A Bandit Approach to Multiple Testing with False Discovery Control

Kevin Jamieson^{*,†}, Lalit Jain^{*} {jamieson,lalitj}@cs.washington.edu *Paul G. Allen School of Computer Science & Engineering, University of Washington, Seattle, WA, and [†]Optimizely, San Francisco, CA

Abstract

We propose an adaptive sampling approach for multiple testing which aims to maximize statistical power while ensuring anytime false discovery control. We consider n distributions whose means are partitioned by whether they are below or equal to a baseline (nulls), versus above the baseline (actual positives). In addition, each distribution can be sequentially and repeatedly sampled. Inspired by the multi-armed bandit literature, we provide an algorithm that takes as few samples as possible to exceed a target true positive proportion (i.e. proportion of actual positives discovered) while giving anytime control of the false discovery proportion (nulls predicted as actual positives). Our sample complexity results match known information theoretic lower bounds and through simulations we show a substantial performance improvement over uniform sampling and an adaptive elimination style algorithm. Given the simplicity of the approach, and its sample efficiency, the method has promise for wide adoption in the biological sciences, clinical testing for drug discovery, and online A/B/n testing problems.

1 Introduction

Consider n possible treatments, say, drugs in a clinical trial, where each treatment either has a positive expected effect relative to a baseline (actual positive), or no difference (null), with a goal of identifying as many actual positive treatments as possible. If evaluating the *i*th trial results in a noisy outcome (e.g. due to variance in the actual measurement or just diversity in the population) then given a total measurement budget of B, it is standard practice to execute and average B/n measurements of each treatment, and then output a set of predicted actual positives based on the measured effect sizes. False alarms (i.e. nulls predicted as actual positives) are controlled by either controlling *family-wise error rate (FWER)*, where one bounds the probability that at least one of the predictions is null, or *false discovery rate (FDR)*, where one bounds the expected proportion of the number of predicted nulls to the number of predictions. FDR is a weaker condition than FWER but is often used in favor of FWER because of its higher *statistical power*: more actual positives are output as predictions using the same measurements.

In the pursuit of even greater statistical power, there has recently been increased interest in the biological sciences to reject the uniform allocation strategy of B/n trials to the *n* treatments in favor of an *adaptive* allocation. Adaptive allocations partition the budget *B* into sequential rounds of measurements in which the measurements taken at one round inform which measurements are taken in the next [1, 2]. Intuitively, if the effect size is relatively large for some treatment, fewer trials will be necessary to identify that treatment as an actual positive relative to the others, and that savings of measurements can be allocated towards treatments with smaller effect sizes to boost the signal. However, both [1, 2] employed ad-hoc heuristics which may not only have sub-optimal

32nd Conference on Neural Information Processing Systems (NeurIPS 2018), Montréal, Canada.

statistical power, but also may even result in more false alarms than expected. As another example, in the domain of A/B/n testing in online environments, the desire to understand and maximize click-through-rate across treatments (e.g., web-layouts, campaigns, etc.) has become ubiquitous across retail, social media, and headline optimization for the news. And in this domain, the desire for statistically rigorous adaptive sampling methods with high statistical power are explicit [3].

In this paper we propose an adaptive measurement allocation scheme that achieves near-optimal statistical power subject to FWER or FDR false alarm control. Perhaps surprisingly, we show that even if the treatment effect sizes of the actual positives are identical, adaptive measurement allocation can still substantially improve statistical power. That is, more actual positives can be predicted using an adaptive allocation relative to the uniform allocation under the same false alarm control.

1.1 Problem Statement

Consider *n* distributions (or arms) and a game where at each time *t*, the player chooses an arm $i \in [n] := \{1, \ldots, n\}$ and immediately observes a reward $X_{i,t} \stackrel{iid}{\sim} \nu_i$ where $X_{i,t} \in [0,1]^1$ and $\mathbb{E}_{\nu_i}[X_{i,t}] = \mu_i$. For a *known* threshold μ_0 , define the sets²

$$\mathcal{H}_1 = \{i \in [n] : \mu_i > \mu_0\}$$
 and $\mathcal{H}_0 = \{i \in [n] : \mu_i = \mu_0\} = [n] \setminus \mathcal{H}_1.$

The value of the means μ_i for $i \in [n]$ and the cardinality of \mathcal{H}_1 are *unknown*. The arms (treatments) in \mathcal{H}_1 have means greater than μ_0 (positive effect) while those in \mathcal{H}_0 have means equal to μ_0 (no effect over baseline). At each time t, after the player plays an arm, she also outputs a set of indices $\mathcal{S}_t \subseteq [n]$ that are interpreted as *discoveries* or rejections of the null-hypothesis (that is, if $i \in \mathcal{S}_t$ then the player believes $i \in \mathcal{H}_1$). For as small a $\tau \in \mathbb{N}$ as possible, the goal is to have the number of true detections $|\mathcal{S}_t \cap \mathcal{H}_1|$ be approximately $|\mathcal{H}_1|$ for all $t \ge \tau$, subject to the number of false alarms $|\mathcal{S}_t \cap \mathcal{H}_0|$ being small uniformly over all times $t \in \mathbb{N}$. We now formally define our notions of false alarm control and true discoveries.

Definition 1 (False Discovery Rate, FDR- δ). Fix some $\delta \in (0, 1)$. We say an algorithm is FDR- δ if for all possible problem instances $(\{\nu_i\}_{i=1}^n, \mu_0)$ it satisfies $\mathbb{E}[\frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t|^{\vee 1}}] \leq \delta$ for all $t \in \mathbb{N}$ simultaneously.

Definition 2 (Family-wise Error Rate, FWER- δ). Fix some $\delta \in (0, 1)$. We say an algorithm is FWER- δ if for all possible problem instances $(\{\nu_i\}_{i=1}^n, \mu_0)$ it satisfies $\mathbb{P}(\bigcup_{t=1}^{\infty} \{\mathcal{S}_t \cap \mathcal{H}_0 \neq \emptyset\}) \leq \delta$.

Note FWER- δ implies FDR- δ , the former being a stronger condition than the latter. Allowing a relatively small number of false discoveries is natural, especially if $|\mathcal{H}_1|$ is relatively large. Because μ_0 is known, there exist schemes that guarantee FDR- δ or FWER- δ even if the arm means μ_i and the cardinality of \mathcal{H}_1 are unknown (see Section 2.1). It is also natural to relax the goal of identifying *all* arms in \mathcal{H}_1 to simply identifying a *large proportion* of them.

Definition 3 (True Positive Rate, TPR- δ, τ). *Fix some* $\delta \in (0, 1)$. *We say an algorithm is TPR-* δ, τ on an instance $(\{\nu_i\}_{i=1}^n, \mu_0)$ if $\mathbb{E}[\frac{|\mathcal{S}_i \cap \mathcal{H}_1|}{|\mathcal{H}_1|}] \ge 1 - \delta$ for all $t \ge \tau$.

Definition 4 (Family-wise Probability of Detection, FWPD- δ , τ). Fix some $\delta \in (0, 1)$. We say an algorithm is FWPD- δ , τ on an instance $(\{\nu_i\}_{i=1}^n, \mu_0)$ if $\mathbb{P}(\mathcal{H}_1 \subseteq \mathcal{S}_t) \ge 1 - \delta$ for all $t \ge \tau$.

Note that FWPD- δ , τ implies TPR- δ , τ , the former being a stronger condition than the latter. Also note $\mathbb{P}(\bigcup_{t=1}^{\infty} \{S_t \cap \mathcal{H}_0 \neq \emptyset\}) \leq \delta$ and $\mathbb{P}(\mathcal{H}_1 \subseteq S_{\tau}) \geq 1 - \delta$ together imply $\mathbb{P}(\mathcal{H}_1 = S_{\tau}) \geq 1 - 2\delta$. We will see that it is possible to control the number of false discoveries $|S_t \cap \mathcal{H}_0|$ regardless of how the player selects arms to play. It is the rate at which S_t includes \mathcal{H}_1 that can be thought of as the statistical power of the algorithm, which we formalize as its *sample complexity*:

Definition 5 (Sample Complexity). Fix some $\delta \in (0, 1)$ and an algorithm \mathcal{A} that is FDR- δ (or FWER- δ) over all possible problem instances. Fix a particular problem instance $(\{\nu_i\}_{i=1}^n, \mu_0)$. At each time $t \in \mathbb{N}$, \mathcal{A} chooses an arm $i \in [n]$ to obtain an observation from, and before proceeding to the next round outputs a set $S_t \subseteq [n]$. The sample complexity of \mathcal{A} on this instance is the smallest time $\tau \in \mathbb{N}$ such that \mathcal{A} is TPR- δ , τ (or FWPD- δ , τ).

The sample complexity and value of τ of an algorithm will depend on the particular instance $(\{\nu_i\}_{i=1}^n, \mu_0)$. For example, if $\mathcal{H}_1 = \{i \in [n] : \mu_i = \mu_0 + \Delta\}$ and $\mathcal{H}_0 = [n] \setminus \mathcal{H}_1$, then we expect the

¹All results without modification apply to unbounded, sub-Gaussian random variables.

²All results generalize to the case when $\mathcal{H}_0 = \{i : \mu_i \leq \mu_0\}.$

	False	alarm control
	$FDR-\delta$	$\mathbf{FWER} ext{-}\delta$
	$\max_{t} \mathbb{E}\left[\frac{ \mathcal{S}_{t} \cap \mathcal{H}_{0} }{ \mathcal{S}_{t} \lor 1}\right] \le \delta$	$\mathbb{P}(\bigcup_{t=1}^{\infty} \{ \mathcal{S}_t \cap \mathcal{H}_0 \neq \emptyset \}) \le \delta$
Detection Probability		
$\text{TPR-}\delta,\tau$	Theorem 2	Theorem 5
$\mathbb{E}\left[\frac{ \mathcal{S}_{\tau} \cap \mathcal{H}_1 }{ \mathcal{H}_1 }\right] \ge 1 - \delta$	$n\Delta^{-2}$	$(n-k)\Delta^{-2} + k\Delta^{-2}\log(n-k)$
FWPD- δ, τ	Theorem 3	Theorem 4
$\mathbb{P}(\mathcal{H}_1 \subseteq \mathcal{S}_\tau) \ge 1 - \delta \big $	$(n-k)\Delta^{-2}\log(k) + k\Delta^{-2}$	$(n-k)\Delta^{-2}\log(k) + k\Delta^{-2}\log(n-k)$

Table 1: Informal summary of sample complexity results proved in this paper for $|\mathcal{H}_1| = k$, constant δ (e.g., $\delta = .05$) and $\Delta = \min_{i \in \mathcal{H}_1} \mu_i - \mu_0$. Uniform sampling across all settings requires at least $n\Delta^{-2} \log(n/k)$ samples, and in the FWER+FWPD setting requires $n\Delta^{-2} \log(n)$. Constants and $\log \log$ factors are ignored.

sample complexity to increase as Δ decreases since at least Δ^{-2} samples are necessary to determine whether an arm has mean μ_0 versus $\mu_0 + \Delta$. The next section will give explicit cases.

Remark 1 (Impossibility of stopping time). We emphasize that just as in the non-adaptive setting, at no time can an algorithm stop and declare that it is TPR- δ , τ or FWPD- δ , τ for any finite $\tau \in \mathbb{N}$. This is because there may be an arm in \mathcal{H}_1 with a mean infinitesimally close to μ_0 but distinct such that no algorithm can determine whether it is in \mathcal{H}_0 or \mathcal{H}_1 . Thus, the algorithm must run indefinitely or until it is stopped externally. However, using an anytime confidence bound (see Section 2) one can always make statements like "either $\mathcal{H}_1 \subseteq S_t$, or $\max_{i \in \mathcal{H}_1 \setminus S_t} \mu_i - \mu_0 \leq \epsilon$ " where the ϵ will depend on the width of the confidence interval.

