with m and n defined as is), $\epsilon_0 = \theta_1 - \theta_0$, $\alpha_0 = \alpha$, and $\hat{\theta}_0 = \theta_0$, then a heavy coin is returned with probability at least $1 - \delta$ and the expected number of samples is bounded by (3). If a tight lower bound is known on either $\epsilon = \theta_1 - \theta_0$ or α , there is only one parameter that is unknown and the “doubling trick”, along with Theorem 4, can be used to identify a heavy coin with just $\frac{\log(\log(\epsilon^{-2})/\delta)}{\alpha \epsilon^2}$ and $\frac{\log(\log(\alpha^{-1})/\delta)}{\alpha \epsilon^2}$ samples, respectively (see Appendix B.3).

Now consider Algorithm 3 that assumes no prior knowledge of $\alpha, \theta_0, \theta_1$, the first result for this setting that we are aware of. We remark that while the placing of “landmarks” (α_k, ϵ_k) throughout the search space as is done in Algorithm 3 appears elementary in hindsight, it is surprising that so few can cover this two dimensional space since one has to balance the exploration of α and ϵ . We believe similar a similar approach may be generalized for more generic infinite armed bandit problems.

Algorithm 3 Adaptive strategy for heavy distribution identification with unknown parameters

Given $\delta > 0$.

Initialize $\ell = 1$, heavy distribution $h = \text{null}$.

Repeat until h is not null:

Set $\gamma_\ell = 2^\ell, \delta_\ell = \delta/(2^{\ell^3})$

Repeat for $k = 0, \dots, \ell$:

Set $\alpha_k = \frac{2^k}{\gamma_\ell}, \epsilon_k = \sqrt{\frac{1}{2\alpha_k \gamma_\ell}}$

Run Algorithm 2 with $\alpha_0 = \alpha_k, \epsilon_0 = \epsilon_k, \delta = \delta_\ell$ and **Set** h to its output.

If h is not null **break**

Set $\ell = \ell + 1$

Output h

Theorem 5 (Unknown $\alpha, \theta_0, \theta_1$) Fix $\delta \in (0, 1)$. If Algorithm 3 is run with δ then with probability at least $1 - \delta$ a heavy distribution is returned and the expected number of total samples taken is bounded by

$$c \frac{\log_2(\frac{1}{\alpha \epsilon^2})}{\alpha \epsilon^2} (\alpha \log_2(\frac{1}{\epsilon^2}) + \log(\log_2(\frac{1}{\alpha \epsilon^2})) + \log(1/\delta))$$

for an absolute constant c .

5 Conclusion

While all prior works have required at least partial knowledge of $\alpha, \theta_0, \theta_1$ to solve the most biased coin problem, our algorithm requires no knowledge of these parameters yet obtain the near-optimal sample complexity. In addition, we have proved lower bounds on the sample complexity of detecting the presence of a mixture distribution when the parameters are known or unknown, with consequences for any estimate-then-explore strategy, an approach previously proposed for an infinite armed bandit problem. Extending our adaptive algorithm to arbitrary arm reservoir distributions is of significant interest. We believe a successful algorithm in this vein could have a significant impact on how researchers think about sequential decision processes in both finite and uncountable action spaces.

Acknowledgments Kevin Jamieson is generously supported by ONR awards N00014-15-1-2620, and N00014-13-1-0129. This research is supported in part by NSF CISE Expeditions Award CCF-1139158, DOE Award SN10040 DE-SC0012463, and DARPA XData Award FA8750-12-2-0331, and gifts from Amazon Web Services, Google, IBM, SAP, The Thomas and Stacey Siebel Foundation, Apple Inc., Arimo, Blue Goji, Bosch, Cisco, Cray, Cloudera, Ericsson, Facebook, Fujitsu, Guavus, HP, Huawei, Intel, Microsoft, Pivotal, Samsung, Schlumberger, Splunk, State Farm and VMware.

References

- [1] Jayadev Acharya, Constantinos Daskalakis, and Gautam C Kamath. Optimal testing for properties of distributions. In *Advances in Neural Information Processing Systems*, pages 3577–3598, 2015.
- [2] Deepak Agarwal. Detecting anomalies in cross-classified streams: a bayesian approach. *Knowledge and Information Systems*, 11(1):29–44, 2006.
- [3] Michael S Bernstein, Joel Brandt, Robert C Miller, and David R Karger. Crowds in two seconds: enabling realtime crowd-powered interfaces. *UIST*, 2011.
- [4] Donald A. Berry, Robert W. Chen, Alan Zame, David C. Heath, and Larry A. Shepp. Bandit problems with infinitely many arms. *Ann. Statist.*, 25(5):2103–2116, 10 1997.
- [5] Thomas Bonald and Alexandre Proutiere. Two-target algorithms for infinite-armed bandits with bernoulli rewards. In C.J.C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems 26*, pages 2184–2192. Curran Associates, Inc., 2013.
- [6] Alexandra Carpentier and Michal Valko. Simple regret for infinitely many armed bandits. *arXiv preprint arXiv:1505.04627*, 2015.
- [7] Karthekeyan Chandrasekaran and Richard Karp. Finding a most biased coin with fewest flips. In *Proceedings of The 27th Conference on Learning Theory*, pages 394–407, 2014.
- [8] Karthekeyan Chandrasekaran and Richard M. Karp. Finding the most biased coin with fewest flips. *CoRR*, abs/1202.3639, 2012. URL <http://arxiv.org/abs/1202.3639>.
- [9] Eleazar Eskin. Anomaly detection over noisy data using learned probability distributions. In *Proceedings of the Seventeenth International Conference on Machine Learning, ICML '00*, pages 255–262, San Francisco, CA, USA, 2000. Morgan Kaufmann Publishers Inc.
- [10] Yoav Freund and Yishay Mansour. Estimating a mixture of two product distributions. In *Proceedings of the twelfth annual conference on Computational learning theory*, pages 53–62. ACM, 1999.
- [11] Daniel Haas, Jiannan Wang, Eugene Wu, and Michael J. Franklin. Clamshell: Speeding up crowds for low-latency data labeling. *Proc. VLDB Endow.*, 9(4):372–383, December 2015. ISSN 2150-8097.
- [12] Moritz Hardt and Eric Price. Sharp bounds for learning a mixture of two gaussians. *ArXiv e-prints*, 1404, 2014.
- [13] Kevin Jamieson and Robert Nowak. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2014.
- [14] Matthew L Malloy, Gongguo Tang, and Robert D Nowak. Quickest search for a rare distribution. In *Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2012.
- [15] MTurk. Amazon Mechanical Turk. <https://www.mturk.com/>.
- [16] David Pollard. Asymptopia. *Manuscript in progress. Available at <http://www.stat.yale.edu/~pollard>*, 2000.
- [17] David Siegmund. *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media, 2013.
- [18] Robert Spira. Calculation of the gamma function by stirling’s formula. *mathematics of computation*, pages 317–322, 1971.
- [19] Gautam Thattai, Urbashi Mitra, and John Heidemann. Parametric methods for anomaly detection in aggregate traffic. *IEEE/ACM Trans. Netw.*, 19(2):512–525, April 2011. ISSN 1063-6692.
- [20] Yizao Wang, Jean yves Audibert, and Rémi Munos. Algorithms for infinitely many-armed bandits. In D. Koller, D. Schuurmans, Y. Bengio, and L. Bottou, editors, *Advances in Neural Information Processing Systems 21*, pages 1729–1736. Curran Associates, Inc., 2009.

A Proofs of Lower Bounds

A.1 Proof of Claim 1

Proof Suppose there exists a δ -probably correct procedure with $\mathbb{P}(N < \infty) > \delta$ on a problem instance $(\alpha, \theta_0, \theta_1)$ when $\alpha = 0$. Then there exists a finite $\hat{n} \in \mathbb{N}$ such that $\mathbb{P}(N \leq \hat{n}) > \delta$. For some $\epsilon \in (0, 1)$ to be defined later, define $\hat{\alpha} = \frac{\log(\frac{1}{1-\epsilon})}{2\hat{n}}$ and note that for this $\hat{\alpha}$, $\mathbb{P}(\bigcap_{i=1}^{\hat{n}} \{\xi_i = 0\}) = (1 - \hat{\alpha})^{\hat{n}} \geq e^{-2\hat{n}\hat{\alpha}} \geq 1 - \epsilon$. Thus, the probability that the procedure terminates with a light distribution under $\alpha = \hat{\alpha}$ is at least

$$\begin{aligned} \mathbb{P}_{\alpha=\hat{\alpha}}(N \leq \hat{n}, \bigcap_{i=1}^{\hat{n}} \{\xi_i = 0\}) &= \mathbb{P}_{\alpha=\hat{\alpha}}(N \leq \hat{n} | \bigcap_{i=1}^{\hat{n}} \{\xi_i = 0\}) \mathbb{P}_{\alpha=\hat{\alpha}}(\bigcap_{i=1}^{\hat{n}} \{\xi_i = 0\}) \\ &= \mathbb{P}_{\alpha=0}(N \leq \hat{n}) \mathbb{P}_{\alpha=\hat{\alpha}}(\bigcap_{i=1}^{\hat{n}} \{\xi_i = 0\}) > \delta(1 - \epsilon). \end{aligned}$$

Because we can make ϵ arbitrarily small, the above display implies that the procedure makes a mistake with probability at least δ , but this is a contradiction as the procedure is δ -probably correct. ■

A.2 Proof of Theorem 2

Proof First, let N be the number of distributions considered at the stopping time T . Note that $T \geq N$. By assumption the procedure satisfies $\mathbb{P}_1(N \geq n | \bigcap_{i=1}^{n-1} \{\xi_i = 0\}) \geq 1 - \delta$ for all $n \in \mathbb{N}$. And

$$\begin{aligned} \mathbb{P}_1(N \geq n) &\geq \mathbb{P}_1(N \geq n, \bigcap_{i=1}^{n-1} \{\xi_i = 0\}) = \mathbb{P}_1(N \geq n | \bigcap_{i=1}^{n-1} \{\xi_i = 0\}) \mathbb{P}_1(\bigcap_{i=1}^{n-1} \{\xi_i = 0\}) \\ &\geq (1 - \delta)(1 - \alpha)^{n-1} \end{aligned}$$

Thus, $\mathbb{E}_1[N] = \sum_{n=1}^{\infty} \mathbb{P}_1(N \geq n) \geq (1 - \delta) \sum_{n=1}^{\infty} (1 - \alpha)^{n-1} = \frac{1-\delta}{\alpha}$ which results in the first argument of the max.

Applying Theorem 2.38 of [17] we have

$$\mathbb{E}_1[N] \chi^2(\mathbb{P}_1 | \mathbb{P}_0) \stackrel{\text{Eqn. (2)}}{\geq} \mathbb{E}_1[N] KL(\mathbb{P}_1 | \mathbb{P}_0) \stackrel{\text{Thm. 2.38}}{\geq} \log\left(\frac{1}{\mathbb{P}_0(N < \infty)}\right) \stackrel{\text{assumption}}{\geq} \log\left(\frac{1}{\delta}\right),$$

which results in the second argument of the max.