1.2 Contributions and Informal Summary of Main Results

In Section 2 we propose an algorithm that handles all four combinations of {FDR- δ , FWER- δ } and {TPR- δ, τ , FWPD- δ, τ }. A reader familiar with the multi-armed bandit literature would expect an adaptive sampling algorithm to have a large advantage over uniform sampling when there is a large diversity in the means of \mathcal{H}_1 since larger means can be distinguished from μ_0 with fewer samples. However, one should note that to declare all of \mathcal{H}_1 as discoveries, one must sample every arm in \mathcal{H}_0 at *least* as many times as the *most sampled* arm in \mathcal{H}_1 , otherwise they are statistically indistinguishable. As discoveries are typically uncovering rare phenomenon, it is common to assume $|\mathcal{H}_1| = n^{\beta}$ for $\beta \in (0,1)$ [4, 5], or $|\mathcal{H}_1| = o(n)$, but this implies that the number of samples taken from the arms in \mathcal{H}_1 , regardless of how samples are allocated to those arms, will almost always be dwarfed by the number of samples allocated to those arms in \mathcal{H}_0 since there are $\Omega(n)$ of them. This line of reasoning, in part, is what motivates us to give our sample complexity results in terms of the quantities that best describe the contributions from those arms in \mathcal{H}_0 , namely, the cardinality $|\mathcal{H}_1| = n - |\mathcal{H}_0|$, the confidence parameter δ (e.g., $\delta = .05$), and the gap $\Delta := \min_{i \in \mathcal{H}_1} \mu_i - \mu_0$ between the means of the arms in \mathcal{H}_0 and the smallest mean in \mathcal{H}_1 . Reporting sample complexity results in terms of Δ also allows us to compare to known lower bounds in the literature [6, 4, 7, 8]. Nevertheless, we do address the case where the means of \mathcal{H}_1 are varied in Theorem 2.

An informal summary of the sample complexity results proven in this work are found in Table 1 for $|\mathcal{H}_1| = k$. For the least strict setting of FDR+TPR, the upper-left quadrant of Table 1 matches the lower bound of [4], a sample complexity of just $\Delta^{-2}n$. In this FDR+TPR setting (which requires the fewest samples of the four settings), uniform sampling which pulls each arm an equal number of times has a sample complexity of at least $n\Delta^{-2}\log(n/|\mathcal{H}_1|)$ (see Theorem 7 in Appendix G), which exceeds all results in Table 1 demonstrating the statistical power gained by adaptive sampling. For the most strict setting of FWER+FWPD, the lower-right quadrant of Table 1 matches the lower bounds of [7, 9, 8], a sample complexity of $(n - k)\Delta^{-2}\log(k) + k\Delta^{-2}\log(n - k)$. Uniform sampling in the FWER+FWPD setting has a sample complexity lower bounded by $n\Delta^{-2}\log(n)$ (see Theorem 8 in Appendix G). The settings of FDR+FWPD and FWER+TPR are sandwiched between these results, and we are unaware of existing lower bounds for these settings.

All the results in Table 1 are novel, and to the best of our knowledge are the first non-trivial sample complexity results for an adaptive algorithm in the *fixed confidence* setting where a desired confidence δ is set, and the algorithm attempts to minimize the number of samples taken to meet the desired conditions. We also derive tools that we believe may be useful outside this work: for always valid *p*-values (c.f. [3, 10]) we show that FDR is controlled for all times using the Benjamini-Hochberg

procedure [11] (see Lemma 1), and also provide an anytime high probability bound on the false discovery proportion (see Lemma 2).

Finally, as a direct consequence of the theoretical guarantees proven in this work and the empirical performance of the FDR+TPR variant of the algorithm on real data, an algorithm faithful to the theory was implemented and is in use in production at a leading A/B testing platform [12].

1.3 Related work

Identifying arms with means above a threshold, or equivalently, multiple testing via rejecting nullhypotheses with small *p*-values, is an ubiquitous problem in the biological sciences. In the standard setup, each arm is given an equal number of measurements (i.e., a uniform sampling strategy), a *p*-value P_i is produced for each arm where $\mathbb{P}(P_i \leq x) \leq x$ for all $x \in (0, 1]$ and $i \in \mathcal{H}_0$, and a procedure is then run on these *p*-values to declare small *p*-values as rejections of the nullhypothesis, or discoveries. For a set of *p*-values $P_1 \leq P_2 \leq \cdots \leq P_n$, the so-called Bonferroni selection rule selects $S_{BF} = \{i : P_i \leq \delta/n\}$. The fact that FWER control implies FDR control, $\mathbb{E}[|S_{BF} \cap \mathcal{H}_0|] \leq \mathbb{P}(\bigcup_{i \in \mathcal{H}_0} \{P_i \leq \delta/n\}) \leq \delta \frac{|\mathcal{H}_0|}{n} \leq \delta$, suggests that greater statistical power (i.e. more discoveries) could be achieved with procedures designed specifically for FDR. The BH procedure [11] is one such procedure to control FDR and is widely used in practice (with its many extensions [6] and performance investigations [5]). Recall that a uniform measurement strategy where every arm is sampled the same number of times requires $n\Delta^{-2} \log(n/k)$ samples in the FDR+TPR setting, and $n\Delta^{-2} \log(n)$ samples in the FWER+FWPD setting (Theorems 7 and 8 in Appendix G), which can be substantially worse than our adaptive procedure (see Table 1).

Adaptive sequential testing has been previously addressed in the *fixed budget* setting: the procedure takes a sampling budget as input, and the guarantee states that if the given budget is larger than a problem dependent constant, the procedure drives the error probability to zero and the detection probability to one. One of the first methods called *distilled sensing* [13] assumed that arms from \mathcal{H}_0 were Gaussian with mean at most μ_0 , and successively discarded arms after repeated sampling by thresholding at μ_0 -at most the median of the null distribution-thereby discarding about half the nulls at each round. The procedure made guarantees about FDR and TPR, which were later shown to be nearly optimal [4]. Specifically, [4, Corollary 4.2] implies that any procedure with $\max\{FDR + (1 - TPR)\} \leq \delta$ requires a budget of at least $\Delta^{-2}n \log(1/\delta)$, which is consistent with our work. Later, another thresholding algorithm for the fixed budget setting addressed the FWER and FWPD metrics [7]. In particular, if their procedure is given a budget exceeding $(n - |\mathcal{H}_1|)\Delta^{-2}\log(|\mathcal{H}_1|) + |\mathcal{H}_1|\Delta^{-2}\log(n - |\mathcal{H}_1|)$ then the FWER is driven to zero, and the FWPD is driven to one. By appealing to the optimality properties of the SPRT (which knows the distributions precisely) it was argued that this is optimal. These previous works mostly focused on the asymptotic regime as $n \to \infty$ and $|\mathcal{H}_1| = o(n)$.

Our paper, in contrast to these previous works considers the *fixed confidence* setting: the procedure takes a desired FDR (or FWER) and TPR (or FWPD) and aims to minimize the number of samples taken before these constraints are met. To the best of our knowledge, our paper is the first to propose a scheme for this problem in the fixed confidence regime with near-optimal sample complexity guarantees.

A related line of work is the threshold bandit problem, where all the means of \mathcal{H}_1 are assumed to be strictly above a given threshold, and the means of \mathcal{H}_0 are assumed to be strictly below the threshold [14, 15]. To identify this partition, each arm must be pulled a number of times inversely proportional to the square of its deviation from the threshold. This contrasts with our work, where the majority of arms may have means *equal* to the threshold and the goal is to identify arms with means greater than the threshold subject to discovery constraints. If the arms in \mathcal{H}_0 are assumed to be strictly below the threshold it is possible to declare arms as in \mathcal{H}_0 . In our setting we can only ever determine that an arm is in \mathcal{H}_1 and not \mathcal{H}_0 , but it is impossible to detect that an arm is in \mathcal{H}_0 and not in \mathcal{H}_1 .

Note that the problem considered in this paper is very related to the top-k identification problem where the objective is to identify the unique k arms with the highest means with high probability [16, 9, 8]. Indeed, if we knew $|\mathcal{H}_1|$, then our FWER+FWPD setting is equivalent to the top-k problem with $k = |\mathcal{H}_1|$. Lower bounds derived for the top-k problem assume the algorithm has knowledge of the values of the means, just not their indices [16, 8]. Thus, these lower bounds also apply to our setting and are what are referenced in Section 1.2.

Algorithm 1 An algorithm for identifying arms with means above a threshold μ_0 using as few samples as possible subject to false alarm and true discovery conditions. The set S_t is designed to control FDR at level δ . The set \mathcal{R}_t is designed to control FWER at level δ .

Input: Threshold μ_0 , confidence $\delta \in (0, e^{-1}]$, confidence interval $\phi(\cdot, \cdot)$ **Initialize:** Pull each arm $i \in [n]$ once and let $T_i(t)$ denote the number of times arm i has been pulled up to time t. Set $S_{n+1} = \emptyset$, $\mathcal{R}_{n+1} = \emptyset$, and **If** TPR $\xi_t = 1$, and $\nu_t = 1 \quad \forall t$ **Else if** FWPD $\xi_t = \max\{2|S_t|, \frac{5}{3(1-4\delta)}\log(1/\delta)\}$, and $\nu_t = \max\{|S_t|, 1\} \quad \forall t$ **For** $t = n + 1, n + 2, \dots$ **Pull arm** $I_t = \arg\max_{i \in [n] \setminus S_t} \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{\xi_t}),$ **Apply** Benjamini-Hochberg [11] selection at level $\delta' = \frac{\delta}{6.4 \log(36/\delta)}$ to obtain δ FDR-controlled set S_t : $s(k) = \{i \in [n] : \hat{\mu}_{i,T_i(t)} - \phi(T_i(t), \delta' \frac{k}{n}) \ge \mu_0\}, \forall k \in [n]$ $S_{t+1} = s(\hat{k})$ where $\hat{k} = \max\{k \in [n] : |s(k)| \ge k\}$ (if $\exists \hat{k}$ set $S_{t+1} = S_t$) **If** FWER and $S_t \neq \emptyset$: **Pull arm** $J_t = \arg\max_{i \in S_t \setminus \mathcal{R}_t} \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{\nu_t})$ **Apply** Bonferroni-like selection to obtain FWER-controlled set \mathcal{R}_t : $\chi_t = n - (1 - 2\delta'(1 + 4\delta'))|S_t| + \frac{4(1 + 4\delta')}{3} \log(5 \log_2(n/\delta')/\delta')$ $\mathcal{R}_{t+1} = \mathcal{R}_t \cup \{i \in S_t : \hat{\mu}_{i,T_i(t)} - \phi(T_i(t), \frac{\delta}{\lambda_t}) \ge \mu_0\}$

As pointed out by [14], both our setting and the threshold bandit problem can be posed as a combinatorial bandits problem as studied in [17, 18], but such generality leads to unnecessary log factors. The techniques used in this work aim to reduce extraneous log factors, a topic of recent interest in the top-1 and top-k arm identification problem [19, 20, 21, 22, 16, 8]. While these works are most similar to exact identification (FWER+FWPD), there also exist examples of *approximate* top-k where the objective is to find any k means that are each within ϵ of the best k means [9]. Approximate recovery is also studied in a ranking context with a symmetric difference metric [23] which is more similar to the FDR and TPR setting, but neither this nor that work subsumes one another.

Finally, maximizing the number of discoveries subject to a FDR constraint has been studied in a sequential setting in the context of A/B testing with uniform sampling [3]. This work popularized the concept of an always valid *p*-value that we employ here (see Section 2). The work of [10] controls FDR over a *sequence* of independent bandit problems that each outputs at most one discovery. While [10] shares much of the same vocabulary as this paper, the problem settings are very different.

2 Algorithm and Discussion

Throughout, we will assume the existence of an *anytime confidence interval*. Namely, if $\hat{\mu}_{i,t}$ denotes the empirical mean of the first t bounded i.i.d. rewards in [0,1] from arm i, then for any $\delta \in (0,1)$ we assume the existence of a function ϕ such that for any δ we have $\mathbb{P}\left(\bigcap_{t=1}^{\infty}\{|\hat{\mu}_{i,t} - \mu_i| \leq \phi(t,\delta)\}\right) \geq 1-\delta$. We assume that $\phi(t,\delta)$ is non-increasing in its second argument and that there exists an absolute constant c_{ϕ} such that $\phi(t,\delta) \leq \sqrt{\frac{c_{\phi} \log(\log_2(2t)/\delta)}{t}}$. It suffices to define ϕ with this upper bound with $c_{\phi} = 4$ but there are much sharper known bounds that should be used in practice (e.g., they may take empirical variance into account), see [21, 24, 25, 26]. Anytime bounds constructed with such a $\phi(t,\delta)$ are known to be tight in the sense that $\mathbb{P}(\bigcup_{t=1}^{\infty} \{|\hat{\mu}_{i,t} - \mu_i| \geq \phi(t,\delta)\}) \leq \delta$ and that there exists an absolute constant $h \in (0, 1)$ such that $\mathbb{P}(\{|\hat{\mu}_{i,t} - \mu_i| \geq h \phi(t,\delta)$ for infinitely many $t \in \mathbb{N}\}) = 1$ by the Law of the Iterated Logarithm [27].

Consider Algorithm 1. Before entering the for loop, time-dependent variables ξ_t and ν_t are defined that should be updated at each time t for different settings. If just FDR control is desired, the algorithm merely loops over the three lines following the for loop, pulling the arm I_t not in S_t that

has the highest upper confidence bound; such strategies are common for pure-exploration problems [21, 10]. But if FWER control is desired then at most one additional arm J_t is pulled per round to provide an extra layer of filtering and evidence before an arm is added to \mathcal{R}_t . Below we describe the main elements of the algorithm and along the way sketch out the main arguments of the analysis, shedding light on the constants ξ_t and ν_t .

2.1 False alarm control

 S_t is FDR-controlled. In addition to its use as a confident bound, we can also use $\phi(t, \delta)$ to construct:

$$P_{i,t} := \sup\{\alpha \in (0,1] : \widehat{\mu}_{i,t} - \mu_0 \le \phi(t,\alpha)\} \le \log_2(2t) \exp(-t(\widehat{\mu}_{i,t} - \mu_0)^2 / c_\phi).$$
(1)

Proposition 1 of [10] (and the proof of our Lemma 1) shows that if $i \in \mathcal{H}_0$ so that $\mu_i = \mu_0$ then $P_{i,t}$ is an *anytime, sub-uniformly distributed p-value* in the sense that $\mathbb{P}(\bigcup_{t=1}^{\infty} \{P_{i,t} \leq x\}) \leq x$. Sequences that have this property are sometimes referred to as *always-valid p*-values [3]. Note that if $i \in \mathcal{H}_1$ so that $\mu_i > \mu_0$, we would intuitively expect the sequence $\{P_{i,t}\}_{t=1}^{\infty}$ to be point-wise smaller than if $\mu_i = \mu_0$ by the property that $\phi(\cdot, \cdot)$ is non-increasing in its second argument. This leads to the intuitive rule to reject the null-hypothesis (i.e., declare $i \notin \mathcal{H}_0$) for those arms $i \in [n]$ where $P_{i,t}$ is very small. The Benjamini-Hochberg (BH) procedure introduced in [11] proceeds by first sorting the *p*-values so that $P_{(1),T_{(1)}(t)} \leq P_{(2),T_{(2)}(t)} \leq \cdots \leq P_{(n),T_{(n)}(t)}$, then defines $\hat{k} = \max\{k : P_{(k),T_{(k)}(t)} \leq \delta \frac{k}{n}\}$, and sets $\mathcal{S}_{BH} = \{i : P_{i,T_i(t)} \leq \delta \frac{k}{n}\}$. Note that this procedure is identical to defining sets

$$s(k) = \{i : P_{i,T_i(t)} \le \delta \frac{k}{n}\} = \{i : \hat{\mu}_{i,T_i(t)} - \phi(T_i(t), \delta \frac{k}{n}) \ge \mu_0\},\$$

setting $\hat{k} = \max\{k : |s(k)| \ge k\}$, and $S_{BH} = s(\hat{k})$, which is exactly the set $S_t = S_{BH}$ in Algorithm 1. Thus, S_t in Algorithm 1 is equivalent to applying the BH procedure at a level $O(\delta/\log(1/\delta))$ to the anytime *p*-values of (1). We now discuss the extra logarithmic factor.

Because the algorithm is pulling arms sequentially, some dependence between the *p*-values may be introduced. Because the anytime *p*-values are not independent, the BH procedure at level δ does not directly guarantee FDR-control at level δ . However, it has been shown [28] that for even arbitrarily dependent *p*-values the BH procedure at level δ controls FDR at level $\delta \log(n)$ (and that it is nearly tight). Similarly, the following theorem, which may be of independent interest, is a significant improvement when applied to our setting.

Theorem 1. Fix $\delta \in (0, e^{-1})$. Let p_1, \ldots, p_n be random variables such that $\{p_i\}_{i \in \mathcal{H}_0}$ are independent and sub-uniformly distributed so that $\max_{i \in \mathcal{H}_0} \mathbb{P}(p_i \leq x) \leq x$. For any $k \in \{0, 1, \ldots, n\}$, let $R_k := \{i : p_i \leq \delta \frac{k}{n}\}$ and $\widehat{FDP}(R_k) := \frac{\max_{p_i \in R_k} p_i}{|R_k| \vee 1}$.

$$\mathbb{E}\left[\max_{k:\widehat{FDP}(R_k)\leq\delta}FDP(R_k)\right]\leq\frac{|\mathcal{H}_0|\delta}{n}\left(2\log(\frac{2n}{|\mathcal{H}_0|\delta})+\log(8e^5\log(\frac{8n}{|\mathcal{H}_0|\delta}))\right)\\<4\delta\log(9/\delta)$$

In other words, any procedure that chooses a set $\{i : p_i \leq \frac{\delta k}{n}\}$ satisfying $|\{i : p_i \leq \frac{\delta k}{n}\}| \geq k$ is FDR controlled at level $O(\delta \log(1/\delta))$.

Recall, if $\hat{k} = \max\{k : \widehat{FDP}(R_k) \leq \delta\}$ then $\mathbb{E}[FDP(R_{\hat{k}})] \leq \delta$ by the standard BH result. When running the algorithm we recommend using BH at level δ , not level $O(\delta/\log(1/\delta))$. As T_i gets very large, $P_{i,T_i(t)} \rightarrow P_{i,*}$ and we know that if BH is run on $P_{i,*}$ at level δ then FDR would be controlled at level δ . We believe this inflation to be somewhat of an artifact of our proofs.

 \mathcal{R}_t is FWER-controlled. A core obstacle in our analysis is the fact that we don't know the cardinality of \mathcal{H}_1 . If we did know $|\mathcal{H}_1|$ (and equivalently know $|\mathcal{H}_0| = n - |\mathcal{H}_1|$) then a FWER+FWPD algorithm is equivalent to the so-called top-k multi-armed bandit problem [9, 8] and controlling FWER would be relatively simple using a Bonferroni correction:

$$\mathbb{P}\Big(\bigcup_{i\in\mathcal{H}_0}\cup_{t=1}^{\infty}\{\widehat{\mu}_{i,t}-\phi(t,\frac{\delta}{n-|\mathcal{H}_1|})\geq\mu_0\}\Big)\leq\sum_{i\in\mathcal{H}_0}\mathbb{P}\left(\bigcup_{t=1}^{\infty}\{\widehat{\mu}_{i,t}-\phi(t,\frac{\delta}{|\mathcal{H}_0|})\geq\mu_0\}\right)\leq|\mathcal{H}_0|\frac{\delta}{|\mathcal{H}_0|}$$

which implies FWER- δ . Comparing the first expression immediately above to the definition of \mathcal{R}_t in the algorithm, it is clear our strategy is to use $|\mathcal{S}_t|$ as a surrogate for $|\mathcal{H}_1|$. Note that we could

use the bound $|\mathcal{H}_0| = n - |\mathcal{H}_1| \leq n$ to guarantee FWER- δ , but this could be very loose and induce an $n \log(n)$ sample complexity. Using $|\mathcal{S}_t|$ as a surrogate for $|\mathcal{H}_1|$ in \mathcal{R}_t is intuitive because by the FDR guarantee, we know $|\mathcal{H}_1| \geq \mathbb{E}[|\mathcal{S}_t \cap \mathcal{H}_1|] = \mathbb{E}[|\mathcal{S}_t|] - \mathbb{E}[|\mathcal{S}_t \cap \mathcal{H}_0|] \geq (1 - \delta)\mathbb{E}[|\mathcal{S}_t|]$, implying that $|\mathcal{H}_0| = n - |\mathcal{H}_1| \leq n - (1 - \delta)\mathbb{E}[|\mathcal{S}_t|]$ which may be much tighter than n if $\mathbb{E}[|\mathcal{S}_t|] \to |\mathcal{H}_1|$. Because we only know $|\mathcal{S}_t|$ and not its expectation, the extra factors in the surrogate expression used in \mathcal{R}_t are used to ensure correctness with high-probability (see Lemma 7).

2.2 Sampling strategies to boost statistical power

The above discussion about controlling false alarms for S_t and \mathcal{R}_t holds for *any* choice of arms I_t and J_t that may be pulled at time t. Thus, I_t and J_t are chosen in order to minimize the amount of time necessary to add arms into S_t and \mathcal{R}_t , respectively, and optimize the sample complexity.

TPR- δ, τ setting implies $\xi_t = \nu_t = 1$. Define the random set $\mathcal{I} = \{i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \ge \mu_i \ \forall t \in \mathbb{N}\}$. Because ϕ is an anytime confidence bound, $\mathbb{E}[|\mathcal{I}|] \ge (1-\delta)|\mathcal{H}_1|$. If $\Delta = \min_{i \in \mathcal{H}_1} \mu_i - \mu_0$, then $\min_{i \in \mathcal{I}} \mu_i \ge \mu_0 + \Delta$ and we claim that with probability at least $1 - O(\delta)$ (Section C)

$$\begin{split} \sum_{t=1}^{\infty} \mathbf{1}\{I_t \in \mathcal{H}_0, \mathcal{I} \not\subseteq \mathcal{S}_t\} &\leq \sum_{t=1}^{\infty} \mathbf{1}\{I_t \in \mathcal{H}_0, \widehat{\mu}_{I_t, T_{I_t}(t)} + \phi(T_{I_t}(t), \delta) \geq \mu_0 + \Delta\} \\ &\leq c |\mathcal{H}_0| \Delta^{-2} \log(\log(\Delta^{-2}/\delta). \end{split}$$

Thus once this number of samples has been taken, either $\mathcal{I} \subseteq S_t$, or arms in \mathcal{I} will be repeatedly sampled until they are added to S_t since each arm $i \in \mathcal{I}$ has its upper confidence bound larger than those arms in \mathcal{H}_0 by definition. It is clear that an arm in \mathcal{H}_1 that is repeatedly sampled will eventually be added to S_t since its anytime *p*-value of (1) approaches 0 at an exponential rate as it is pulled, and BH selects for low *p*-values. A similar argument holds for J_t and adding arms to \mathcal{R}_t .