If $\tilde{\Theta} = \{\theta_0\}$ then $\chi^2(\mathbb{P}_1 | \mathbb{P}_0) = \chi^2((1 - \alpha)f_{\theta_0} + \alpha f_{\theta_1} | f_{\theta_0})$ and

$$\chi^2((1 - \alpha)f_{\theta_0} + \alpha f_{\theta_1} | f_{\theta_0}) = \int \frac{((1 - \alpha)f_{\theta_0}(x) + \alpha f_{\theta_1}(x) - f_{\theta_0}(x))^2}{f_{\theta_0}(x)} dx = \alpha^2 \chi^2(f_{\theta_1} | f_{\theta_0})$$

Thus, $\mathbb{E}_1[N] \geq \frac{\log(\frac{1}{\delta})}{\alpha^2 \chi^2(f_{\theta_1} | f_{\theta_0})}$ which results in the second part of the theorem. ■

A.3 Proof of Corollary 1

Proof For $k = 0, 1$ let g_{θ_k} be a Bernoulli distribution with parameter θ_k and let $f_{\theta_k} = g_{\theta_k} \otimes \cdots \otimes g_{\theta_k}$ be a product distribution composed of m g_{θ_k} distributions. Then

$$\chi^2(g_{\theta_1} | g_{\theta_0}) = \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)} \leq e^{\frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}} - 1$$

and

$$\chi^2(f_{\theta_1} | f_{\theta_0}) = (1 + \chi^2(g_{\theta_1} | g_{\theta_0}))^m - 1 \leq e^{m \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}} - 1.$$

Moreover, $e^{m \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}} - 1 \leq m \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}$ whenever $m \leq \frac{\theta_0(1 - \theta_0)}{2(\theta_1 - \theta_0)^2}$ since $e^{x/2} - 1 \leq x$ for all $x \in [0, 1]$. Applying Theorem 2 obtains

$$\mathbb{E}[N_m] \geq \max \left\{ \frac{1 - \delta}{\alpha}, \frac{\log(\frac{1}{\delta})}{\alpha^2 (e^{m \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}} - 1)} \right\} \geq \frac{\theta_0(1 - \theta_0) \log(\frac{1}{\delta})}{m \alpha^2 (\theta_1 - \theta_0)^2} \mathbf{1}_{m \leq \frac{\theta_0(1 - \theta_0)}{2(\theta_1 - \theta_0)^2}}.$$

The claimed result follows from loosening the integer constraint on m and minimizing the lower bound on $\mathbb{E}[N_m]$ multiplied by m . To perform the minimization, we note that the function $\max\{\frac{1-\delta}{\alpha}, 2 \log(\frac{1}{\delta}) / [\alpha^2 (e^{m \frac{(\theta_1 - \theta_0)^2}{\theta_0(1 - \theta_0)}} - 1)]\}$ reaches its minimum at the intersection of the two arguments and solve for m at that point. ■

A.4 Proof of Theorem 3

We begin by restating the theorem with the problem dependent parameters explicitly defined.

Theorem 6 Suppose f_θ for $\theta \in \Theta \subset \mathbb{R}$ is a single parameter exponential family so that $f_\theta(x) = h(x) \exp(\eta(\theta)x - b(\eta(\theta)))$ for some scalar functions h, b, η where η is strictly increasing. If $\mathbb{E}_\theta[X] = \int x f_\theta(x) dx$ then let $M_k(\theta) = \int (x - \mathbb{E}_\theta[X])^k f_\theta(x) dx$ denote the k th centered moment under distribution f_θ . Define

$$\begin{aligned}\theta_* &= \eta^{-1}((1 - \alpha)\eta(\theta_0) + \alpha\eta(\theta_1)) \\ \theta_- &= \eta^{-1}(\eta(\theta_0) - \alpha(\eta(\theta_1) - \eta(\theta_0))) \\ \theta_+ &= \eta^{-1}(\eta(\theta_1) + (1 - \alpha)(\eta(\theta_1) - \eta(\theta_0)))\end{aligned}$$

and assume there exist finite κ, γ such that

$$\sup_{y \in [\theta_0, \theta_1]} b(2\eta(y) - \eta(\theta_*)) - [2b(\eta(y)) - b(\eta(\theta_*))] \leq \kappa, \quad \text{and} \quad \sup_{x \in [\dot{b}(\eta(\theta_-)), \dot{b}(\eta(\theta_+))]} \phi_x(\dot{b}^{-1}(x)) \leq \gamma,$$

where $\phi_x(\eta(\theta)) = f_\theta(x)$. Then

$$\chi^2((1 - \alpha)f_{\theta_0}(x) + \alpha f_{\theta_1}(x) | f_{\theta_*}(x)) \leq c \left(\frac{1}{2} \alpha(1 - \alpha)(\eta(\theta_1) - \eta(\theta_0))^2 \right)^2$$

where if $\Delta = \dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-))$

$$c = e^\kappa \left(\sup_{\theta \in [\theta_0, \theta_1]} M_2(\theta)^2 (2 + \gamma\Delta) + 8M_4(\theta_-) + 8M_4(\theta_+) + 16\Delta^4 + \frac{2}{5}\gamma\Delta^5 \right).$$

Thus, if $\tilde{\Theta} = \{\theta_*\}$ and N is the stopping time of any procedure that satisfies $\mathbb{P}_0(N < \infty) \leq \delta$ and $\mathbb{P}_1(\cup_{i=1}^N \{\xi_i = 1\}) \geq 1 - \delta$, then

$$\mathbb{E}_1[N] \geq \max \left\{ \frac{1 - \delta}{\alpha}, \frac{\log(\frac{1}{\delta})}{c \left(\frac{1}{2} \alpha(1 - \alpha)(\eta(\theta_1) - \eta(\theta_0))^2 \right)^2} \right\}.$$

Proof Define $\phi_x(\eta) = h(x) \exp(\eta x - b(\eta))$. By the properties of scalar exponential families, note that $b'(\eta)$ and $b''(\eta) \geq 0$ represent the mean and variance of the distribution. We deduce that b' is monotonically increasing. Define $\eta_0 = \eta(\theta_0)$, $\eta_1 = \eta(\theta_1)$, and $\mu = (1 - \alpha)\eta_0 + \alpha\eta_1$. Noting that

$$\chi^2((1 - \alpha)\phi_x(\eta_0) + \alpha\phi_x(\eta_1) | \phi_x(\mu)) = \int \phi_x(\mu) \left(\frac{(1 - \alpha)\phi_x(\eta_0) + \alpha\phi_x(\eta_1) - \phi_x(\mu)}{\phi_x(\mu)} \right)^2 dx$$

we will use a technique that was used in [16] to approximate the divergence between a single Gaussian distribution and a mixture of them. Essentially, we will take the Taylor series of each $\phi_x(\cdot)$ centered at μ and bound. We have

$$\begin{aligned}\phi_x(\eta) &= h(x) \exp(\eta x - b(\eta)) \\ \phi'_x(\eta) &= (x - b'(\eta))\phi_x(\eta) \\ \phi''_x(\eta) &= (-b''(\eta) + (x - b'(\eta))^2)\phi_x(\eta)\end{aligned}$$

so that

$$\phi_x(y) = \phi_x(\mu) \left[1 + (x - b'(\mu))(y - \mu) + \frac{1}{2}(-b''(\mu) + (x - b'(\mu))^2)(y - \mu)^2 \dots \right].$$

Noting that $(\eta_0 - \mu) = -\alpha(\eta_1 - \eta_0)$, $(\eta_1 - \mu) = (1 - \alpha)(\eta_1 - \eta_0)$, and $(1 - \alpha)\alpha^2 + \alpha(1 - \alpha)^2 = \alpha(1 - \alpha)$, we have

$$\begin{aligned}& \left| \frac{(1 - \alpha)\phi_x(\eta_0) + \alpha\phi_x(\eta_1) - \phi_x(\mu)}{\phi_x(\mu)} \right| \\ &= \left| \frac{\phi'_x(\mu)}{\phi_x(\mu)} [(1 - \alpha)(\eta_0 - \mu) + \alpha(\eta_1 - \mu)] + \frac{1}{2} \frac{\phi''_x(\mu)}{\phi_x(\mu)} [(1 - \alpha)(\eta_0 - \mu)^2 + \alpha(\eta_1 - \mu)^2] + \dots \right| \\ &= \left| \frac{1}{2} \frac{\phi''_x(\mu)}{\phi_x(\mu)} \alpha(1 - \alpha)(\eta_1 - \eta_0)^2 + \dots \right| \\ &\leq \sup_{z \in [\eta_0, \eta_1]} \frac{|\phi''_x(z)|}{\phi_x(\mu)} \frac{1}{2} \alpha(1 - \alpha)(\eta_1 - \eta_0)^2.\end{aligned}$$

Thus,

$$\begin{aligned}\chi^2((1-\alpha)\phi_x(\eta_0) + \alpha\phi_x(\eta_1)|\phi_x(\mu)) &= \int \phi_x(\mu) \left(\frac{(1-\alpha)\phi_x(\eta_0) + \alpha\phi_x(\eta_1) - \phi_x(\mu)}{\phi_x(\mu)} \right)^2 dx \\ &\leq \left(\frac{1}{2}\alpha(1-\alpha)(\eta_1 - \eta_0)^2 \right)^2 \int \sup_{z \in [\eta_0, \eta_1]} \frac{|\phi_x''(z)|^2}{\phi_x(\mu)^2} \phi_x(\mu) dx.\end{aligned}$$

By distributing the square and noting that $b''(\eta) \geq 0$, we have

$$\begin{aligned}\int \sup_{z \in [\eta_0, \eta_1]} \frac{|\phi_x''(z)|^2}{\phi_x(\mu)^2} \phi_x(\mu) dx &= \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 (-b''(z) + (x - b'(z))^2) \phi_x(\mu) dx \\ &\leq \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 b''(z)^2 \phi_x(\mu) dx + \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 (x - b'(z))^4 \phi_x(\mu) dx \\ &\leq \sup_{y \in [\eta_0, \eta_1]} b''(y)^2 \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 \phi_x(\mu) dx + \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 (x - b'(z))^4 \phi_x(\mu) dx.\end{aligned}$$

The remainder of the proof bounds the integrals. Define $\eta_- = 2\eta_0 - \mu = \eta(\theta_-)$ and $\eta_+ = 2\eta_1 - \mu = \eta(\theta_+)$. Observe that

$$\begin{aligned}&\sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 \phi_x(\mu) \\ &= \sup_{z \in [\eta_0, \eta_1]} h(x) \exp((2z - \mu)x - (2b(z) - b(\mu))) \\ &= \sup_{z \in [\eta_0, \eta_1]} h(x) \exp((2z - \mu)x - b(2z - \mu)) \exp(b(2z - \mu) - (2b(z) - b(\mu))) \\ &\leq e^\kappa \sup_{z \in [\eta_0, \eta_1]} h(x) \exp((2z - \mu)x - b(2z - \mu)) \\ &= e^\kappa \sup_{z \in [2\eta_0 - \mu, 2\eta_1 - \mu]} h(x) \exp(zx - b(z)) \\ &= e^\kappa \sup_{z \in [\eta_-, \eta_+]} h(x) \exp(zx - b(z)) \\ &\leq e^\kappa \left(\phi_x(\eta_-) + \phi_x(\eta_+) + \phi_x(\dot{b}^{-1}(x)) \mathbf{1}_{x \in [\dot{b}(\eta_-), \dot{b}(\eta_+)]} \right) \\ &\leq e^\kappa \left(\phi_x(\eta_-) + \phi_x(\eta_+) + \gamma \mathbf{1}_{x \in [\dot{b}(\eta_-), \dot{b}(\eta_+)]} \right)\end{aligned}$$

where the second inequality follows by observing that the maximum of the function $\phi_x(z)$ will occur either at an endpoint of the interval $z \in [\eta(\theta_-), \eta(\theta_+)]$ or at the point where $\frac{\partial}{\partial z} g(z) = 0$ (if that point occurs inside the interval), and loosely bounding the maximum by simply adding the function values at all three points.

Consequently,

$$\sup_{y \in [\eta_0, \eta_1]} b''(y)^2 \int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 \phi_x(\mu) dx \leq \sup_{\theta \in [\theta_0, \theta_1]} M_2(\theta)^2 e^\kappa \left(2 + \gamma(\dot{b}(\eta_+) - \dot{b}(\eta_-)) \right).$$

By Jensen's inequality, $(a + b)^4 = 16(\frac{1}{2}a + \frac{1}{2}b)^4 \leq 8(a^4 + b^4)$, so

$$\begin{aligned}\int \sup_{z \in [\eta_0, \eta_1]} \phi_x(\eta_-)(x - \dot{b}(z))^4 dx &= \int \sup_{z \in [\eta_0, \eta_1]} \phi_x(\eta_-)(x - \dot{b}(\eta_-) + \dot{b}(\eta_-) - \dot{b}(z))^4 dx \\ &\leq \int 8\phi_x(\eta_-)[(x - \dot{b}(\eta_-))^4 + \sup_{z \in [\eta_0, \eta_1]} (\dot{b}(\eta_-) - \dot{b}(z))^4] dx \\ &\leq \int 8\phi_x(\eta_-)[(x - \dot{b}(\eta_-))^4 + (\dot{b}(\eta_-) - \dot{b}(\eta_1))^4] dx \\ &= 8[M_4(\theta_-) - (\dot{b}(\eta_-) - \dot{b}(\eta_1))^4].\end{aligned}$$

Repeating an analogous series of steps for η_+ , we have

$$\begin{aligned}&\int \sup_{z \in [\eta_0, \eta_1]} \left(\frac{\phi_x(z)}{\phi_x(\mu)} \right)^2 (x - b'(z))^4 \phi_x(\mu) dx \\ &\leq e^\kappa \int \left(\phi_x(\eta_-) + \phi_x(\eta_+) + \gamma \mathbf{1}_{x \in [\dot{b}(\eta_-), \dot{b}(\eta_+)]} \right) \sup_{z \in [\eta_0, \eta_1]} (x - \dot{b}(z))^4 dx \\ &\leq e^\kappa \left(8M_4(\theta_-) + 8(\dot{b}(\eta_1) - \dot{b}(\eta_-))^4 + 8M_4(\theta_+) + 8(\dot{b}(\eta_+) - \dot{b}(\eta_0))^4 + \frac{2}{5}\gamma(\dot{b}(\eta_+) - \dot{b}(\eta_-))^5 \right) \\ &\leq e^\kappa \left(8M_4(\theta_-) + 8M_4(\theta_+) + 16(\dot{b}(\eta_+) - \dot{b}(\eta_-))^4 + \frac{2}{5}\gamma(\dot{b}(\eta_+) - \dot{b}(\eta_-))^5 \right).\end{aligned}$$

The final result holds by Theorem 2.

A.5 Proof of Corollary 2

Proof A binomial distribution for fixed m is an exponential family $f_\theta(x) = h(x) \exp(\eta(\theta)x - b(\eta(\theta)))$ with $h(x) = \binom{m}{x}$, $\eta(\theta) = \log(\frac{\theta}{1-\theta})$, and $b(\tau) = m \log(1 + e^\tau)$. Note that η is monotonically increasing, b is m -Lipschitz, and $\dot{b}(\tau) = m(1 + e^{-\tau})^{-1}$ so that $\dot{b}(\eta(\theta)) = m\theta$.

Step 1: Relating θ_+, θ_- to θ_1, θ_0

We will make repeated use of the fact that if f is convex then $f(y) \geq f(x) + f'(x)^T(y - x)$. Since $\frac{x}{1-x}$ and $\frac{1-x}{x}$ are both convex, we have

$$\frac{y}{1-y} \geq \frac{x}{1-x} + \frac{y-x}{(1-x)^2} \quad \text{and} \quad \frac{1-y}{y} \geq \frac{1-x}{x} - \frac{y-x}{x^2}$$

for all $x, y \in [0, 1]$.

To begin, note $\eta^{-1}(\nu) = (1 + e^{-\nu})^{-1}$ so that for any θ we have $\theta(1 - \theta) = \eta^{-1}(\eta(\theta))(1 - \eta^{-1}(\eta(\theta))) = \frac{e^{-\eta(\theta)}}{(1 + e^{-\eta(\theta)})^2}$. Observe that

$$\frac{1}{4}e^{-|\eta(\theta)|} \leq \frac{e^{-\eta(\theta)}}{(1 + e^{-\eta(\theta)})^2} \leq e^{-|\eta(\theta)|}$$

and recalling that $\theta_* = \eta^{-1}((1 - \alpha)\theta_0 + \alpha\theta_1) \in [\theta_0, \theta_1]$ we have

$$\begin{aligned} \theta_+(1 - \theta_+) &\geq \frac{1}{4}e^{-|\eta(\theta_+)|} = \frac{1}{4}e^{-|2\eta(\theta_1) - \eta(\theta_*)|} \\ &= \frac{1}{4}\mathbf{1}_{\theta_+ \leq 1/2} \left(\frac{\theta_1}{1 - \theta_1} \right)^2 \left(\frac{1 - \theta_*}{\theta_*} \right) + \frac{1}{4}\mathbf{1}_{\theta_+ > 1/2} \left(\frac{1 - \theta_1}{\theta_1} \right)^2 \left(\frac{\theta_*}{1 - \theta_*} \right) \\ &\geq \frac{1}{4}\mathbf{1}_{\theta_+ \leq 1/2} \left(\frac{\theta_1}{1 - \theta_1} \right)^2 \left(\frac{1 - \theta_1}{\theta_1} \right) + \frac{1}{4}\mathbf{1}_{\theta_+ > 1/2} \left(\frac{1 - \theta_1}{\theta_1} \right)^2 \left(\frac{\theta_0}{1 - \theta_0} \right) \\ &\geq \frac{1}{4}\mathbf{1}_{\theta_+ \leq 1/2} \left(\frac{\theta_1}{1 - \theta_1} \right) + \frac{1}{4}\mathbf{1}_{\theta_+ > 1/2} \left(\frac{1 - \theta_1}{\theta_1} \right)^2 \left(\frac{\theta_1}{1 - \theta_1} - \frac{\theta_1 - \theta_0}{(1 - \theta_1)^2} \right) \\ &\geq \frac{1}{4}\mathbf{1}_{\theta_+ \leq 1/2} \left(\frac{\theta_1}{1 - \theta_1} \right) + \frac{1}{8}\mathbf{1}_{\theta_+ > 1/2} \left(\frac{1 - \theta_1}{\theta_1} \right) \geq \frac{1}{8}\theta_1(1 - \theta_1) \end{aligned}$$

where the last line follows from the assumption that $\theta_1(1 - \theta_1) \geq 2(\theta_1 - \theta_0)$. Analogously,

$$\begin{aligned} \theta_-(1 - \theta_-) &\geq \frac{1}{4}e^{-|\eta(\theta_-)|} = \frac{1}{4}e^{-|2\eta(\theta_0) - \eta(\theta_*)|} \\ &= \frac{1}{4}\mathbf{1}_{\theta_- \leq 1/2} \left(\frac{\theta_0}{1 - \theta_0} \right)^2 \left(\frac{1 - \theta_*}{\theta_*} \right) + \frac{1}{4}\mathbf{1}_{\theta_- > 1/2} \left(\frac{1 - \theta_0}{\theta_0} \right)^2 \left(\frac{\theta_*}{1 - \theta_*} \right) \\ &\geq \frac{1}{4}\mathbf{1}_{\theta_- \leq 1/2} \left(\frac{\theta_0}{1 - \theta_0} \right)^2 \left(\frac{1 - \theta_1}{\theta_1} \right) + \frac{1}{4}\mathbf{1}_{\theta_- > 1/2} \left(\frac{1 - \theta_0}{\theta_0} \right)^2 \left(\frac{\theta_0}{1 - \theta_0} \right) \\ &\geq \frac{1}{4}\mathbf{1}_{\theta_- \leq 1/2} \left(\frac{\theta_0}{1 - \theta_0} \right)^2 \left(\frac{1 - \theta_0}{\theta_0} - \frac{\theta_1 - \theta_0}{\theta_0^2} \right) + \frac{1}{4}\mathbf{1}_{\theta_- > 1/2} \left(\frac{1 - \theta_0}{\theta_0} \right) \\ &\geq \frac{1}{8}\mathbf{1}_{\theta_- \leq 1/2} \left(\frac{\theta_0}{1 - \theta_0} \right) + \frac{1}{4}\mathbf{1}_{\theta_- > 1/2} \left(\frac{1 - \theta_0}{\theta_0} \right) \geq \frac{1}{8}\theta_0(1 - \theta_0) \end{aligned}$$

where the last line follows from the assumption that $\theta_0(1 - \theta_0) \geq 2(\theta_1 - \theta_0)$. We conclude that

$$\inf_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \geq \frac{1}{8} \inf_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta). \quad (4)$$

Conversely,

$$\sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \leq \mathbf{1}_{1/2 \in [\theta_-, \theta_+]} \frac{1}{4} + \theta_+(1 - \theta_+)\mathbf{1}_{\theta_+ \leq 1/2} + \theta_-(1 - \theta_-)\mathbf{1}_{\theta_- > 1/2}.$$

We consider these three cases in turn. If $\theta_+ \leq 1/2$:

$$\begin{aligned}
\theta_+(1 - \theta_+) &\leq e^{-|\eta(\theta_+)|} = e^{-|2\eta(\theta_1) - \eta(\theta_*)|} \\
&= \left(\frac{\theta_1}{1 - \theta_1}\right)^2 \left(\frac{1 - \theta_*}{\theta_*}\right) \leq \left(\frac{\theta_1}{1 - \theta_1}\right)^2 \left(\frac{1 - \theta_0}{\theta_0}\right) \leq \left(\frac{\theta_1}{1 - \theta_1}\right)^2 \left(\frac{1 - \theta_1}{\theta_1} + \frac{\theta_1 - \theta_0}{\theta_0^2}\right) \\
&= \left(\frac{\theta_1}{1 - \theta_1}\right) \left(1 + \frac{\theta_1(\theta_1 - \theta_0)}{(1 - \theta_1)\theta_0^2}\right) \leq \left(\frac{\theta_1}{1 - \theta_1}\right) \left(1 + \frac{\theta_1(1 - \theta_0)}{2(1 - \theta_1)\theta_0}\right) \\
&= \left(\frac{\theta_1}{1 - \theta_1}\right) \left(1 + \frac{\theta_0(1 - \theta_0) + (\theta_1 - \theta_0)(1 - \theta_0)}{2(1 - \theta_1)\theta_0}\right) \\
&\leq \left(\frac{\theta_1}{1 - \theta_1}\right) \left(1 + \frac{\theta_0(1 - \theta_0) + \theta_0(1 - \theta_0)^2/2}{2(1 - \theta_1)\theta_0}\right) \leq \frac{5}{2} \left(\frac{\theta_1}{1 - \theta_1}\right) \leq 10\theta_1(1 - \theta_1)
\end{aligned}$$

using the convexity of $\frac{1-x}{x}$, the assumption that $2(\theta_1 - \theta_0) \leq \theta_0(1 - \theta_0)$, that $\theta_1 \leq \theta_+ \leq 1/2$, and that $1 - \theta_0 \leq 1$. If $\theta_- > 1/2$:

$$\begin{aligned}
\theta_-(1 - \theta_-) &\leq e^{-|\eta(\theta_-)|} = e^{-|2\eta(\theta_0) - \eta(\theta_*)|} \\
&= \left(\frac{1 - \theta_0}{\theta_0}\right)^2 \left(\frac{\theta_*}{1 - \theta_*}\right) \leq \left(\frac{1 - \theta_0}{\theta_0}\right)^2 \left(\frac{\theta_1}{1 - \theta_1}\right) \leq \left(\frac{1 - \theta_0}{\theta_0}\right)^2 \left(\frac{\theta_0}{1 - \theta_0} + \frac{\theta_1 - \theta_0}{(1 - \theta_1)^2}\right) \\
&\leq \left(\frac{1 - \theta_0}{\theta_0}\right) \left(1 + \frac{(1 - \theta_0)(\theta_1 - \theta_0)}{\theta_0(1 - \theta_1)^2}\right) \leq \left(\frac{1 - \theta_0}{\theta_0}\right) \left(1 + \frac{(1 - \theta_0)\theta_1/2}{\theta_0(1 - \theta_1)}\right) \\
&= \left(\frac{1 - \theta_0}{\theta_0}\right) \left(1 + \frac{(1 - \theta_1)\theta_1 + (\theta_1 - \theta_0)\theta_1}{2\theta_0(1 - \theta_1)}\right) \\
&\leq \left(\frac{1 - \theta_0}{\theta_0}\right) \left(1 + \frac{(1 - \theta_1)\theta_1 + (1 - \theta_1)\theta_1^2/2}{2\theta_0(1 - \theta_1)}\right) \leq \frac{5}{2} \left(\frac{1 - \theta_0}{\theta_0}\right) \leq 10\theta_0(1 - \theta_0)
\end{aligned}$$

using the same methods as above. From these two cases, we can conclude that if $1/2 \notin [\theta_-, \theta_+]$,

$$\sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \leq 10 \sup_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta). \quad (5)$$

The remaining case, when $1/2 \in [\theta_-, \theta_+]$, also satisfies (5), which we now demonstrate. When $\theta_+ = 1/2$ we have $1/4 = \theta_+(1 - \theta_+) \leq 10\theta_1(1 - \theta_1)$ so that $\theta_1(1 - \theta_1) \geq 1/40$. Because θ_1 is monotonically increasing in θ_+ and $\sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \leq 1/4$ we conclude that (5) holds whenever $\theta_1 \leq 1/2$. A similar argument follows for all $\theta_0 \geq 1/2$. Finally, if $1/2 \in [\theta_0, \theta_1]$, it must be true that $\sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \leq \sup_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta)$ because $\theta_- \leq \theta_0 \leq \frac{1}{2} \leq \theta_1 \leq \theta_+$ and the function $\theta(1 - \theta)$ is concave taking its maximum at $\frac{1}{2}$. Thus, (5) holds for all θ_-, θ_+ .