Remark 2. While the main objective of Algorithm 1 is to identify all arms with means above a given threshold, we note that prior to adding an arm to S_t in the TPR setting (i.e., when $\xi_t = 1$) Algorithm 1 behaves identically to the nearly optimal best-arm identification algorithm lil'UCB of [21]. Thus, whether the goal is best-arm identification or to identify all arms with means above a certain threshold, Algorithm 1 is applicable.

FWPD- δ, τ setting is more delicate and uses inflated values of ξ_t and ν_t . This time, we must ensure that $\{\mathcal{H}_1 \not\subseteq \mathcal{S}_t\} \implies \max_{i \in \mathcal{H}_1 \cap S_t^c} \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \ge \min_{i \in \mathcal{H}_1 \cap S_t^c} \mu_i \ge \mu_0 + \Delta$. Because then we could argue that either $\mathcal{H}_1 \subset \mathcal{S}_t$, or only arms in \mathcal{H}_1 are sampled until they are added to \mathcal{S}_t (mirroring the TPR argument). As in the FWER setting above, if we knew the value of $|\mathcal{H}_1|$ the we could set $\xi_t \ge |\mathcal{H}_1|$ to observe that

$$\mathbb{P}(\bigcup_{i\in\mathcal{H}_1}\cup_{t=1}^{\infty}\{\widehat{\mu}_{i,t}+\phi(t,\frac{\delta}{\xi_t})<\mu_i\})\leq\sum_{i\in\mathcal{H}_1}\mathbb{P}\left(\cup_{t=1}^{\infty}\{\widehat{\mu}_{i,t}+\phi(t,\frac{\delta}{\xi_t})<\mu_i\}\right)\leq|\mathcal{H}_1|\frac{\delta}{\xi_t}$$

which is less than δ , to guarantee such a condition. But we don't know $|\mathcal{H}_1|$ so we use $|\mathcal{S}_t|$ as a surrogate, resulting in the inflated definitions of ξ_t and ν_t relative to the TPR setting. The key argument is that either $\mathcal{I} \not\subseteq \mathcal{S}_t$ so that $\max_{i \in \mathcal{I} \cap \mathcal{S}_t^c} \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{\xi_t}) \ge \mu_0 + \Delta$ by the definition of \mathcal{I} (since $\xi_t \ge 1$), or $\mathcal{I} \subset \mathcal{S}_t$ and $|\mathcal{S}_t| \ge \frac{1}{2} |\mathcal{H}_1|$ with high probability which implies $\xi_t = \max\{2|\mathcal{S}_t|, \frac{5}{3(1-4\delta)}\log(1/\delta)\} \ge |\mathcal{H}_1|$ and the union bound of the display above holds.

3 Main Results

In what follows, we say $f \leq g$ if there exists a c > 0 that is independent of all problem parameters and $f \leq cg$. The theorems provide an upper bound on the sample complexity $\tau \in \mathbb{N}$ as defined in Section 1.1 for TPR- δ , τ or FWER- δ , τ that holds with probability at least $1 - c\delta$ for different values of c^3 . We begin with the least restrictive setting, resulting in the smallest sample complexity of all the results presented in this work. Note the slight generalization in the below theorem where the means of \mathcal{H}_0 are assumed to be no greater than μ_0 .

Theorem 2 (FDR, TPR). Let $\mathcal{H}_1 = \{i \in [n] : \mu_i > \mu_0\}$, $\mathcal{H}_0 = \{i \in [n] : \mu_i \le \mu_0\}$. Define $\Delta_i = \mu_i - \mu_0$ for $i \in \mathcal{H}_1$, $\Delta = \min_{i \in \mathcal{H}_1} \Delta_i$, and $\Delta_i = \min_{j \in \mathcal{H}_1} \mu_j - \mu_i = \Delta + (\mu_0 - \mu_i)$ for

³ Each theorem relies on different events holding with high probability, and consequently a different c for each. To have c = 1 for each of the four settings, we would have had to define different constants in the algorithm for each setting. We hope the reader forgives us for this attempt at minimizing clutter.

 $i \in \mathcal{H}_0$. For all $t \in \mathbb{N}$ we have $\mathbb{E}[\frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t| \vee 1}] \leq \delta$. Moreover, with probability at least $1 - 2\delta$ there exists a T such that

$$T \lesssim \min\left\{n\Delta^{-2}\log(\log(\Delta^{-2})/\delta),\right.$$
$$\sum_{i \in \mathcal{H}_0} \Delta_i^{-2}\log(\log(\Delta_i^{-2})/\delta) + \sum_{i \in \mathcal{H}_1} \Delta_i^{-2}\log(n\log(\Delta_i^{-2})/\delta)\right\}$$

and $\mathbb{E}\left[\frac{|\mathcal{S}_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right] \ge 1 - \delta$ for all $t \ge T$. Neither argument of the minimum follows from the other.

If the means of \mathcal{H}_1 are very diverse so that $\max_{i \in \mathcal{H}_1} \mu_i - \mu_0 \gg \min_{i \in \mathcal{H}_1} \mu_i - \mu_0$ then the second argument of the min in Theorem 2 can be tighter than the first. But as discussed above, this advantage is inconsequential if $|\mathcal{H}_1| = o(n)$. The remaining theorems are given in terms of just Δ . The $\log \log(\Delta^{-2})$ dependence is due to inverting the ϕ confidence interval and is unavoidable on at least one arm when Δ is unknown a priori due to the law of the iterated logarithm [27, 21, 22].

Informally, Theorem 2 states that if just most true detections suffice while not making too many mistakes, then O(n) samples suffice. The first argument of the min is known to be tight in a minimax sense up to doubly logarithmic factors due to the lower bound of [4]. As a consequence of this work, an algorithm inspired by Algorithm 1 in this setting is now in production at one of the largest A/B testing platforms on the web. The full proof of Theorem 2 (and all others) is given in the Appendix due to space.

Theorem 3 (FDR, FWPD). For all $t \in \mathbb{N}$ we have $\mathbb{E}[\frac{|S_t \cap \mathcal{H}_0|}{|S_t| \vee 1}] \leq \delta$. Moreover, with probability at least $1 - 5\delta$, there exists a T such that

$$T \lesssim (n - |\mathcal{H}_1|)\Delta^{-2}\log(\max\{|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta) + |\mathcal{H}_1|\Delta^{-2}\log(\log(\Delta^{-2})/\delta)$$

and $\mathcal{H}_1 \subset \mathcal{S}_t$ for all t > T.

Here T roughly scales like $(n - |\mathcal{H}_1|) \max\{\log(|\mathcal{H}_1|), \log\log\log(n/\delta)\} + |\mathcal{H}_1|$ where the $\log \log \log (n/\delta)$ term comes from a high probability bound on the false discovery proportion for anytime p-values (in contrast to just expectation) in Lemma 2 that may be of independent interest. While negligible for all practical purposes, it appears unnatural and we suspect that this is an artifact of our analysis. We note that if $|\mathcal{H}_1| = \Omega(\log(n))$ then the sample complexity sheds this awkwardness⁴.

The next two theorems are concerned with controlling FWER on the set \mathcal{R}_t and determining how long it takes before the claimed detection conditions are satisfied on the set \mathcal{R}_t . Note we still have that FDR is controlled on the set S_t but now this set feeds into \mathcal{R}_t .

Theorem 4 (FWER, FWPD). For all t we have $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_0|}{|S_t| \vee 1}\right] \leq \delta$. Moreover, with probability at least $1 - 6\delta$, we have $\mathcal{H}_0 \cap \mathcal{R}_t = \emptyset$ for all $t \in \mathbb{N}$ and there exists a T such that

$$T \lesssim (n - |\mathcal{H}_1|) \Delta^{-2} \log(\max\{|\mathcal{H}_1|, \log\log(n/\delta)\} \log(\Delta^{-2})/\delta) + |\mathcal{H}_1| \Delta^{-2} \log(\max\{n - (1 - 2\delta(1 + 4\delta))|\mathcal{H}_1|, \log\log(n/\delta)\} \log(\Delta^{-2})/\delta)$$

and $\mathcal{H}_1 \subseteq \mathcal{R}_t$ for all $t \geq T$. Note, together this implies $\mathcal{H}_1 = \mathcal{R}_t$ for all $t \geq T$.

Theorem 4 has the strongest conditions, and therefore the largest sample complexity. Ignoring $\log \log \log(n) \text{ factors, } T \text{ roughly scales as } (n - |\mathcal{H}_1|) \log(|\mathcal{H}_1|) + |\mathcal{H}_1| \log(n - (1 - 2\delta(1 + 4\delta))|\mathcal{H}_1|).$ Inspecting the top-k lower bound of [8] where the arms' means in \mathcal{H}_1 are equal to $\mu_0 + \Delta$, the arms' means in \mathcal{H}_0 are equal to μ_0 , and the algorithm has knowledge of the cardinality of \mathcal{H}_1 , a necessary sample complexity of $(n - |\mathcal{H}_1|) \log(|\mathcal{H}_1|) + |\mathcal{H}_1| \log(n - |\mathcal{H}_1|)$ is given. It is not clear whether this small difference of $\log(n - (1 - 2\delta(1 + 4\delta))|\mathcal{H}_1|)$ versus $\log(n - |\mathcal{H}_1|)$ is an artifact of our analysis, or a fundamental limitation when the cardinality $|\mathcal{H}_1|$ is unknown. We now state our final theorem.


Theorem 5 (FWER, TPR). For all t we have $\mathbb{E}\left[\frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t| \vee 1}\right] \leq \delta$. Moreover, with probability at least $1 - 7\delta$ we have $\mathcal{H}_0 \cap \mathcal{R}_t = \emptyset$ for all $t \in \mathbb{N}$ and there exists a T such that

$$T \lesssim (n - |\mathcal{H}_1|)\Delta^{-2}\log(\log(\Delta^{-2})/\delta) \\ + |\mathcal{H}_1|\Delta^{-2}\log(\max\{n - (1 - \eta)|\mathcal{H}_1|, \log\log(n\log(1/\delta)/\delta)\}\log(\Delta^{-2})/\delta) \\ \underbrace{and \mathbb{E}[\frac{|\mathcal{R}_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}] \ge 1 - \delta \text{ for all } t \ge T, \text{ where } \eta = (1 - 3\delta - \sqrt{2\delta\log(1/\delta)/|\mathcal{H}_1|}). \\ \hline {}^4 \text{In the asymptotic } n \text{ regime, it is common to study the case when } |\mathcal{H}_1| = n^{\beta} \text{ for } \beta \in (0, 1) [4, 13]. \end{cases}$$

4 Experiments

The distribution of each arm equals $\nu_i = \mathcal{N}(\mu_i, 1)$ where $\mu_i = \mu_0 = 0$ if $i \in \mathcal{H}_0$, and $\mu_i > 0$ if $i \in \mathcal{H}_1$. We consider three algorithms: *i*) uniform allocation with anytime BH selection as done in Algorithm 1, *ii*) successive elimination (SE) (see Appendix G)⁵ that performs uniform allocation on only those arms that have not yet been selected by BH, and *iii*) Algorithm 1 (UCB). Algorithm

1 and the BH selection rule for all algorithms use $\phi(t, \delta) = \sqrt{\frac{2 \log(1/\delta) + 6 \log \log(1/\delta) + 3 \log(\log(et/2))}{t}}$ from [25, Theorem 8]. In addition, we ran BH at level δ instead of $\delta/(6.4 \log(36/\delta))$ as discussed in section 3. Here we present the sample complexity for TPR+FDR with $\delta = 0.05$ and different parameterizations of μ , n, $|\mathcal{H}_1|$.