We now turn our attention to bounding $\theta_+ - \theta_-$. Let $g(y) = \eta^{-1}(y)$ then $g(y) = (1 + e^{-y})^{-1}$ and $\dot{g}(y) = e^{-y}(1 + e^{-y})^{-2}$. Observing that $\dot{g}(\eta(\theta)) = \theta(1 - \theta)$ we have by Taylor's remainder theorem

$$\begin{aligned}
\theta_+ - \theta_- &= \eta^{-1}(\eta(\theta_+)) - \eta^{-1}(\eta(\theta_-)) \leq (\eta(\theta_+) - \eta(\theta_-)) \sup_{y \in [\eta(\theta_-), \eta(\theta_+)]} e^{-y}(1 + e^{-y})^{-2} \\
&= (\eta(\theta_+) - \eta(\theta_-)) \sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) = 2(\eta(\theta_1) - \eta(\theta_0)) \sup_{\theta \in [\theta_-, \theta_+]} \theta(1 - \theta) \\
&\leq 20(\eta(\theta_1) - \eta(\theta_0)) \sup_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta).
\end{aligned}$$

Since $\eta(\theta) = \log\left(\frac{\theta}{1-\theta}\right)$ and $\eta'(\theta) = \frac{1}{\theta} + \frac{1}{1-\theta} = \frac{1}{\theta(1-\theta)}$, we have

$$\theta_+ - \theta_- \leq 20(\eta(\theta_1) - \eta(\theta_0)) \sup_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta) \leq 20(\theta_1 - \theta_0) \frac{\sup_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta)}{\inf_{\theta \in [\theta_0, \theta_1]} \theta(1 - \theta)}.$$

If $\theta_1(1 - \theta_1) \geq \theta_0(1 - \theta_0)$:

$$\begin{aligned}
\frac{\theta_1(1 - \theta_1)}{\theta_0(1 - \theta_0)} &= \frac{\theta_0(1 - \theta_1) + (\theta_1 - \theta_0)(1 - \theta_1)}{\theta_0(1 - \theta_0)} \\
&\leq \frac{\theta_0(1 - \theta_1) + \theta_0(1 - \theta_0)(1 - \theta_1)/2}{\theta_0(1 - \theta_0)} \leq 1 + (1 - \theta_1)/2 \leq 3/2,
\end{aligned}$$

else if $\theta_0(1 - \theta_0) \geq \theta_1(1 - \theta_1)$

$$\begin{aligned}
\frac{\theta_0(1 - \theta_0)}{\theta_1(1 - \theta_1)} &= \frac{\theta_0(1 - \theta_1) + \theta_0(\theta_1 - \theta_0)}{\theta_1(1 - \theta_1)} \\
&\leq \frac{\theta_0(1 - \theta_1) + \theta_0\theta_1(1 - \theta_1)/2}{\theta_1(1 - \theta_1)} \leq 1 + \theta_0/2 \leq 3/2.
\end{aligned}$$

Finally, if $1/2 \in [\theta_0, \theta_1]$ then $\sup_{\theta \in [\theta_0, \theta_1]} \theta(1-\theta) = 1/4$ taking its maximum at $1/4$. To maximize the ratio of the sup to the inf, it suffices to just consider the case when $\theta_0 = 1/2$ or $\theta_1 = 1/2$. Thus, the above two bounds suffice for this case and we observe that

$$\frac{\sup_{\theta \in [\theta_0, \theta_1]} \theta(1-\theta)}{\inf_{\theta \in [\theta_0, \theta_1]} \theta(1-\theta)} \leq 3/2. \quad (6)$$

Thus, putting the pieces together, we conclude that

$$\theta_+ - \theta_- \leq 30(\theta_1 - \theta_0). \quad (7)$$

Step 2: Bounding γ, κ, c

In what follows, define $\theta_h = \arg \sup_{\theta \in [\theta_0, \theta_1]} \theta(1-\theta)$ and $\theta_l = \arg \inf_{\theta \in [\theta_0, \theta_1]} \theta(1-\theta)$. We now continue to bound the terms of the theorem. Note

$$\begin{aligned} \sup_{x \in [\dot{b}(\eta(\theta_-)), \dot{b}(\eta(\theta_+))]} \phi_x(\dot{b}^{-1}(x)) &= \sup_{x \in [m\theta_-, m\theta_+]} \phi_x(\eta(x/m)) \\ &\leq \sup_{x \in [m\theta_-, m\theta_+]} \sup_{y \in [0,1]} \phi_x(\eta(y)) \\ &= \sup_{x \in [m\theta_-, m\theta_+]} \sup_{y \in [0,1]} \frac{\Gamma(m+1)}{\Gamma(m-x+1)\Gamma(x+1)} y^x (1-y)^{m-x} \\ &= \sup_{\theta \in [\theta_-, \theta_+]} \sup_{y \in [0,1]} \frac{\Gamma(m+1)}{\Gamma(m(1-\theta)+1)\Gamma(m\theta+1)} y^{m\theta} (1-y)^{m(1-\theta)} \\ &\leq \sup_{\theta \in [\theta_-, \theta_+]} \sup_{y \in [0,1]} \frac{e/2\pi}{\sqrt{m\theta(1-\theta)}} \frac{y^{m\theta} (1-y)^{m(1-\theta)}}{\theta^{m\theta} (1-\theta)^{m(1-\theta)}} \\ &= \sup_{\theta \in [\theta_-, \theta_+]} \frac{e/2\pi}{\sqrt{m\theta(1-\theta)}} \leq \frac{2}{\sqrt{m\theta_l(1-\theta_l)}} =: \gamma \end{aligned}$$

by Stirling's approximation: $\sqrt{2\pi} \leq \frac{\Gamma(s+1)}{e^{-s} s^{s+1/2}} \leq e$ [18] and (4). And for any $y \in [\theta_0, \theta_1]$

$$\begin{aligned} &b(2\eta(y) - \eta(\theta_*)) - (2b(\eta(y)) - b(\eta(\theta_*))) \\ &= m \log(1 + e^{2\eta(y) - \eta(\theta_*)}) - 2m \log(1 + e^{\eta(y)}) + m \log(1 + e^{\eta(\theta_*)}) \\ &= m \log \left(\frac{(1 + e^{2\eta(y) - \eta(\theta_*)})(1 + e^{\eta(\theta_*)})}{(1 + e^{\eta(y)})^2} \right) \\ &= m \log \left(\left(1 + \left(\frac{y}{1-y} \right)^2 \frac{1-\theta_*}{\theta_*} \right) \left(\frac{1}{1-\theta_*} \right) (1-y)^2 \right) \\ &= m \log \left((1-y)^2 \frac{1}{1-\theta_*} + y^2 \frac{1}{\theta_*} \right) \\ &= m \log \left((1-2y+y^2) \frac{\theta_*}{\theta_*(1-\theta_*)} + y^2 \frac{1-\theta_*}{\theta_*(1-\theta_*)} \right) \\ &= m \log \left((1-2y) \frac{\theta_*}{\theta_*(1-\theta_*)} + y^2 \frac{1}{\theta_*(1-\theta_*)} \right) \\ &= m \log \left(1 + \frac{(y-\theta_*)^2}{\theta_*(1-\theta_*)} \right) \end{aligned}$$

so

$$\begin{aligned} &\sup_{y \in [\theta_0, \theta_1]} b(2\eta(y) - \eta(\theta_*)) - (2b(\eta(y)) - b(\eta(\theta_*))) \\ &\leq \sup_{y \in [\theta_0, \theta_1]} m \log \left(1 + \frac{(y-\theta_*)^2}{\theta_*(1-\theta_*)} \right) \leq m \left(\frac{(\theta_1 - \theta_0)^2}{\theta_*(1-\theta_*)} \right) =: \kappa. \end{aligned}$$

Noting that $M_2(\theta) = m\theta(1-\theta)$,

$$\begin{aligned} &\sup_{y \in [\theta_0, \theta_1]} M_2(y)^2 (2 + \gamma(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)))) \leq m^2 (\theta_h(1-\theta_h))^2 (2 + \gamma m(\theta_+ - \theta_-)) \\ &\leq m^2 (\theta_h(1-\theta_h))^2 \left(2 + \frac{2m}{\sqrt{m\theta_l(1-\theta_l)}} 30(\theta_1 - \theta_0) \right) \\ &\leq m^2 (\theta_h(1-\theta_h))^2 \left(2 + 60 \sqrt{m \frac{(\theta_1 - \theta_0)^2}{\theta_l(1-\theta_l)}} \right). \end{aligned}$$

Since for any $\theta \in [0, 1]$

$$M_4(\theta) = m\theta(1-\theta)(3\theta(1-\theta)(m-2)+1) < 3m^2(\theta(1-\theta))^2 + m\theta(1-\theta),$$

we have

$$\begin{aligned} & 8M_4(\theta_-) + 8M_4(\theta_+) + 16 \left(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)) \right)^4 + \frac{2}{5}\gamma \left(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)) \right)^5 \\ & \leq 24m^2(\theta_-(1-\theta_-))^2 + 8m\theta_-(1-\theta_-) + 24m^2(\theta_+(1-\theta_+))^2 + 8m\theta_+(1-\theta_+) \\ & \quad + 16m^4(\theta_+ - \theta_-)^4 + \frac{4/5}{\sqrt{m\theta_i(1-\theta_i)}} m^5(\theta_+ - \theta_-)^5 \\ & \leq 4800m^2(\theta_h(1-\theta_h))^2 + 160m\theta_h(1-\theta_h) \\ & \quad + 3240000m^4(\theta_1 - \theta_0)^4 + 19440000\sqrt{m\frac{(\theta_1 - \theta_0)^2}{\theta_i(1-\theta_i)}} m^4(\theta_1 - \theta_0)^4 \end{aligned}$$

where we have applied (5) and (7). Finally, recall from above that

$$\eta(\theta_1) - \eta(\theta_0) \leq \frac{\theta_1 - \theta_0}{\theta_i(1-\theta_i)} \leq \frac{3}{2} \frac{\theta_1 - \theta_0}{\theta_*(1-\theta_*)}.$$

Step 3: Putting the pieces together

Noting that $\theta_i(1-\theta_i) \leq \theta_*(1-\theta_*) \leq \theta_h(1-\theta_h)$ and $\frac{\theta_h(1-\theta_h)}{\theta_i(1-\theta_i)} \leq 3/2$ by (6), we can use $\theta_*(1-\theta_*)$ throughout at the cost of a constant. Putting it altogether, if $m\frac{(\theta_1 - \theta_0)^2}{\theta_*(1-\theta_*)} \leq 1$ then $\kappa \leq 1$ and

$$\begin{aligned} c & \leq c' \left(m^2(\theta_*(1-\theta_*))^2 + m\theta_*(1-\theta_*) + m^4(\theta_1 - \theta_0)^4 \right) \\ & \leq c' \left(m^2(\theta_*(1-\theta_*))^2 + m\theta_*(1-\theta_*) \right) \end{aligned}$$

for some absolute constant c' . Thus,

$$\begin{aligned} & c \left(\frac{1}{2}\alpha(1-\alpha)(\eta(\theta_1) - \eta(\theta_0))^2 \right)^2 \\ & \leq c' \left(m^2(\theta_*(1-\theta_*))^2 + m\theta_*(1-\theta_*) \right) \left(\frac{9}{8}\alpha(1-\alpha)\frac{(\theta_1 - \theta_0)^2}{(\theta_*(1-\theta_*))^2} \right)^2 \\ & \leq c' \left(m^2 + \frac{m}{\theta_*(1-\theta_*)} \right) \left(\frac{9}{8}\alpha(1-\alpha)\frac{(\theta_1 - \theta_0)^2}{\theta_*(1-\theta_*)} \right)^2 \\ & \leq 2c'm \left(\min\left\{ \frac{1}{m}, \theta_*(1-\theta_*) \right\} \right)^{-1} \left(\frac{9}{8}\alpha(1-\alpha)\frac{(\theta_1 - \theta_0)^2}{\theta_*(1-\theta_*)} \right)^2. \end{aligned}$$

■

B Proofs of Upper Bounds

B.1 Fixed sample size strategy

Consider a fixed sample strategy that knows $\alpha, \theta_0, \theta_1$ and flips each coin exactly m times until it declares a coin as heavy (the algorithm can use the parameters to choose m). Note the result below is within a $\log(1/\delta)$ factor of the lower bound proved in Corollary 1 in general and is tight when $\alpha \leq \delta$.

Theorem 7 (Fixed sample size, known $\alpha, \theta_0, \theta_1$) Fix $\delta \in (0, 1/4)$ and set $\hat{n} = \lceil \frac{1}{\alpha} \log(\frac{2}{\delta}) \rceil$ and $m = \lceil \frac{2 \log(4\hat{n}/\delta)}{(\theta_1 - \theta_0)^2} \rceil$. There exists a fixed sample size strategy with stopping time $N_m \leq \hat{n}$ that is δ -probably correct and satisfies

$$\mathbb{E}[mN_m] \leq 3 \frac{\log(1/\alpha) + \log(12 \log(6/\delta)/\delta)}{\alpha(\theta_1 - \theta_0)^2} \leq 12 \frac{\log(\frac{2}{\delta\alpha})}{\alpha(\theta_1 - \theta_0)^2}.$$

Proof Let $\hat{\mu}_i$ be the empirical mean of the i th distribution sampled m times with mean $\mu_i \in \{\theta_0, \theta_1\}$. Let N be the minimum of \hat{n} and the first $i \in \mathbb{N}$ such that $\hat{\mu}_i \geq \frac{\theta_0 + \theta_1}{2}$. Declare distribution N to be heavy. The total number of flips this procedure makes equals mN .

Define the events

$$\xi_1 = \bigcup_{i=1}^{\hat{n}} \{\mu_i = \theta_1\}, \quad \text{and} \quad \xi_2 = \bigcap_{i=1}^{\hat{n}} \{|\hat{\mu}_i - \mu_i| < \frac{\theta_1 - \theta_0}{2}\}.$$

Note that $\mathbb{P}(\xi_1^c) = \mathbb{P}(\mu_1 = \theta_0)^{\widehat{n}} = (1 - \alpha)^{\widehat{n}} \leq \exp(-\alpha\widehat{n}) \leq \delta/2$. And, by a union bound and Chernoff's inequality $\mathbb{P}(\xi_2^c) \leq 2\widehat{n}e^{-m(\theta_1 - \theta_0)^2/2} \leq \delta/2$. Thus, the probability that ξ_1 or ξ_2 fail to occur is less than δ , so in what follows assume they succeed.

Under ξ_1 at least one of the \widehat{n} distributions is heavy. Under ξ_2 , for any $i \in [\widehat{n}]$ with $\mu_i = \theta_0$ we have $\widehat{\mu}_i < \mu_i + \frac{\theta_1 - \theta_0}{2} = \frac{\theta_0 + \theta_1}{2}$ which implies that the procedure will never exit with a light distribution unless $N = \widehat{n}$. On the other hand, for the first $i \in [\widehat{n}]$ with $\mu_i = \theta_1$ we have $\widehat{\mu}_i > \mu_i - \frac{\theta_1 - \theta_0}{2} = \frac{\theta_0 + \theta_1}{2}$ which means the algorithm will output distribution i at time $N = i$. Thus, N is equal to the first distribution that is heavy and

$$\begin{aligned} \mathbb{E}[N] &= \sum_{n=1}^{\widehat{n}} \mathbb{P}(N \geq n) = \sum_{n=1}^{\widehat{n}} \mathbb{P}(N \geq n, \max_{i=1, \dots, n-1} \mu_i \neq \theta_1) + \mathbb{P}(N \geq n, \max_{i=1, \dots, n-1} \mu_i = \theta_1) \\ &\leq \sum_{n=1}^{\widehat{n}} \mathbb{P}(\max_{i=1, \dots, n-1} \mu_i \neq \theta_1) + \mathbb{P}(\cup_{i=1}^{n-1} \{|\widehat{\mu}_i - \mu_i| > \frac{\theta_1 - \theta_0}{2}\} | \max_{i=1, \dots, n-1} \mu_i = \theta_1) \mathbb{P}(\max_{i=1, \dots, n-1} \mu_i = \theta_1) \\ &\leq \sum_{n=1}^{\widehat{n}} \mathbb{P}(\max_{i=1, \dots, n-1} \mu_i \neq \theta_1) + \mathbb{P}(\cup_{i=1}^{n-1} \{|\widehat{\mu}_i - \mu_i| > \frac{\theta_1 - \theta_0}{2}\}) \\ &\leq \sum_{n=1}^{\widehat{n}} (1 - \alpha)^{n-1} + \frac{n-1}{\widehat{n}} \frac{\delta}{2} \leq \frac{1}{\alpha} + \widehat{n}\delta/4 = \frac{1}{\alpha} (1 + \frac{\delta \log(2e/\delta)}{4}) \leq \frac{3/2}{\alpha}. \end{aligned}$$

Multiplying $\mathbb{E}[N]$ by m yields the result. ■

B.2 Proof of Theorem 4

First, we prove several technical lemmas necessary to analyze our algorithm.

Lemma 1 For $i \in \mathbb{N}$, let $X_i \in [a_i, b_i]$ for $|b_i - a_i| \leq 1$ be a random variable with $\mathbb{E}[X_i] = 0$. Then

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} \left\{ \sum_{i=1}^n X_i \geq \alpha n + \beta \right\}\right) \leq 7 \exp(-\alpha\beta/2)$$

whenever $\alpha\beta \geq 1$.

Proof First we will break the bound into two pieces:

$$\mathbb{P}\left(\bigcup_{n=1}^{\infty} \left\{ \sum_{i=1}^n X_i \geq \alpha n + \beta \right\}\right) \leq \min_{n_0} \mathbb{P}\left(\bigcup_{n=1}^{n_0} \left\{ \sum_{i=1}^n X_i \geq \beta \right\}\right) + \mathbb{P}\left(\bigcup_{n=n_0+1}^{\infty} \left\{ \sum_{i=1}^n X_i \geq \alpha n \right\}\right)$$

where $\mathbb{P}\left(\bigcup_{n=1}^{n_0} \left\{ \sum_{i=1}^n X_i \geq \beta \right\}\right) \leq \exp(-2\beta^2/n_0)$ by Doob-Hoeffding's maximal inequality. For any fixed $k \in \mathbb{N}$:

$$\mathbb{P}\left(\sum_{i=1}^{2^k} X_i \geq \alpha 2^k / 2\right) \leq \exp(-\alpha^2 2^k / 2)$$

and

$$\begin{aligned} \mathbb{P}\left(\bigcup_{n=2^k+1}^{2^{k+1}} \left\{ \sum_{i=2^k+1}^n X_i \geq \alpha n / 2 \right\}\right) &\leq \mathbb{P}\left(\bigcup_{n=2^k+1}^{2^{k+1}} \left\{ \sum_{i=2^k+1}^n X_i \geq \alpha 2^k / 2 \right\}\right) \\ &= \mathbb{P}\left(\bigcup_{\ell=1}^{2^k} \left\{ \sum_{i=1}^{\ell} X_i \geq \alpha 2^k / 2 \right\}\right) \leq \exp(-\alpha^2 2^k / 2) \end{aligned}$$

by Hoeffding's and Doob-Hoeffding's maximal inequality, respectively. Thus

$$\begin{aligned}
\mathbb{P} \left(\bigcup_{n=n_0}^{\infty} \left\{ \sum_{i=1}^n X_i \geq \alpha n \right\} \right) &= \mathbb{P} \left(\bigcup_{n=n_0}^{\infty} \left\{ \sum_{i=1}^{2^{\lceil \log_2(n) \rceil}} X_i + \sum_{i=2^{\lceil \log_2(n) \rceil}+1}^n X_i \geq \alpha n \right\} \right) \\
&= \mathbb{P} \left(\bigcup_{k=\log_2(n_0)}^{\infty} \bigcup_{n=2^{k+1}}^{2^{k+1}} \left\{ \sum_{i=1}^{2^k} X_i + \sum_{i=2^k+1}^n X_i \geq \alpha n \right\} \right) \\
&\leq \sum_{k=\log_2(n_0)}^{\infty} \mathbb{P} \left(\sum_{i=1}^{2^k} X_i \geq \alpha 2^k / 2 \right) + \mathbb{P} \left(\bigcup_{n=2^{k+1}}^{2^{k+1}} \left\{ \sum_{i=2^k+1}^n X_i \geq \alpha n / 2 \right\} \right) \\
&\leq \sum_{k=\log_2(n_0)}^{\infty} 2 \exp(-\alpha^2 2^k / 2) \leq 2 \int_{\log_2(n_0)}^{\infty} \exp(-(\alpha/2)^2 2^x) dx \\
&= \frac{2}{\log(2)} \int_{n_0}^{\infty} u^{-1} \exp(-(\alpha/2)^2 u) du \leq \frac{8 \exp(-(\alpha/2)^2 n_0)}{n_0 \alpha^2 \log(2)}.
\end{aligned}$$

Putting the pieces together we have

$$\begin{aligned}
\mathbb{P} \left(\bigcup_{n=1}^{\infty} \left\{ \sum_{i=1}^n X_i \geq \alpha n + \beta \right\} \right) &\leq \min_{n_0} \exp(-2\beta^2/n_0) + \frac{8 \exp(-(\alpha/2)^2 n_0)}{n_0 \alpha^2 \log(2)} \\
&\leq \exp(-\beta\alpha) + \frac{4 \exp(-\beta\alpha/2)}{\beta\alpha \log(2)} \leq 7 \exp(-\beta\alpha/2)
\end{aligned}$$

where the last inequality holds with $\beta\alpha \geq 1$. ■

Lemma 2 Given $\theta_1 - \hat{\gamma} \geq \frac{2B}{m}$,

$$\mathbb{P} \left(\max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B \mid \mu_i = \theta_1 \right) \geq 1 - \exp(-m(\theta_1 - \hat{\gamma})^2/2).$$

Similarly, given $\hat{\gamma} - \theta_0 \geq \frac{2|A|}{m}$,

$$\mathbb{P} \left(\min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) < A \mid \mu_i = \theta_0 \right) \geq 1 - \exp(-m(\hat{\gamma} - \theta_0)^2/2).$$

Proof We analyze the left hand side of the lemma:

$$\begin{aligned}
&\mathbb{P} \left(\max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B \mid \mu_i = \theta_1 \right) \\
&= \mathbb{P} \left(\bigcup_{j=1}^m \left\{ \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B \mid \mu_i = \theta_1 \right\} \right) \\
&\geq \mathbb{P} \left(\sum_{s=1}^m (X_{i,s} - \hat{\gamma}) > B \mid \mu_i = \theta_1 \right) \\
&= 1 - \mathbb{P} \left(\frac{1}{m} \sum_{s=1}^m (X_{i,s} - \mu_i) \leq \frac{B}{m} - (\mu_i - \hat{\gamma}) \mid \mu_i = \theta_1 \right) \\
&= 1 - \mathbb{P} \left(\frac{1}{m} \sum_{s=1}^m (\mu_i - X_{i,s}) \geq (\mu_i - \hat{\gamma}) - \frac{B}{m} \mid \mu_i = \theta_1 \right) \\
&\geq 1 - \exp \left(-2m \left[(\theta_1 - \hat{\gamma}) - \frac{B}{m} \right]^2 \right) \\
&\geq 1 - \exp \left(\frac{-m(\theta_1 - \hat{\gamma})^2}{2} \right)
\end{aligned}$$

Where the second to last statement holds by Hoeffding's inequality, and the last uses the bound on B/m given in the lemma. A nearly identical argument yields the second half of the lemma. ■

Lemma 3 If $\theta_1 - \hat{\gamma} \geq \frac{2B}{m}$ then

$$\begin{aligned} & \mathbb{P} \left(\bigcup_{i=1}^n \left\{ \mu_i = \theta_1, \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \right\} \right) \\ & \geq 1 - \exp[-\alpha n(1 - \exp(-B(\theta_1 - \hat{\gamma})) - 7 \exp(-|A|(\hat{\gamma} - \theta_0)/2))] \end{aligned}$$

Proof Consider iid events Ω_i for $i = 1, \dots, n$. Then $\mathbb{P}(\bigcup_{i=1}^n \Omega_i) = 1 - \mathbb{P}(\bigcap_{i=1}^n \Omega_i^c) = 1 - \mathbb{P}(\Omega_1^c)^n = 1 - (1 - \mathbb{P}(\Omega_1))^n \geq 1 - \exp(-n\mathbb{P}(\Omega_1))$. We follow the same line of reasoning:

$$\begin{aligned} & \mathbb{P} \left(\bigcup_{i=1}^n \left\{ \mu_i = \theta_1, \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \right\} \right) \\ & = 1 - \left(1 - \mathbb{P} \left(\mu_i = \theta_1, \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \right) \right)^n \\ & = 1 - \left(1 - \alpha \mathbb{P} \left(\max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \mid \mu_i = \theta_1 \right) \right)^n \\ & = 1 - \left(1 - \alpha \left(1 - \mathbb{P} \left(\max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) < B \mid \mu_i = \theta_1 \right) - \mathbb{P} \left(\min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) < A \mid \mu_i = \theta_1 \right) \right) \right)^n \\ & \geq 1 - (1 - \alpha(1 - \exp(-m(\theta_1 - \hat{\gamma})^2/2) - 7 \exp(-|A|(\hat{\gamma} - \theta_0)/2)))^n \\ & \geq 1 - \exp[-\alpha n(1 - \exp(-m(\theta_1 - \hat{\gamma})^2/2) - 7 \exp(-|A|(\hat{\gamma} - \theta_0)/2))] \\ & \geq 1 - \exp[-\alpha n(1 - \exp(-B(\theta_1 - \hat{\gamma})) - 7 \exp(-|A|(\hat{\gamma} - \theta_0)/2))] \end{aligned}$$

Where the third-to-last inequality applies Lemmas 1 and 2. ■

Now, we are ready to prove Theorem 4.