The first panel shows an empirical estimate of $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right]$ at each time t for each algorithm, averaged over 1000 trials. The black dashed line on the first panel denotes the level $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right] = 1 - \delta = .95$, and corresponds to the dashed black line on the second panel. The right four panels show the number of samples each algorithm takes before the true positive rate exceeds $1 - \delta = .95$, relative to the number of samples taken by UCB, for various parameterizations. Panels two, three, and four have $\Delta_i = \Delta$ for $i \in \mathcal{H}_1$ while panel five is a case where the Δ_i 's are linear for $i \in \mathcal{H}_1$. While the differences are most clear on the second panel when $|\mathcal{H}_1| = 2 = o(n)$, over all cases UCB uses at least ≈ 3 times fewer samples than uniform and SE. For FDR+TPR, Appendix G shows uniform sampling roughly has a sample complexity that scales like $n\Delta^{-2}\log(\frac{n}{|\mathcal{H}_1|})$ while SE's is upper bounded by $\min\{n\Delta^{-2}\log(\frac{n}{|\mathcal{H}_1|}), (n - |\mathcal{H}_1|)\Delta^{-2}\log(\frac{n}{|\mathcal{H}_1|}) + \sum_{i \in \mathcal{H}_1}\Delta_i^{-2}\log(n)\}$. Comparing with Theorem 2 for the difference cases (i.e., $|\mathcal{H}_1| = 2, \sqrt{n}, n/5$) provides insight into the relative difference between UCB, uniform, and SE on the different panels.

Acknowledgments

This work was informed and inspired by early discussions with Aaditya Ramdas on methods for controlling the false discovery rate (FDR) in multiple testing; we are grateful to have learned from a leader in the field. We also thank him for his careful reading and feedback. We'd also like to thank Martin J. Zhang for his input. We also thank the leading experimentation and A/B testing platform on the web, *Optimizely*, for its support, insight into its customers' needs, and for committing engineering time to implementing this research into their platform [12]. In particular, we thank Whelan Boyd, Jimmy Jin, Pete Koomen, Sammy Lee, Ajith Mascarenhas, Sonesh Surana, and Hao Xia at Optimizely for their efforts.

⁵Inspired by the best-arm identification literature [19].

References

- Linhui Hao, Akira Sakurai, Tokiko Watanabe, Ericka Sorensen, Chairul A Nidom, Michael A Newton, Paul Ahlquist, and Yoshihiro Kawaoka. Drosophila rnai screen identifies host genes important for influenza virus replication. *Nature*, 454(7206):890, 2008.
- [2] GJ Rocklin, TM Chidyausiku, I Goreshnik, A Ford, S Houliston, A Lemak, L Carter, R Ravichandran, VK Mulligan, A Chevalier, CH Arrowsmith, and D Baker. Global analysis of protein folding using massively parallel design, synthesis, and testing. *Science*, 357:168–175, 2017.
- [3] Ramesh Johari, Leo Pekelis, and David J Walsh. Always valid inference: Bringing sequential analysis to a/b testing. *arXiv preprint arXiv:1512.04922*, 2015.
- [4] Rui M Castro. Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli*, 20(4):2217–2246, 2014.
- [5] Maxim Rabinovich, Aaditya Ramdas, Michael I Jordan, and Martin J Wainwright. Optimal rates and tradeoffs in multiple testing. *arXiv preprint arXiv:1705.05391*, 2017.
- [6] A. Ramdas, R. Foygel Barber, M. J. Wainwright, and M. I. Jordan. A Unified Treatment of Multiple Testing with Prior Knowledge. *ArXiv e-prints*, March 2017.
- [7] Matthew L Malloy and Robert D Nowak. Sequential testing for sparse recovery. *IEEE Transactions on Information Theory*, 60(12):7862–7873, 2014.
- [8] Max Simchowitz, Kevin Jamieson, and Benjamin Recht. The simulator: Understanding adaptive sampling in the moderate-confidence regime. In *Conference on Learning Theory*, pages 1794– 1834, 2017.
- [9] Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pages 655–662, 2012.
- [10] Fanny Yang, Aaditya Ramdas, Kevin G Jamieson, and Martin J Wainwright. A framework for multi-a(rmed)/b(andit) testing with online fdr control. In *Advances in Neural Information Processing Systems*, pages 5959–5968, 2017.
- [11] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)*, pages 289–300, 1995.
- [12] Optimizely. Accelerating experimentation through machine learning, https://help. optimizely.com/Build_Campaigns_and_Experiments/Stats_Accelerator, 2018.
- [13] Jarvis Haupt, Rui M Castro, and Robert Nowak. Distilled sensing: Adaptive sampling for sparse detection and estimation. *IEEE Transactions on Information Theory*, 57(9):6222–6235, 2011.
- [14] Andrea Locatelli, Maurilio Gutzeit, and Alexandra Carpentier. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning*, pages 1690–1698, 2016.
- [15] Hideaki Kano, Junya Honda, Kentaro Sakamaki, Kentaro Matsuura, Atsuyoshi Nakamura, and Masashi Sugiyama. Good arm identification via bandit feedback. *arXiv preprint arXiv:1710.06360*, 2017.
- [16] Lijie Chen, Jian Li, and Mingda Qiao. Nearly instance optimal sample complexity bounds for top-k arm selection. In *Artificial Intelligence and Statistics*, pages 101–110, 2017.
- [17] Shouyuan Chen, Tian Lin, Irwin King, Michael R Lyu, and Wei Chen. Combinatorial pure exploration of multi-armed bandits. In Advances in Neural Information Processing Systems, pages 379–387, 2014.
- [18] Tongyi Cao and Akshay Krishnamurthy. Disagreement-based combinatorial pure exploration: Efficient algorithms and an analysis with localization. *arXiv preprint arXiv:1711.08018*, 2017.

- [19] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(Jun):1079–1105, 2006.
- [20] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, pages 1238–1246, 2013.
- [21] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil'ucb: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pages 423–439, 2014.
- [22] Lijie Chen, Jian Li, and Mingda Qiao. Towards instance optimal bounds for best arm identification. In *Conference on Learning Theory*, pages 535–592, 2017.
- [23] Reinhard Heckel, Max Simchowitz, Kannan Ramchandran, and Martin Wainwright. Approximate ranking from pairwise comparisons. In *International Conference on Artificial Intelligence* and Statistics, pages 1057–1066, 2018.
- [24] A. Balsubramani. Sharp Finite-Time Iterated-Logarithm Martingale Concentration. *ArXiv e-prints*, May 2014.
- [25] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1– 42, 2016.
- [26] Ervin Tanczos, Robert Nowak, and Bob Mankoff. A kl-lucb algorithm for large-scale crowd-sourcing. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 5896–5905. 2017.
- [27] Philip Hartman and Aurel Wintner. On the law of the iterated logarithm. *American Journal of Mathematics*, 63(1):169–176, 1941.
- [28] Yoav Benjamini and Daniel Yekutieli. The control of the false discovery rate in multiple testing under dependency. *Annals of statistics*, pages 1165–1188, 2001.
- [29] Maxim Raginsky and Alexander Rakhlin. Lower bounds for passive and active learning. In *Advances in Neural Information Processing Systems*, pages 1026–1034, 2011.
- [30] Alexandre B Tsybakov. Introduction to nonparametric estimation, 2009.
- [31] Pascal Massart et al. The tight constant in the dvoretzky-kiefer-wolfowitz inequality. *The annals of Probability*, 18(3):1269–1283, 1990.

A Analysis Preliminaries

Recall that for any $i \in [n]$ we assume the existence of a function ϕ such that for any δ we have $\mathbb{P}(\bigcap_{t=1}^{\infty} \{ |\hat{\mu}_{i,t} - \mu_i| \le \phi(t, \delta) \}) \ge 1 - \delta$. We assume that $\phi(t, \delta)$ is **non-increasing in its second**

argument and that there exists an absolute constant c_{ϕ} such that $\phi(t, \delta) \leq \sqrt{\frac{c_{\phi} \log(\log_2(2t)/\delta)}{t}}$. It suffices to take $c_{\phi} = 4$ but there are much sharper known bounds that should be used in practice, see [21, 24, 25, 26]. Moreover, define its **inverse** $\phi^{-1}(\epsilon, \delta) = \min\{t : \phi(t, \delta) \leq \epsilon\}$. For the same constant c_{ϕ} , it can be shown that $\phi^{-1}(\epsilon, \delta) \leq c_{\phi}\epsilon^{-2}\log(2\log(\frac{ec_{\phi}\epsilon^{-2}}{\delta})/\delta) \leq c\epsilon^{-2}\log(\log(\epsilon^{-2})/\delta)$ for a sufficiently large constant c (and any $\epsilon, \delta < 1/4$).

The technical challenges in this work revolve around arguments that *avoid union bounds*. By union bounding over all n arms we have

$$\mathbb{P}\left(\bigcup_{i=1}^{n}\bigcup_{t=1}^{\infty}\{|\widehat{\mu}_{i,t}-\mu_{i}| \le \phi(t,\frac{\delta}{n})\}\right) \le \sum_{i=1}^{n} \mathbb{P}\left(\bigcup_{t=1}^{\infty}\{|\widehat{\mu}_{i,t}-\mu_{i}| \le \phi(t,\frac{\delta}{n})\}\right) \le n\frac{\delta}{n} = \delta$$

which says that with probability at least $1 - \delta$, the deviations of *all* arms after any number of samples s are bounded by $\phi(s, \delta/n)$. This is attractive because we can easily upper bound the number of times we need to sample an arm i before its empirical mean is within ϵ of its true mean by $\phi^{-1}(\epsilon, \frac{\delta}{n}) \approx \epsilon^{-2} \log(\frac{n}{\delta} \log(\epsilon^{-2}))$. However, we note that this number of samples scales as $\log(n)$, and this $\log(n)$ will persist in any sample complexity result in any analysis that uses such a union bounding technique. That is, using union bounds in a naive way leads to sample complexities of at least $\Delta^{-2}n \log(n)$, which is no better than uniform sampling when considering FWER+FWPD (see Theorem 8)!

To avoid union bounds, we use techniques developed for best arm and top-k identification for multi-armed bandits [21, 8]. Define the random variable

$$\rho_{i} = \sup\left\{\rho \in (0,1] : \bigcap_{t=1}^{\infty} \{|\hat{\mu}_{i,t} - \mu_{i}| \le \phi(t,\rho)\}\right\}$$
(2)

and note that by the definition of ϕ , we have $\mathbb{P}(\rho_i \leq x) \leq x$ for any $x \in (0, 1]$. That is, each ρ_i is an **independent sub-uniformly distributed random variable**. However, note that the ρ_i random variables are *not p*-values, that is, they are unrelated to $P_{i,t}$. We will often make use of the fact that $\mu_i - \phi(t, \rho_i) \leq \hat{\mu}_{i,t} \leq \mu_i + \phi(t, \rho_i)$ which is simply sandwiching a random quantity by two other random quantities. Furthermore, by definition $|\hat{\mu}_{i,t} - \mu_i| \leq \phi(t, \delta)$ implies $\rho_i \geq \delta$. While a union bound would be equivalent to saying $\mathbb{P}(\bigcup_{i=1}^n \{\rho_i \leq \frac{\delta}{n}\}) \leq \delta$, we avoid union bounds by reasoning about only the statements we need. For instance, instead of guaranteeing *all* the ρ_i are bounded by a single value, we will make guarantees about how *most* behave, and argue that this is sufficient. Examples of such statements are given in events $\mathcal{E}_3, \mathcal{E}_{4,0}, \mathcal{E}_{4,1}$ defined below. This strategy has been applied successfully to remove extraneous log factors (for instance, [8] improved the top-*k* algorithm and analysis initially proposed in [9]).