Proof First, we consider the estimation of $\hat{\theta}_0$ of Algorithm 2, then consider the sample complexity of the algorithm, and then prove correctness.

Let $\xi_0 = \{\hat{\theta}_0 - \theta_0 \geq -\frac{\epsilon_0}{4}\}$ and $\xi_1 = \{\hat{\theta}_0 - \theta_0 \leq \frac{\epsilon_0}{4}\}$ be the events that we accurately estimate the parameter θ_0 . We will show that $\mathbb{P}(\xi_0) \geq 1 - \delta'$ and $\mathbb{P}(\xi_1) \geq 3/4$ where $\delta' = \min\{\delta/8, \frac{1}{m\epsilon_0^2}\}$. Let $k_1 = 5$ and $k_2 = 8\epsilon_0^{-2} \log(\frac{2k_1}{\delta'})$. First note that

$$\mathbb{P} \left(\bigcup_{i=1}^{k_1} \left\{ |\hat{\mu}_{i,k_2} - \mu_i| \geq \frac{\epsilon_0}{4} \right\} \right) \leq 2k_1 \exp(-2k_2(\epsilon_0/4)^2) \leq \delta'$$

so that with probability at least $1 - \delta'$ we have $\hat{\theta}_0 = \min_{i=1, \dots, k_1} \hat{\mu}_{i,k_2} \geq \min_{i=1, \dots, k_1} \mu_i - \epsilon_0/4 \geq \theta_0 - \epsilon_0/4$, and in particular, $\mathbb{P}(\xi_0) \geq 1 - \delta'$. Let $\mathcal{E} = \{\bigcup_{i=1}^{k_1} \{\mu_i = \theta_0\}\}$ be the event that at least one of the distributions is light. Then

$$\mathbb{P}(\mathcal{E}) = 1 - \alpha^{k_1} \geq 1 - 2^{-k_1} \geq 31/32,$$

so that under $\mathcal{E} \cap \xi_0$, we have $\hat{\theta}_0 = \min_{i=1, \dots, k_1} \hat{\mu}_{i,k_2} \leq \min_{i=1, \dots, k_1} \mu_i + \epsilon_0/4 = \theta_0 + \epsilon_0/4$ which means $\mathbb{P}(\xi_1^c) \leq \mathbb{P}(\xi_0^c \cup \mathcal{E}^c) \leq \delta/8 + 1/32 \leq 1/16$. Moreover, the total number of samples is bounded by $k_1 k_2 = c\epsilon_0^{-2} \log(1/\delta') \leq c\epsilon_0^{-2} \log(\max\{\frac{1}{\delta}, \log(\frac{1}{\alpha_0 \delta})\})$ which is clearly dominated by $\frac{\log(1/\delta)}{\alpha_0 \epsilon_0^2}$.

We now turn our attention to the sample complexity. By Wald's identity [17, Proposition 2.18],

$$\mathbb{E}[T] = \mathbb{E} \left[\sum_{i=1}^N M_i \right] = \mathbb{E}[N] \mathbb{E}[M_1] = \mathbb{E}[N] ((1 - \alpha) \mathbb{E}[M_1 | \mu_1 = \theta_0] + \alpha \mathbb{E}[M_1 | \mu_1 = \theta_1]).$$

Trivially, $\mathbb{E}[N] \leq n$ and $\mathbb{E}[M_1 | \mu_1 = \theta_1] \leq m$, so we only need to bound $\mathbb{E}[M_1 | \mu_1 = \theta_0]$. Clearly we have that

$$\mathbb{E}[M_1 | \mu_1 = \theta_0] = \mathbb{E}[M_1 | \xi_0, \mu_1 = \theta_0] \mathbb{P}(\xi_0) + \mathbb{E}[M_1 | \xi_0^c, \mu_1 = \theta_0] \mathbb{P}(\xi_0^c) \leq \mathbb{E}[M_1 | \xi_0, \mu_1 = \theta_0] + \delta' m$$

so

$$\begin{aligned}
\mathbb{E}[M_1 | \xi_0, \mu_1 = \theta_0] &\leq \sum_{t=1}^{\infty} \mathbb{P} \left(\arg \min_j \left\{ \sum_{s=1}^j (X_{1,s} - \hat{\gamma}) < A \mid \xi_0, \mu_1 = \theta_0 \right\} \geq t \right) \\
&= \sum_{t=1}^{\infty} 1 - \mathbb{P} \left(\min_{j=1, \dots, t-1} \sum_{s=1}^j (X_{1,s} - \hat{\gamma}) < A \mid \xi_0, \mu_1 = \theta_0 \right) \\
&= \sum_{t=0}^{\infty} 1 - \mathbb{P} \left(\min_{j=1, \dots, t} \sum_{s=1}^j (X_{1,s} - \hat{\gamma}) < A \mid \xi_0, \mu_1 = \theta_0 \right) \\
&\leq \sum_{t=0}^{\infty} 1 - \mathbf{1}_{\hat{\gamma} - \theta_0 \geq \frac{2|A|}{t}} (1 - \exp(-t(\hat{\gamma} - \theta_0)^2/2)) \\
&\leq \frac{2|A|}{\hat{\gamma} - \theta_0} + 2e^{1/2}(\hat{\gamma} - \theta_0)^{-2} \exp(-|A|(\hat{\gamma} - \theta_0)) \leq \frac{3|A|}{\hat{\gamma} - \theta_0} \leq \frac{293}{\epsilon_0^2}.
\end{aligned}$$

where the second inequality follows by applying Lemma 2 and the last inequality holds by ξ_0 and the value of $|A|$ since if ξ_0 holds, $\hat{\gamma} - \theta_0 = \hat{\theta}_0 - \theta_0 + \frac{\epsilon_0}{2} \geq \frac{\epsilon_0}{4}$. Thus

$$\mathbb{E}[M_1] \leq (1 - \alpha) \left[\left(\frac{293}{\epsilon_0^2} \right) + \delta' m \right] + \alpha m \leq \delta' m + \frac{1}{\epsilon_0^2} (293 + 64\alpha \log(\frac{14n}{\delta})) \leq \frac{c\alpha \log(\frac{1}{\alpha_0 \delta})}{\epsilon_0^2}$$

for some c where we use the fact that $\delta' m \leq \epsilon_0^{-2}$. So we have

$$\mathbb{E}[T] \leq n\mathbb{E}[M_1] \leq \frac{c' \alpha \log(1/\alpha_0) + c'' \log(\frac{1}{\delta})}{\alpha_0 \epsilon_0^2}.$$

Now, we analyze the correctness claims. Under ξ_0 , $\hat{\gamma} - \theta_0 \geq \frac{\epsilon_0}{4}$. Note that this event fails to occur with probability less than $\delta/2$, and if it is used in conjunction with some other event that fails to occur with probability $\delta/2$, we may conclude that either of these events fail with probability less than δ .

To justify Claim 1, we apply Lemma 1 to observe that the probability that we output a light distribution is no greater than

$$\begin{aligned}
&\mathbb{P}(\xi_0^c) + \mathbb{P} \left(\bigcup_{i=1}^n \left\{ \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \mu_i = \theta_0 \right\} \mid \xi_0 \right) \mathbb{P}(\xi_0) \\
&\leq \mathbb{P}(\xi_0^c) + n(1 - \alpha) \mathbb{P} \left(\max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B \mid \mu_i = \theta_0, \xi_0 \right) \\
&\leq \delta/2 + 7n \exp(-B(\hat{\gamma} - \theta_0)/2) \leq \delta
\end{aligned}$$

where we have used $\hat{\gamma} - \theta_0 \geq \frac{\epsilon_0}{4}$ and plugged in the values of B and n .

To justify Claim 2, assume $\alpha_0 \leq \alpha$ and $\epsilon_0 \leq \theta_1 - \theta_0$. We apply Lemma 3 to observe that the probability that we return a heavy distribution is at least

$$\begin{aligned}
&\mathbb{P} \left(\xi_0 \cap \xi_1 \cap \bigcup_{i=1}^n \left\{ \mu_i = \theta_1, \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \right\} \right) \\
&= \mathbb{P}(\xi_0 \cap \xi_1) \mathbb{P} \left(\bigcup_{i=1}^n \left\{ \mu_i = \theta_1, \max_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > B, \min_{j=1, \dots, m} \sum_{s=1}^j (X_{i,s} - \hat{\gamma}) > A \right\} \mid \xi_0, \xi_1 \right) \\
&\geq \mathbb{P}(\xi_0 \cap \mathcal{E}) (1 - \exp[-\alpha n (1 - \exp(-B(\epsilon_0/4)) - 7 \exp(-|A|(\epsilon_0/4)/2))]) \\
&\geq (15/16) (1 - \exp[-\alpha n (1 - (\frac{\delta}{14n})^2 - 1/3)]) \geq (15/16)(8/9) \geq 4/5
\end{aligned}$$

where we have used $\mathbb{P}(\xi_0 \cap \mathcal{E}) \geq 1 - \mathbb{P}(\xi_0^c) - \mathbb{P}(\mathcal{E}^c) \geq 15/16$, $(\frac{\delta}{14n})^2 \leq 1/6$, $\alpha n \geq 2 \log(9)$ and plugged in the values for A and B .

To justify Claim 3, we simply observe that the algorithm always terminates after $n \times m$ steps. ■

B.3 Either α or ϵ is unknown, but not both

<p>Algorithm 4 Algorithm for unknown $\theta_1 - \theta_0$.</p> <p>Given $\delta \in (0, 1), \alpha \in (0, 1/2)$.</p> <p>Initialize $k = 1$</p> <p>While Algorithm 2 run with inputs $\delta/(2k^2)$, $\alpha_0 = \alpha, \epsilon_0 = 2^{-k}$ returns null: Set $k = k + 1$.</p> <p>Output distribution k.</p>	<p>Algorithm 5 Algorithm for unknown α.</p> <p>Given $\delta \in (0, 1), \epsilon \in (0, 1]$.</p> <p>Initialize $k = 1$</p> <p>While Algorithm 2 run with inputs $\delta/(2k^2)$, $\alpha_0 = 2^{-k}, \epsilon_0 = \epsilon$ returns null: Set $k = k + 1$.</p> <p>Output distribution k.</p>
---	--

First we consider the case when α is known but a lower bound on $\theta_1 - \theta_0$ is not, and then opposite.