Our proofs may upper bound an expression by another with a leading constant c, which may become c' on the next line, then c'', and so on. This is meant to indicate that the constant is changing from line to line, but the next display may revert back to using c and this is not the same c as before. All constants are independent of problem parameters. This strategy is principally used to hide the ugliness of inverting the ϕ function, which arises only in the analysis and not the algorithm itself.

The proofs that follow are meant to be read sequentially, as some of the proofs will reuse the same lemmas. All four theorems claim that FDR is controlled for the set S_t in the algorithm. This is independent of the sampling scheme because of the defined anytime *p*-values.

B Anytime FDR for Anytime *p*-values

We begin our analysis by proving a few general statements about anytime *p*-values and FDR control for the set S_t .

For *p*-values p_1, \ldots, p_n and null set $\mathcal{H}_0 \subseteq [n]$ define

$$R_k = \{i : p_i \le \alpha \frac{k}{n}\}$$

$$FDP(R_k) = \frac{\sum_{i \in \mathcal{H}_0} 1\{p_i \le \alpha \frac{k}{n}\}}{|R_k| \lor 1}$$

$$\widehat{FDP}(R_k) = \frac{n \max_{i \in R_k} p_i}{|R_k| \lor 1}$$

The celebrated result of Benjamini-Hochberg says that under the assumption that the null p-values are independent and sub-uniformly distributed, if $\hat{k} = \max\{k : \widehat{FDP}(R_k) \leq \alpha\}$ then $\mathbb{E}[FDP(R_{\hat{k}})] \leq \alpha$. The following theorem provides a bound on the expected false discovery proportion for *any* k such that $\widehat{FDP}(R_k) \leq \alpha$.

Theorem 1. Fix $\alpha \in (0, e^{-1})$. For $n \ge 2$ let p_1, \ldots, p_n be random variables such that $\{p_i\}_{i \in \mathcal{H}_0}$ are independent and sub-uniformly distributed so that $\max_{i \in \mathcal{H}_0} \mathbb{P}(p_i \le x) \le x$. For any $k \in \{0, 1, \ldots, n\}$, if $R_k = \{i : p_i \le \alpha \frac{k}{n}\}$ then

$$\mathbb{E}\left[\max_{k:\widehat{FDP}(R_k)\leq\alpha}FDP(R_k)\right] \leq \frac{|\mathcal{H}_0|\alpha}{n}\left(2\log(\frac{2n}{|\mathcal{H}_0|\alpha}) + \log(8e^5\log(\frac{8n}{|\mathcal{H}_0|\alpha}))\right) \leq 4\alpha\log(9/\alpha).$$

In other words, any procedure that chooses a set $\{i : p_i \leq \frac{\alpha k}{n}\}$ for any k satisfying $|\{i : p_i \leq \frac{\alpha k}{n}\}| \geq k$ is FDR controlled at level $O(\alpha \log(1/\alpha))$.

Proof. Define

$$p_{\ell}^{0} = \min\{t : \sum_{i \in \mathcal{H}_{0}} 1\{p_{i} \le t\} \ge \ell\}.$$

so that p_{ℓ}^0 is the value of the ℓ th largest *p*-value in \mathcal{H}_0 . Note that

$$\max_{k:\widehat{FDP}(R_k)\leq\alpha} FDP(R_k) = \max_{k:\widehat{FDP}(R_k)\leq\alpha} \frac{\sum_{i\in\mathcal{H}_0} 1\{p_i \leq \alpha\frac{k}{n}\}}{|R_k|}$$
$$= \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=1}^{|\mathcal{H}_0|} \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} 1\{p_i \leq \alpha\frac{k}{n}\} \right\}$$
$$\leq \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=1}^{\ell_0} \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} 1\{p_i \leq \alpha\frac{k}{n}\} \right\}$$
$$+ \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} 1\{p_i \leq \alpha\frac{k}{n}\} \right\}$$

for some $\ell_0 \in \{1, \ldots, |\mathcal{H}_0|\}$ to be defined later, where the last line follows from the fact that $\max_k(f(k) + g(k)) \leq \max_k f(k) + \max_k g(k)$. If $\ell = \sum_{i \in \mathcal{H}_0} 1\{p_i \leq \alpha \frac{k}{n}\}$ then $\max_{i \in R_k} p_i$ is at least as large as the ℓ th largest *p*-value in $\{p_i : i \in \mathcal{H}_0\}$ (i.e., p_ℓ^0) and

$$\frac{1\left\{\ell = \sum_{i \in \mathcal{H}_0} 1\left\{p_i \le \alpha \frac{k}{n}\right\}\right\}}{\max_{i \in R_k} p_i} \le \frac{1\left\{\ell = \sum_{i \in \mathcal{H}_0} 1\left\{p_i \le \alpha \frac{k}{n}\right\}\right\}}{p_\ell^0}$$

Let $u_1, \ldots, u_{|\mathcal{H}_0|}$ be iid uniformly distributed random variables on [0, 1]. Note that for every $i \in \mathcal{H}_0$ and $j = 1, \ldots, |\mathcal{H}_0|$ we have $\mathbb{P}(p_i \leq x) \leq x = \mathbb{P}(u_j \leq x)$ where we have used the assumption that each p_i is sub-uniform randomly distributed. Let $u_{(i)}$ be the *i*th largest u_i such that $u_{(1)} \leq u_{(2)} \leq \cdots \leq u_{(|\mathcal{H}_0|)}$. Note that for any $\ell = 1, \ldots, |\mathcal{H}_0|$ we have

$$\mathbb{P}\left(p_{\ell}^{0} \leq x\right) = \mathbb{P}\left(\sum_{i \in \mathcal{H}_{0}} \mathbb{1}\left\{p_{i} \leq x\right\} \geq \ell\right) \leq \mathbb{P}\left(\sum_{i=1}^{|\mathcal{H}_{0}|} \mathbb{1}\left\{u_{i} \leq x\right\} \geq \ell\right) = \mathbb{P}\left(u_{(\ell)} \leq x\right)$$

since each $\mathbf{1}\{p_i \leq x\}$ is an independent Bernoulli random variable with parameter $\mathbb{P}(p_i \leq x) \leq x = \mathbb{P}(u_j \leq x)$ for any $j \in \{1, \ldots, |\mathcal{H}_0|\}$.

We then observe that

$$\mathbb{P}\left(u_{(\ell)} \le x\right) \le \mathbb{P}\left(\sum_{i=1}^{|\mathcal{H}_0|} \mathbf{1}\left\{u_i \le x\right\} \ge \ell\right)$$
$$\le \exp\left(-|\mathcal{H}_0|KL(\frac{\ell}{|\mathcal{H}_0|}, x)\right)$$
$$\le \exp\left(-|\mathcal{H}_0|\frac{\left(\frac{\ell}{|\mathcal{H}_0|} - x\right)^2}{2\ell/|\mathcal{H}_0|}\right)$$
$$= \exp\left(-\frac{\ell}{2}\left(1 - \frac{|\mathcal{H}_0|x}{\ell}\right)^2\right)$$

where the second inequality is Chernoff's bound, and the next line follows from $KL(x,y) \ge \frac{(x-y)^2}{2x}$ for x > y. For each $\ell = 1, 2, \ldots, |\mathcal{H}_0|$ pick $x = \frac{\ell}{2|\mathcal{H}_0|}$ so that $\mathbb{P}\left(p_\ell^0 \le \frac{\ell}{2|\mathcal{H}_0|}\right) \le e^{-\ell/8}$. Also, if we define the event

$$\mathcal{E}_{\ell} := \left\{ p_{\ell}^0 \ge \frac{\ell}{2|\mathcal{H}_0|} \right\}$$

then $\mathbb{P}(\mathcal{E}_{\ell}) \geq 1 - e^{-\ell/8}$. If $\mathcal{E} = \bigcap_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \mathcal{E}_{\ell}$ then

$$\begin{split} \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{\ell \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{|R_k|} \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{\ell \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{|R_k|} \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{\ell \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{n\max_{i\in R_k} p_i} \alpha \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{\ell \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{np_\ell^0} \alpha \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \max_k \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{2|\mathcal{H}_0|\mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{n} \alpha \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \max_k \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \frac{2|\mathcal{H}_0|\mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}}{n} \alpha \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \frac{2|\mathcal{H}_0|}{n} \alpha \max_k \sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\} \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \frac{2|\mathcal{H}_0|}{n} \alpha \exp_k \left\{\sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}\right\} \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \mathbf{1}\left\{\mathcal{E}\right\} \frac{2|\mathcal{H}_0|}{n} \alpha \exp_k \left\{\sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|} \mathbf{1}\left\{\ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\left\{p_i \leq \alpha \frac{k}{n}\right\}\right\}\right\} \\ &\leq \mathbf{1}\left\{\mathcal{E}^c\right\} + \frac{2|\mathcal{H}_0|}{n} \alpha. \end{split}$$

Note that the only randomness on the right-hand-side is $1\{\mathcal{E}^c\}$. After taking expectations on both sides we observe that

$$\mathbb{E}[\mathbf{1}\{\mathcal{E}^{c}\}] = \mathbb{E}\left[\mathbf{1}\{\bigcup_{\ell=\ell_{0}+1}^{|\mathcal{H}_{0}|}\mathcal{E}_{\ell}^{c}\}\right] = \sum_{\ell=\ell_{0}+1}^{|\mathcal{H}_{0}|}\mathbb{P}(\mathcal{E}_{\ell}^{c}) \le \sum_{\ell=\ell_{0}+1}^{|\mathcal{H}_{0}|}e^{-\ell/8} \le 8e^{-\ell_{0}/8}$$

Thus,

$$\mathbb{E}\left[\max_{k:\widehat{FDP}(R_k)\leq\alpha}\sum_{\ell=\ell_0+1}^{|\mathcal{H}_0|}\frac{\ell\mathbf{1}\left\{\ell=\sum_{i\in\mathcal{H}_0}1\{p_i\leq\alpha\frac{k}{n}\}\right\}}{|R_k|}\right]\leq 8e^{-\ell_0/8}+\frac{2|\mathcal{H}_0|}{n}\alpha.$$

On the other hand, for any k with $\widehat{FDP}(R_k) = \frac{n \max_{i \in R_k} p_i}{|R_k| \vee 1} \leq \alpha$ and $\ell = \sum_{i \in \mathcal{H}_0} 1\{p_i \leq \alpha \frac{k}{n}\}$ we have that the ℓ th p-value is in R_k and $\frac{np_\ell^0}{\alpha} \leq \frac{n \max_{i \in R_k} p_i}{\alpha} \leq |R_k|$ where $|R_k| \geq 1$ necessarily. Consequently, $|R_k| \geq \lceil np_\ell^0/\alpha \rceil$ so that

$$\begin{split} \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=1}^{\ell_0} \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\{p_i \leq \alpha \frac{k}{n}\} \right\} \leq \max_{k:\widehat{FDP}(R_k)\leq\alpha} \sum_{\ell=1}^{\ell_0} \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\{p_i \leq \alpha \frac{k}{n}\} \right\} \\ \leq \sum_{\ell=1}^{\ell_0} \max_k \frac{\ell}{|R_k|} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\{p_i \leq \alpha \frac{k}{n}\} \right\} \\ \leq \sum_{\ell=1}^{\ell_0} \max_k \frac{\ell}{\lceil np_\ell^0/\alpha \rceil} \mathbf{1} \left\{ \ell = \sum_{i\in\mathcal{H}_0} \mathbf{1}\{p_i \leq \max_{i\in R_k} p_i\} \right\} \\ = \sum_{\ell=1}^{\ell_0} \frac{\ell}{\lceil np_\ell^0/\alpha \rceil} \end{split}$$

As above, we use the fact that each p_ℓ^0 is stochastically dominated by $u_{(\ell)}$ so that

$$\begin{split} \mathbb{E}\left[\frac{\ell}{\lceil np_{\ell}^{0}/\alpha\rceil}\right] &= \int_{x=0}^{\infty} \mathbb{P}\left(\frac{\ell}{\lceil np_{\ell}^{0}/\alpha\rceil} \ge x\right) dx \\ &= \int_{x=0}^{\infty} \mathbb{P}\left(\lceil np_{\ell}^{0}/\alpha\rceil \le \ell/x\right) dx \\ &= \int_{x=0}^{\infty} \mathbb{P}\left(np_{\ell}^{0}/\alpha \le \ell/x\right) dx \\ &\le \int_{x=0}^{\infty} \mathbb{P}\left(nu_{(\ell)}/\alpha \le \ell/x\right) dx \\ &= \int_{x=0}^{\infty} \mathbb{P}\left(\lceil nu_{(\ell)}/\alpha\rceil \le \ell/x\right) dx \\ &= \mathbb{E}\left[\frac{\ell}{\lceil nu_{(\ell)}/\alpha\rceil}\right] \end{split}$$

since $\lceil a \rceil > b \iff a > \lfloor b \rfloor$. Thus,

$$\mathbb{E}\left[\max_{k:\widehat{FDP}(R_k)\leq\alpha}\sum_{\ell=1}^{\ell_0}\frac{\ell}{|R_k|}\mathbf{1}\left\{\ell=\sum_{i\in\mathcal{H}_0}\mathbf{1}\{p_i\leq\alpha\frac{k}{n}\}\right\}\right]\leq\sum_{\ell=1}^{\ell_0}\mathbb{E}\left[\frac{\ell}{\lceil np_\ell^0/\alpha\rceil}\right]\leq\sum_{\ell=1}^{\ell_0}\mathbb{E}\left[\frac{\ell}{\lceil nu_{(\ell)}/\alpha\rceil}\right].$$

For notational ease, let $m = |\mathcal{H}_0|$. Recall that the PDF of the ℓ -th order statistic of m iid uniform random variables on [0, 1] is given by,

$$\frac{d\mathbb{P}(u_{(\ell)} \le x)}{dx} = \ell \binom{m}{\ell} (1-x)^{m-\ell} x^{\ell-1}.$$

The following simple bound will be useful for small ℓ and y:

$$\mathbb{P}(u_{(\ell)} \le y) = \int_{x=0}^{y} \ell\binom{m}{\ell} (1-x)^{m-\ell} x^{\ell-1} dx \le \binom{m}{\ell} y^{\ell} \le (my)^{\ell}.$$

First we consider the case when $\ell > 1$.

$$\begin{split} \mathbb{E}\left[\frac{\ell}{\lceil nu_{(\ell)}/\alpha\rceil}\right] &= \sum_{k=1}^{\infty} \frac{1}{k} \mathbb{P}\left(\frac{(k-1)\alpha}{n} \le u_{(\ell)} \le \frac{k\alpha}{n}\right) \\ &\leq \mathbb{P}\left(u_{(\ell)} \le \frac{\alpha}{n}\right) + \int_{x=\frac{\alpha}{n}}^{1} \frac{\alpha}{nx} d\mathbb{P}(u_{(\ell)} \le x) \\ &= \mathbb{P}\left(u_{(\ell)} \le \frac{\alpha}{n}\right) + \frac{\alpha}{n} \ell\binom{m}{\ell} \int_{x=\frac{\alpha}{n}}^{1} (1-x)^{m-\ell} x^{\ell-2} dx \\ &\leq \left(\frac{\alpha m}{n}\right)^{\ell} + \frac{\alpha}{n} \frac{\ell\binom{m}{\ell}}{(\ell-1)\binom{m-1}{\ell-1}} \int_{\alpha/n}^{1} (\ell-1)\binom{m-1}{\ell-1} (1-x)^{m-\ell} x^{\ell-2} dx \\ &\leq \left(\frac{\alpha m}{n}\right)^{\ell} + \frac{\alpha}{n} \frac{m!}{(\ell-1)!(m-\ell)!} \frac{(\ell-2)!(m-\ell)!}{(m-1)!} \\ &= \left(\frac{\alpha m}{n}\right)^{\ell} + \frac{\alpha}{n} \frac{m}{\ell-1} \end{split}$$

where going from the third line to the fourth, the integrand is just the PDF of a Beta-distribution so the integral can be bounded away by 1. When $\ell = 1$, the argument is slightly different. Firstly note that the PDF is a decreasing function, hence

$$\mathbb{P}\left(\frac{(k-1)\alpha}{n} \le u_{(1)} \le \frac{k\alpha}{n}\right) \le \frac{\alpha}{n}m\left(1 - \frac{\alpha(k-1)}{n}\right)^{m-1} \\ \le \frac{\alpha}{n}me^{-(m-1)\frac{\alpha(k-1)}{n}}$$

using the fact that $e^{-x} > 1 - x$. Hence,

$$\mathbb{E}\left[\frac{1}{\lceil nu_{(1)}/\alpha\rceil}\right] = \sum_{k=1}^{\lfloor n/\alpha\rfloor} \frac{1}{k} \mathbb{P}\left(\frac{(k-1)\alpha}{n} \le u_{(1)} \le \frac{k\alpha}{n}\right)$$
$$\le \frac{m\alpha}{n} \sum_{k=1}^{\infty} \frac{1}{k} e^{-(m-1)\frac{\alpha(k-1)}{n}}$$
$$= \frac{m\alpha}{n} \left(1 + \sum_{k=1}^{\infty} \frac{1}{k+1} e^{-\frac{(m-1)\alpha}{n}k}\right)$$
$$= \frac{m\alpha}{n} e^{\frac{(m-1)\alpha}{n}} \log\left(\frac{1}{1-e^{-\frac{(m-1)\alpha}{n}}}\right)$$
$$\le 2\frac{m\alpha}{n} \log(\frac{n}{(m-1)\alpha})$$

where the last line holds for $\alpha(m-1)/n \leq e^{-1}$ by observing that $e^x \log(1-e^{-x})/\log(x)$ is increasing and less than 2 at $x = e^{-1}$. On the other hand, if m = 1 then in the above display $\sum_{k=1}^{\lceil n/\alpha \rceil} 1/k \leq \log(\lceil n/\alpha \rceil e) \leq 2 \log(n/\alpha)$ for $\alpha \leq e^{-1}$ and $n \geq 2$. Thus,

$$\mathbb{E}\left[\frac{1}{\lceil nu_{(1)}/\alpha\rceil}\right] = \sum_{k=1}^{\lceil n/\alpha\rceil} \frac{1}{k} \mathbb{P}\left(\frac{(k-1)\alpha}{n} \le u_{(1)} \le \frac{k\alpha}{n}\right)$$
$$\le \begin{cases} 2\frac{m\alpha}{n} \log(\frac{n}{(m-1)\alpha}) & \text{if } m > 1\\ 2\frac{\alpha}{n} \log(\frac{n}{\alpha}) & \text{if } m = 1\\ \le 2\frac{m\alpha}{n} \log(\frac{2n}{m\alpha}) \end{cases}$$

Plugging it all in we get

$$\begin{split} \sum_{\ell=1}^{\ell_0} \mathbb{E}\left[\frac{\ell}{\lceil nu_{(\ell)}/\alpha\rceil}\right] &\leq 2\frac{m\alpha}{n}\log(\frac{2n}{m\alpha}) + \sum_{\ell=2}^{\ell_0}\left(\left(\frac{\alpha m}{n}\right)^\ell + \frac{\alpha}{n}\frac{m}{\ell-1}\right) \\ &\leq 2\frac{m\alpha}{n}\log(\frac{2n}{m\alpha}) + \left(\frac{\alpha m}{n}\right)^2/(1-\frac{\alpha m}{n}) + \frac{\alpha m}{n}\log(\ell_0 e) \end{split}$$

Putting it all together, we conclude that

$$\mathbb{E}\left[\max_{k:\widehat{FDP}(R_k)\leq\alpha}FDP(R_k)\right]\leq 8e^{-\ell_0/8}+\frac{2|\mathcal{H}_0|}{n}\alpha+\sum_{\ell=1}^{\ell_0}\mathbb{E}\left[\frac{\ell}{\lceil nu_{(\ell)}/\alpha\rceil}\right]$$
$$\leq 8e^{-\ell_0/8}+\frac{|\mathcal{H}_0|}{n}\alpha\left(2\log(\frac{2n}{|\mathcal{H}_0|\alpha})+\log(\ell_0e^4)\right)$$
$$\leq \frac{|\mathcal{H}_0|\alpha}{n}\left(2\log(\frac{2n}{|\mathcal{H}_0|\alpha})+\log(8e^5\log(\frac{8n}{|\mathcal{H}_0|\alpha}))\right)$$
$$=O\left(\frac{|\mathcal{H}_0|\alpha}{n}\log(\frac{n}{|\mathcal{H}_0|\alpha})\right)$$
e last lines have taken $\ell_0=8\log(\frac{8n}{|\mathcal{H}_0|}).$

Where the last lines have taken $\ell_0 = 8 \log(\frac{8n}{|\mathcal{H}_0|\alpha})$.

In this section, consider a hypothesis test between \mathcal{H}_0 and \mathcal{H}_1 . Let $\{P_{i,t}\}_{t=1}^{\infty}, P_{i,t} \in (0, 1]$ be a collection of random variables such that for $i \in \mathcal{H}_0$, $\{P_{i,t}\}_{i=1}^{\infty}$ are sub-uniformly distributed anytime *p*-values, i.e. $\mathbb{P}(\bigcup_{t=1}^{\infty} \{P_{i,t} \leq x\}) \leq x$. Let \mathcal{S}_t be the set of discoveries following the Benjamini-Hochberg procedure at confidence level δ at time *t* on the *p*-values $P_{i,T_i(t)}$, $1 \leq i \leq n$. The following lemma employees the previous Theorem to guarantees FDR-control.

Lemma 1. For all times $t \ge 1$, FDR is controlled at level δ so that $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_0|}{|S_t| \vee 1}\right] \le \delta$.

Proof. For any $i \in \mathcal{H}_0$ so that $\mu_i \leq \mu_0$ define $P_{i,*} = \inf_{t \geq 1} P_{i,t}$. Then for any $i \in \mathcal{H}_0$ and $\delta \in (0,1)$

$$\mathbb{P}(P_{i,*} \leq \delta) = \mathbb{P}\left(\bigcup_{t=1}^{\infty} \{P_{i,t} \leq \delta\}\right) = \mathbb{P}\left(\bigcup_{t=1}^{\infty} \{\widehat{\mu}_{i,t} - \phi(t,\delta) \geq \mu_0\}\right)$$
$$\leq \mathbb{P}\left(\bigcup_{t=1}^{\infty} \{\widehat{\mu}_{i,t} - \phi(t,\delta) \geq \mu_i\}\right) \leq \mathbb{P}\left(\bigcup_{t=1}^{\infty} \{|\widehat{\mu}_{i,t} - \mu_i| \geq \phi(t,\delta)\right) \leq \delta.$$

We observe that $\{P_{i,*}\}_{i \in \mathcal{H}_0}$ are sub-uniformly distributed random variables but are also *independent* since they only depend on the rewards of arm $i \in \mathcal{H}_0$. Thus, we may apply the above proposition at time t with $\{P_{i,*}: i \in \mathcal{H}_0\} \cup \{P_{i,T_i(t)}: i \in \mathcal{H}_1\}$ to conclude that FDR would be controlled for any of the prescribed values of k (including the largest, which would be equivalent to BH). We need to show that the BH procedure controls FDR at time t with $\{P_{i,T_i(t)}: i \in \mathcal{H}_0\} \cup \{P_{i,T_i(t)}: i \in \mathcal{H}_1\}$.