Theorem 8 (Known α , unknown θ_0, θ_1) Fix $\delta \in (0, 1)$. If Algorithm 4 is run with δ, α then with probability at least $1 - \delta$ a heavy distribution is returned and the expected number of total samples taken is no more than

$$\frac{c \log \left(\log \left(\frac{1}{(\theta_1 - \theta_0)^2} \right) / \delta \right)}{\alpha (\theta_1 - \theta_0)^2}.$$

for an absolute constant c .

Proof On each stage k , Algorithm 2 is called with $\delta/(2k^2)$. By the guarantees of Theorem 4, the probability that Algorithm 4 ever outputs a light distribution is less than $\sum_{k=1}^{\infty} \delta/(2k^2) \leq \delta$. Thus, if a distribution is output, it is heavy with probability at least $1 - \delta$. We now show that the expected number of samples taken before outputting a distribution is bounded.

Let K be the random stage in which Algorithm 4 outputs a distribution and let k_* be the smallest $k \in \mathbb{N}$ that satisfies $2^{-k} \leq \theta_1 - \theta_0$. By the guarantees of Theorem 4 and the independence of the stages k , $\mathbb{P}(K \geq k_* + i) \leq \sum_{\ell=i}^{\infty} (\frac{1}{5})^\ell = (\frac{5}{4})(\frac{1}{5})^i$. Moreover, if M_k is the number of measurements taken at stage k , then by Wald's identity the expected number of measurements is bounded by

$$\begin{aligned} \mathbb{E} \left[\sum_{k=1}^K M_k \right] &= \sum_{k=1}^{\infty} \mathbb{E}[N_k] \mathbb{P}(K \geq k) \leq \sum_{k=1}^{\infty} \frac{c' \alpha \log(1/\alpha) + c'' \log \left(\frac{2k^2}{\delta} \right)}{\alpha 2^{-2k}} \max \{ 1, (\frac{5}{4})(\frac{1}{5})^{k-k_*} \} \\ &\leq \sum_{k=1}^{k_*} \frac{c''' \log \left(\frac{2k_*^2}{\delta} \right)}{\alpha} 4^k + 5^{k_*} \frac{c'''}{\alpha} \sum_{k=k_*+1}^{\infty} (2 \log(k) + \log(\frac{2}{\delta})) \left(\frac{4}{5} \right)^k \\ &\leq \frac{c''' \log \left(\frac{2k_*^2}{\delta} \right)}{\alpha} 4^{k_*+1} + 5^{k_*} \frac{c'''}{\alpha} \sum_{k=k_*+1}^{\infty} (2 \log(k) + \log(\frac{2}{\delta})) \left(\frac{4}{5} \right)^k \leq \frac{c'''' \log \left(\frac{k_*}{\delta} \right)}{\alpha} (2^{k_*})^2 \end{aligned}$$

since $\sup_{\alpha} \alpha \log(1/\alpha) \leq e^{-1}$ and

$$\begin{aligned} \sum_{k=k_*}^{\infty} \log(k) \left(\frac{4}{5} \right)^k &= \sum_{k=k_*}^{2k_*-1} \log(k) \left(\frac{4}{5} \right)^k + \sum_{k=2k_*}^{\infty} \log(k) \left(\frac{4}{5} \right)^{k/2} \left(\frac{4}{5} \right)^{k/2} \\ &\leq \log(2k_*) \sum_{k=k_*}^{2k_*-1} \left(\frac{4}{5} \right)^k + \sum_{k=2k_*}^{\infty} \left(\frac{4}{5} \right)^{k/2} \leq (\log(2k_*) + 2) \sum_{k=k_*}^{\infty} \left(\frac{4}{5} \right)^k = 5 \log(2e^2 k_*) \left(\frac{4}{5} \right)^{k_*} \end{aligned}$$

since $\sup_k \log(k) \left(\frac{4}{5} \right)^{k/2} \leq 1$. Noting that $k_* \leq \log_2 \left(\frac{1}{\theta_1 - \theta_0} \right) + 1$ completes the proof. \blacksquare

Theorem 9 (Unknown α , known θ_0, θ_1) Fix $\delta \in (0, 1)$. If Algorithm 5 is run with $\delta, \theta_1 - \theta_0$ then with probability at least $1 - \delta$ a heavy distribution is returned and the the expected number of total samples taken is no more than

$$\frac{c \log \left(\log \left(\frac{1}{\alpha} \right) / \delta \right)}{\alpha (\theta_1 - \theta_0)^2}$$

for an absolute constant c .

Proof The proof of this result is nearly identical to that of Theorem 8 except the following changes. Let K be the random stage in which Algorithm 5 outputs a distribution and let k_* be the smallest $k \in \mathbb{N}$ that satisfies $2^{-k} \leq \alpha$. Moreover, if M_k is the number of measurements taken at stage k , then by Wald's identity expected number of measurements is bounded by

$$\begin{aligned} \mathbb{E} \left[\sum_{k=1}^K M_k \right] &= \sum_{k=1}^{\infty} \mathbb{E}[N_k] \mathbb{P}(K \geq k) \leq \sum_{k=1}^{\infty} \frac{c' \alpha \log(2^k) + c'' \log\left(\frac{2k^2}{\delta}\right)}{2^{-k} \epsilon^2} \max\{1, (\frac{5}{4})(\frac{1}{5})^{k-k_*}\} \\ &\leq \sum_{k=1}^{k_*} \frac{c''' \log\left(\frac{2k^2}{\delta}\right)}{\epsilon^2} 2^k + 5^{k_*} \frac{c'''}{\epsilon^2} \sum_{k=k_*+1}^{\infty} (\alpha k \log(2) + 2 \log(k) + \log(\frac{2}{\delta})) \left(\frac{2}{5}\right)^k \\ &\leq \frac{c'''' (\alpha k_* + \log\left(\frac{k_*}{\delta}\right))}{\epsilon^2} 2^{k_*} \leq \frac{c'''' \log(\log(1/\alpha)/\delta)}{\alpha \epsilon^2} \end{aligned}$$

by the same series of steps as the proof of Theorem 8 and the fact that $\sum_{k=n}^{\infty} k a^k \leq \frac{n a^n}{(1-a)^2}$ for any $a \in (0, 1)$. The final inequality follows from $k_* \leq \log_2(1/\alpha) + 1$ and that $\alpha k_* = \alpha \log_2(2/\alpha) \leq 2$. ■

B.4 Proof of Theorem 5

Proof The proof is broken up into a few steps, summarized as follows. For any given α_0, ϵ_0 , Theorem 4 takes just $O\left(\frac{\alpha \log(1/\alpha_0) + \log(1/\delta)}{\alpha_0 \epsilon_0^2}\right)$ samples in expectation and the procedure makes an error (i.e. returns a light distribution) with probability less than δ . Define $\epsilon = \theta_1 - \theta_0$. In addition, if $\epsilon = \theta_1 - \theta_0$, $\alpha \geq \alpha_0$, and $\epsilon \geq \epsilon_0$ then with probability at least $4/5$ a heavy distribution is returned after the same expected number of samples. We will leverage this result to show that if we are given an upper bound γ_0 such that $\frac{1}{\alpha \epsilon^2} \leq \gamma_0$ then it is possible to identify a heavy distribution with probability at least $4/5$ using just $O(\log_2(\gamma_0) \gamma_0 [\alpha \log_2(\gamma_0) + \log(\log_2(\gamma_0)/\delta)])$ samples in expectation. Finally, we apply the ‘‘doubling trick’’ to γ so that even though the tightest γ is not known a priori, we can adapt to it using only twice the number of samples as if we had known it. Because each of the stages is independent of one another, the probability that the procedure requires more than $\ell_* + i$ stages is less than $(1/5)^i$, which yields our expected sample complexity.

For all $\ell \in \mathbb{N}$ define $\delta_\ell = \frac{\delta}{2^\ell}$ and $\gamma_\ell = 2^\ell$. Fix some ℓ and consider the set $\{(\alpha, \epsilon) : \frac{1}{\alpha \epsilon^2} = \gamma_\ell\}$. Clearly, in this set, $\alpha \in [1/\gamma_\ell, 1/2]$. For all $k \in \{0, \dots, \ell - 1\}$, define $\alpha_k = \frac{2^k}{\gamma_\ell}$ and $\epsilon_k = \sqrt{\frac{1}{2\alpha_k \gamma_\ell}}$. The key observation is that

$$\{(\alpha, \epsilon) : \frac{1}{\alpha \epsilon^2} \leq \gamma_\ell\} \subseteq \bigcup_{k=0}^{\log_2 \gamma_\ell - 1} \{(\alpha, \epsilon) : \alpha \geq \alpha_k, \epsilon \geq \epsilon_k\}. \quad (8)$$

To see this, fix any (α', ϵ') such that $\frac{1}{\alpha' \epsilon'^2} \leq \gamma_\ell$. Let k_* be the integer that satisfies $\alpha_{k_*} \leq \alpha' < 2\alpha_{k_*}$. Such a k_* must exist since $\alpha_{\ell-1} = \frac{1}{2} \geq \alpha' \geq \frac{1}{\gamma_\ell \epsilon'^2} \geq \frac{1}{\gamma_\ell} = \alpha_0$. Then $\gamma_\ell \geq \frac{1}{\alpha' \epsilon'^2} \geq \frac{1}{2\alpha_{k_*} \epsilon'^2}$ which means $\epsilon' \geq \sqrt{\frac{1}{2\alpha_{k_*} \gamma_\ell}} = \epsilon_{k_*}$ which proves the claim of (8). Consequently, even if no information about α or ϵ individually is known but $\frac{1}{\alpha \epsilon^2} \leq \gamma_\ell$, one can cover the entire range of valid (α, ϵ) with just $\log_2(\gamma_\ell) = \ell$ landmarks (α_k, ϵ_k) .

For any $\ell \in \mathbb{N}$ and $k \in \{0, \dots, \ell - 1\}$, if Algorithm 2 is used with $\alpha_0 = \alpha_k, \epsilon_0 = \epsilon_k$ and $\delta = \delta_\ell$ then the probability that a light distribution is returned, declared heavy is less than δ_ℓ . And the probability that a light distribution is returned, declared heavy for any $\ell \in \mathbb{N}$ and $k \in \{0, \dots, \ell - 1\}$ is less than $\sum_{\ell=1}^{\infty} \ell \delta_\ell = \delta \sum_{\ell=1}^{\infty} \ell / (2^{\ell^3}) \leq \delta$. Thus, given that Algorithm 3 terminates with a non-null distribution h , h is heavy with probability at least $1 - \delta$. This proves correctness. We next bound the expected number of samples taken before the procedure terminates.

With the inputs given in the last paragraph for any k, ℓ , Algorithm 2 takes an expected number samples bounded by $c \gamma_\ell (\alpha \log(1/\alpha_k) + \log(1/\delta_\ell))$. Let $L \in \mathbb{N}$ be the random stage at which Algorithm 3 terminates with a non-null distribution h . Let ℓ_* be the first integer such that there exists a $k \in \{0, \dots, \ell_* - 1\}$ with $\alpha \geq \alpha_k$ and $\epsilon \geq \epsilon_k$ (recall that in this case $\frac{1}{\alpha_k \epsilon_k^2} \leq \gamma_{\ell_*}$). Then by the end of stage $\ell \geq \ell_*$, at most $c \ell \gamma_\ell (\alpha \log(\gamma_\ell) + \log(1/\delta_\ell))$ samples in expectation were taken on stage ℓ and with probability at least $4/5$ the procedure terminated with a heavy coin. By the independence of samples between rounds, observe that $\mathbb{P}(L \geq \ell_* + i) = \sum_{j=i}^{\infty} \mathbb{P}(L = \ell_* + j) \leq (\frac{5}{4})(\frac{1}{5})^i$. Thus, if M_ℓ is the number of samples taken at stage ℓ then

by Wald's identify, the total expected number of samples taken before termination is bounded by

$$\begin{aligned}
& \mathbb{E} \left[\sum_{\ell=1}^L c\ell\gamma_\ell(\alpha \log(\gamma_\ell) + \log(1/\delta_\ell)) \right] = \sum_{\ell=1}^{\infty} \mathbb{E}[M_\ell] \mathbb{P}(L \geq \ell) \leq \sum_{\ell=1}^{\infty} c\ell\gamma_\ell(\alpha \log(\gamma_\ell) + \log(1/\delta_\ell)) \mathbb{P}(L \geq \ell) \\
& \leq \sum_{\ell=1}^{\ell_*} c\ell\gamma_\ell(\alpha \log(\gamma_\ell) + \log(1/\delta_\ell)) + \sum_{\ell=\ell_*+1}^{\infty} c\ell\gamma_\ell(\alpha \log(\gamma_\ell) + \log(1/\delta_\ell)) \left(\frac{5}{4}\right) \left(\frac{1}{5}\right)^{\ell-\ell_*} \\
& \leq \sum_{\ell=1}^{\ell_*} c\ell 2^\ell (\alpha\ell + \log(2\ell^3/\delta)) + \sum_{\ell=\ell_*+1}^{\infty} c\ell 2^\ell (\alpha\ell + \log(2\ell^3/\delta)) \left(\frac{5}{4}\right) \left(\frac{1}{5}\right)^{\ell-\ell_*} \\
& \leq c\ell_* (\alpha\ell_* + \log(2\ell_*^3/\delta)) \sum_{\ell=1}^{\ell_*} 2^\ell + c\left(\frac{5}{4}\right) 5^{\ell_*} \sum_{\ell=\ell_*+1}^{\infty} \left(\alpha\ell^2 \left(\frac{2}{5}\right)^\ell + 3\ell \log(\ell) \left(\frac{2}{5}\right)^\ell + \log(2/\delta) \ell \left(\frac{2}{5}\right)^\ell \right) \\
& \leq 2c\ell_* 2^{\ell_*} (\alpha\ell_* + \log(2\ell_*^3/\delta)) \\
& \quad + c\left(\frac{5}{4}\right) 5^{\ell_*} \left(2\alpha(\ell_* + 1)^2 \left(\frac{2}{5}\right)^{\ell_*} + 12 \log(2e^2\ell_*) (\ell_* + 1) \left(\frac{2}{5}\right)^{\ell_*} + 4 \log(2/\delta) (\ell_* + 1) \left(\frac{2}{5}\right)^{\ell_*} \right) \\
& \leq c' \ell_* 2^{\ell_*} (\alpha\ell_* + \log(\ell_*) + \log(1/\delta))
\end{aligned}$$

for some absolute constant c' since $\sum_{k=n}^{\infty} k a^k \leq \frac{na^n}{(1-a)^2}$, $\sum_{k=n}^{\infty} k^2 a^k \leq \frac{n^2 a^n}{(1-a)^3}$, and

$$\begin{aligned}
\sum_{\ell=\ell_*+1}^{\infty} \ell \log(\ell) \left(\frac{2}{5}\right)^\ell & \leq \log(2\ell_*) \sum_{\ell_*+1}^{2\ell_*} \ell \left(\frac{2}{5}\right)^\ell + \sum_{2\ell_*+1}^{\infty} \ell \left(\frac{2}{5}\right)^{\ell/2} \left(\log(\ell) \left(\frac{2}{5}\right)^{\ell/2} \right) \\
& \leq \log(2e^2\ell_*) \sum_{\ell_*+1}^{\infty} \ell \left(\frac{2}{5}\right)^\ell \leq 4 \log(2e^2\ell_*) (\ell_* + 1) \left(\frac{2}{5}\right)^{\ell_*}
\end{aligned}$$

since $\max_{x \geq 1} \log(x) \left(\frac{2}{5}\right)^{x/2} \leq 1$. Noting that $\ell_* \leq \log_2\left(\frac{1}{\alpha\epsilon^2}\right) + 1$, we have that the total number of samples, in expectation, is bounded by

$$\begin{aligned}
c' \ell_* 2^{\ell_*} (\alpha\ell_* + \log(\ell_*) + \log(1/\delta)) & \leq c'' \frac{\log_2\left(\frac{1}{\alpha\epsilon^2}\right)}{\alpha\epsilon^2} \left(\alpha \log_2\left(\frac{1}{\alpha\epsilon^2}\right) + \log\left(\log_2\left(\frac{1}{\alpha\epsilon^2}\right)\right) + \log(1/\delta) \right) \\
& \leq c''' \frac{\log_2\left(\frac{1}{\alpha\epsilon^2}\right)}{\alpha\epsilon^2} \left(\alpha \log_2\left(\frac{1}{\epsilon^2}\right) + \log\left(\log_2\left(\frac{1}{\alpha\epsilon^2}\right)\right) + \log(1/\delta) \right)
\end{aligned}$$

where we've used the fact that $\sup_{\alpha \in [0,1]} \alpha \log(1/\alpha) \leq e^{-1}$. ■

C Gaussians

C.1 On the detection of a mixture of Gaussians

For known σ^2 , consider the hypothesis test of Problem P1. In what follows, let $\chi^2(\theta_1, \theta_0)$ and $KL(\theta_1, \theta_0)$ be the chi-squared and KL divergences of the two distributions of \mathbf{H}_1 . Note that for $\frac{(\theta_1 - \theta_0)^2}{\sigma^2} \leq 1$, we have that $\chi^2(\theta_1, \theta_0) = e^{\frac{(\theta_1 - \theta_0)^2}{\sigma^2}} - 1 \leq 2 \frac{(\theta_1 - \theta_0)^2}{\sigma^2} = 4KL(\theta_1, \theta_0)$

Theorem 2 says that for $\frac{(\theta_1 - \theta_0)^2}{\sigma^2} \leq 1$, a procedure that has maximum probability of error less than δ requires at least $\max \left\{ \frac{1-\delta}{\alpha}, \frac{\log(1/\delta)}{4\alpha^2 KL(\theta_1, \theta_0)} \right\}$ samples to decide the above hypothesis test, even if $\alpha, \theta_0, \theta_1$ are known. The next subsection shows that if $\alpha, \theta_0, \theta_1$ are unknown then one requires at least $\frac{\log(1/\delta)}{2[\alpha KL(\theta_1, \theta_0)]^2}$ samples to decide the above hypothesis test correctly with probability at least $1 - \delta$. This is likely achievable using the method of moments [12].

C.2 Lower bounds

Theorem 10 *For known σ^2 , consider the hypothesis test of Problem P1. If $\theta_* = (1 - \alpha)\theta_0 + \alpha\theta_1$ and $\frac{\theta_1 - \theta_0}{\sigma} \leq 1$ then*

$$\chi^2((1 - \alpha)f_{\theta_0}(x) + \alpha f_{\theta_1}(x) | f_{\theta_*}(x)) \leq c' \left(\alpha(1 - \alpha) \frac{(\theta_1 - \theta_0)^2}{\sigma^2} \right)^2$$

for some absolute constant c' .

Proof If $f_\theta = \mathcal{N}(\theta, \sigma^2)$ then $f_\theta(x) = h(x) \exp(\eta(\theta)x - b(\theta))$ where $h(x) = \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{x^2}{2\sigma^2}}$, $\eta(\theta) = \frac{\theta}{\sigma^2}$, and $b(\eta(\theta)) = \frac{\eta(\theta)^2 \sigma^2}{2} = \frac{\theta^2}{2\sigma^2}$. Thus,

$$\theta_* = \eta^{-1}((1 - \alpha)\eta(\theta_0) + \alpha\eta(\theta_1)) = (1 - \alpha)\theta_0 + \alpha\theta_1$$

and

$$\sup_{y \in [\theta_0, \theta_1]} b(2\eta(y) - \eta(\theta_*)) - (2b(\eta(y)) - b(\eta(\theta_*))) = \sup_{y \in [\theta_0, \theta_1]} \frac{(y - \theta_*)^2}{\sigma^2} \leq \frac{(\theta_1 - \theta_0)^2}{\sigma^2} =: \kappa$$

and

$$\sup_{x \in [\dot{b}(\eta(\theta_-)), \dot{b}(\eta(\theta_+))]} f_{\dot{b}^{-1}(x)}(x) = \sup_{x \in [\theta_-, \theta_+]} \sup_{\theta \in \mathbb{R}} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\theta)^2}{2\sigma^2}} \leq \frac{1}{\sqrt{2\pi\sigma^2}} =: \gamma.$$

Note that for any $\theta < \theta'$ we have $\dot{b}(\eta(\theta')) - \dot{b}(\eta(\theta)) = \theta' - \theta$, $M_2(\theta) = \sigma^2$, and $M_4(\theta) = 3\sigma^4$. Plugging these values into the theorem we have

$$\begin{aligned} c &= e^\kappa \left(\sup_{\theta \in [\theta_0, \theta_1]} M_2(\theta)^2 \left(2 + \gamma \left(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)) \right) \right) \right. \\ &\quad \left. + 8M_4(\theta_-) + 8M_4(\theta_+) + 16 \left(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)) \right)^4 + \frac{2}{5} \gamma \left(\dot{b}(\eta(\theta_+)) - \dot{b}(\eta(\theta_-)) \right)^5 \right) \\ &= e^{\frac{(\theta_1 - \theta_0)^2}{\sigma^2}} \left(\sigma^4 \left(2 + \frac{2(\theta_1 - \theta_0)}{\sqrt{2\pi}\sigma} \right) + 48\sigma^4 + 256(\theta_1 - \theta_0)^4 + \frac{64}{5\sqrt{2\pi}} \frac{(\theta_1 - \theta_0)^5}{\sigma} \right) \end{aligned}$$

noting that $\theta_+ - \theta_- = 2(\theta_1 - \theta_0)$. If $\frac{\theta_1 - \theta_0}{\sigma} \leq 1$ then $c = c'\sigma^4$ for some absolute constant c' and $(\eta(\theta_1) - \eta(\theta_0))^2 = \frac{(\theta_1 - \theta_0)^2}{\sigma^4}$ which yields the final result. \blacksquare

C.3 Gaussian Upper bound for known $\alpha, \theta_0, \theta_1$

For known σ^2 , consider the hypothesis test of Problem P1 with $\theta = \theta_0$. We observe a sample X_1, \dots, X_n and are trying to establish whether it came from \mathbf{H}_0 or \mathbf{H}_1 .

Consider the test

$$\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i > \theta_1} \underset{\mathbf{H}_0}{\overset{\mathbf{H}_1}{\geq}} \frac{\mathbb{P}_1(X_1 > \theta_1) + \mathbb{P}_0(X_1 > \theta_1)}{2} =: \gamma.$$

If $\epsilon = \mathbb{P}_1(X_1 > \theta_1) - \mathbb{P}_0(X_1 > \theta_1)$ then

$$\mathbb{P}_1 \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i > \theta_1} \leq \gamma \right) = \mathbb{P}_1 \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i > \theta_1} \leq \mathbb{P}_1(X_1 > \theta_1) - \epsilon/2 \right) \leq e^{-n\epsilon^2/2}$$

and

$$\mathbb{P}_0 \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i > \theta_1} \geq \gamma \right) = \mathbb{P}_0 \left(\frac{1}{n} \sum_{i=1}^n \mathbf{1}_{X_i > \theta_1} \geq \mathbb{P}_0(X_1 > \theta_1) + \epsilon/2 \right) \leq e^{-n\epsilon^2/2}$$

by sub-Gaussian tail bounds. If $Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz$ and $\Delta = \frac{\theta_1 - \theta_0}{\sigma}$ then

$$\begin{aligned} \mathbb{P}_0(X_1 > \theta_1) &= Q(\Delta) \\ \mathbb{P}_1(X_1 > \theta_1) &= (1 - \alpha)Q(\Delta) + \alpha \frac{1}{2} \end{aligned}$$

so

$$\epsilon = \alpha \left(\frac{1}{2} - Q(\Delta) \right) = \alpha \int_0^\Delta \frac{1}{\sqrt{2\pi}} e^{-x^2/2} dx \geq \min \left\{ \frac{\alpha\Delta}{4\sqrt{2\pi}}, \frac{1}{4}\alpha \right\}.$$

Thus, the test fails with probability at most

$$\exp \left[-n\alpha^2 \min \left\{ \frac{(\theta_1 - \theta_0)^2}{64\pi\sigma^2}, \frac{1}{32} \right\} \right].$$

We conclude that if $\Delta = \frac{\theta_1 - \theta_0}{\sigma} \leq 1$ and $n \geq \frac{(\theta_1 - \theta_0)^2 \log(1/\delta)}{64\pi\alpha^2\sigma^2} = \frac{KL(\mathbb{P}_{\theta_1}, \mathbb{P}_{\theta_0}) \log(1/\delta)}{64\pi\alpha^2}$ the correct hypothesis is selected. The $1/\alpha$ sufficiency result holds for large enough Δ since one merely needs to observe just one sample since the probability of it coming from θ_0 is negligible.