As described in Section 2.1, the BH procedure selects the \hat{k} smallest *p*-values where

$$\hat{k} = \max\{k : |\{i : P_{i,T_i(t)} \le \delta \frac{k}{n}\}| \ge k\}.$$

By definition

$$\widehat{k} = |\{i \in \mathcal{H}_0 : P_{i,T_i(t)} \le \delta \frac{k}{n}\} \cup \{i \in \mathcal{H}_1 : P_{i,T_i(t)} \le \delta \frac{k}{n}\}|$$
$$\le |\{i \in \mathcal{H}_0 : P_{i,*} \le \delta \frac{k}{n}\} \cup \{i \in \mathcal{H}_1 : P_{i,T_i(t)} \le \delta \frac{k}{n}\}|$$

since $P_{i,*} \leq P_{i,T_i(t)}$ for all $i \in \mathcal{H}_0$. Thus, since the $\{P_{i,*}\}_{i \in \mathcal{H}_0}$ are independent and sub-uniformly distributed and at least \hat{k} of $\{i \in \mathcal{H}_0 : P_{i,*} \leq \delta \frac{\hat{k}}{n}\} \cup \{i \in \mathcal{H}_1 : P_{i,T_i(t)} \leq \delta \frac{\hat{k}}{n}\}$ are below the threshold $\frac{\widehat{k\delta}}{n}$ we apply Theorem1 to conclude that the FDR of $\{i \in \mathcal{H}_0 : P_{i,*} \leq \delta \frac{\widehat{k}}{n}\} \cup \{i \in \mathcal{H}_1 : P_{i,T_i(t)} \leq \delta \frac{\widehat{k}}{n}\}$ is bounded by δ .

We observe that

$$\begin{split} FDR(\{i \in [n] : P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\}) &= \mathbb{E}\left[\frac{\sum_{i \in \mathcal{H}_{0}} \mathbf{1}\{P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\}}{\sum_{i \in \mathcal{H}_{0}} \mathbf{1}\{P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\} + \sum_{i \in \mathcal{H}_{1}} \mathbf{1}\{P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\}}\right] \\ &\leq \mathbb{E}\left[\frac{\sum_{i \in \mathcal{H}_{0}} \mathbf{1}\{P_{i,*} \leq \delta \frac{\widehat{k}}{n}\}}{\sum_{i \in \mathcal{H}_{0}} \mathbf{1}\{P_{i,*} \leq \delta \frac{\widehat{k}}{n}\} + \sum_{i \in \mathcal{H}_{1}} \mathbf{1}\{P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\}}\right] \\ &= FDR(\{i \in \mathcal{H}_{0} : P_{i,*} \leq \delta \frac{\widehat{k}}{n}\} \cup \{i \in \mathcal{H}_{1} : P_{i,T_{i}(t)} \leq \delta \frac{\widehat{k}}{n}\}) \\ &\leq \delta \end{split}$$

which is precisely what we wished to prove, where the first inequality follows from $\frac{a}{a+c} \leq \frac{a+b}{a+b+c}$ for positive numbers a, b, c, and the fact that $\sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,T_i(t)} \leq \delta_n^{\hat{k}}\} \leq \sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,*} \leq \delta_n^{\hat{k}}\}$ due to $P_{i,*} \leq P_{i,T_i(t)}$ for all $i \in \mathcal{H}_0$.

The following lemma is a stronger result, providing an anytime high-probability bound on the false discovery proportion and the size of the discovered set. The proof follows from considering not the p-values, $P_{i,t}$ at any given time, but rather the random variable which is the worst-case p-values over all time.

Lemma 2. Let $T_i : \mathbb{N} \to \mathbb{N}$ be an arbitrary function for any *i*. Recall $\delta' = \delta/(6.4 \log(36/\delta))$. Then with probability greater than $\geq 1 - \delta'$,

$$\bigcap_{t=1}^{\infty} \left\{ |\mathcal{S}_t| \le \frac{1}{1 - 2\delta'(1 + 4\delta')} |\mathcal{H}_1| + \frac{4(1 + 4\delta')/3}{1 - 2\delta'(1 + 4\delta')} \log(\frac{5 \log_2(n/\delta')}{\delta'}) \right\}$$

and

$$\bigcap_{t=1}^{\infty} \left\{ \frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t| \vee 1} \le \delta' \frac{|\mathcal{H}_0|}{n} + (1+4\delta') \sqrt{\frac{4\delta' \frac{|\mathcal{H}_0|}{n} \log(\frac{\log_2(n/\delta)}{\delta'})}{|\mathcal{S}_t| \vee 1}} + \frac{(1+4\delta') \log(\frac{\log_2(n/\delta')}{\delta'})}{3(|\mathcal{S}_t| \vee 1)} \right\}$$

in particular, these events hold on \mathcal{E}_3 (defined in the proof), which holds with probability greater than $1 - \delta'$.

Proof. As in the proof of Lemma 1, let $P_{i,*} = \inf_{t \ge 1} P_{i,t}$ for all $i \in \mathcal{H}_0$. Define the event,

$$\begin{aligned} \mathcal{E}_3 &:= \left\{ \forall s \in (0,1] : \sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,*} \le s\} \\ &\le s |\mathcal{H}_0| + (1+4s)\sqrt{2\max\{2s, 2\delta'/n\}} |\mathcal{H}_0| \log(\frac{\log_2(n/\delta')}{\delta'}) + \frac{1+4s}{3}\log(\frac{\log_2(n/\delta')}{\delta'}) \right\} \end{aligned}$$

We have that $\mathbb{P}(\mathcal{E}_3) \ge 1 - \delta'$ by applying Lemma 9 found in the appendix with $c = 2\delta'/n$ and $X_i = P_{i,*}$ for $i \in \mathcal{H}_0$ so that $m = |\mathcal{H}_0|$. Note that there exists some threshold $\tau_t \in [0, 1]$ such that BH selects all indices with $P_{i,T_i(t)} \le \tau_t$ so that $\mathcal{S}_t = \mathcal{S}(\tau_t) := \{i \in [n] : P_{i,T_i(t)} \le \tau_t\}$. Then

$$|\mathcal{S}_t \cap \mathcal{H}_0| = \sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,T_i(t)} \le \tau_t\} \le \sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,*} \le \tau_t\}$$

By definition, $\tau_t = \sup\{s \leq 1 : \frac{|\mathcal{S}(s)|}{n}\delta' \geq s\}$, otherwise we take it to be 0, so that $\tau_t \leq \delta'|\mathcal{S}(\tau_t)|/n$. We apply this inequality and \mathcal{E}_3 to observe

$$\begin{split} |\mathcal{S}_{t} \cap \mathcal{H}_{0}| &\leq \tau_{t} |\mathcal{H}_{0}| + (1+4\tau_{t}) \sqrt{2 \max\{2\delta'/n, 2\tau_{t}\}} |\mathcal{H}_{0}| \log(\frac{\log_{2}(n/\delta')}{\delta'}) + \frac{1+4\tau_{t}}{3} \log(\frac{\log_{2}(n/\delta')}{\delta'}) \\ &\leq \delta' |\mathcal{S}(\tau_{t})| \frac{|\mathcal{H}_{0}|}{n} + (1+4\delta') \sqrt{2 \max\{2\delta'/n, 2\delta'|\mathcal{S}(\tau_{t})|\frac{1}{n}\}} |\mathcal{H}_{0}| \log(\frac{\log_{2}(n/\delta')}{\delta'}) + \frac{1+4\delta'}{3} \log(\frac{\log_{2}(n/\delta')}{\delta'}) \\ &= \delta' |\mathcal{S}(\tau_{t})| \frac{|\mathcal{H}_{0}|}{n} + 2(1+4\delta') \sqrt{\delta|\mathcal{S}(\tau_{t})|\frac{|\mathcal{H}_{0}|}{n}} \log(\frac{\log_{2}(n/\delta')}{\delta'}) + \frac{1+4\delta'}{3} \log(\frac{\log_{2}(n/\delta')}{\delta'}) \\ &\leq 2\delta'(1+4\delta') |\mathcal{S}(\tau_{t})|\frac{|\mathcal{H}_{0}|}{n} + \frac{4(1+4\delta')}{3} \log(\frac{\log_{2}(n)}{\delta'}) \end{split}$$

where the last inequality follows from $a + 2\sqrt{ab} + b = (\sqrt{a} + \sqrt{b})^2 \le 2(a+b)$. After rearranging, because $|S_t| = |S_t \cap \mathcal{H}_1| + |S_t \cap \mathcal{H}_0| \le |\mathcal{H}_1| + |S_t \cap \mathcal{H}_0|$ we obtain the theorem. \Box

We will refer to event \mathcal{E}_3 defined in the proof above in what follows.

C Proof of Theorem 2

We will need the following event. Let $a \in \mathbb{R}^n_+$ be a fixed vector to be defined later in the proof and define

$$\mathcal{E}_{4,j} := \left\{ \sum_{i \in \mathcal{H}_j} a_i \log(1/\rho_i) \le 5 \log(1/\delta) \sum_{i \in \mathcal{H}_j} a_i \right\} \quad j \in \{0,1\}$$

Lemma 3. $\min\{\mathbb{P}(\mathcal{E}_{4,0}), \mathbb{P}(\mathcal{E}_{4,1})\} \ge 1 - \delta.$

Proof. The proof is the same for \mathcal{H}_0 and \mathcal{H}_1 so we prove it for $i = 1, \ldots, n$. Because for $i = 1, \ldots, n$ the ρ_i are independent, sub-uniformly distributed random variables, we have that $Z_i = a_i \log(1/\rho_i)$ are independent random variables satisfying $\mathbb{P}(Z_i \ge t) \le \exp(-t/a_i)$. To see this, $\mathbb{P}(\rho_i \le x) \le x$ by definition, and thus $\mathbb{P}(a_i \log(1/\rho_i) \ge a_i \log(1/x)) \le x$. Set $t = a_i \log(1/x)$ and solve for x. The result follows by applying Lemma 8.

We restate the theorem using the above events.

Theorem 2 (FDR, TPR). Let $\mathcal{H}_1 = \{i \in [n] : \mu_i > \mu_0\}$, $\mathcal{H}_0 = \{i \in [n] : \mu_i \le \mu_0\}$. Define $\Delta_i = \mu_i - \mu_0$ for $i \in \mathcal{H}_1$, $\Delta_i = \min_{j \in \mathcal{H}_1} \mu_j - \mu_i$ for $i \in \mathcal{H}_0$, and $\Delta = \min_{i \in \mathcal{H}_1} \Delta_i$. For all t we have $\mathbb{E}[\frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t|}] \le \delta$. Moreover, on $\mathcal{E}_{4,0} \cap \mathcal{E}_{4,1}$ (which holds with probability at least $1 - 2\delta$), there exists a T such that

$$T \lesssim \sum_{i \in \mathcal{H}_0} \Delta_i^{-2} \log(\log(\Delta_i^{-2})/\delta) + \sum_{i \in \mathcal{H}_1} \Delta_i^{-2} \log(n \log(\Delta_i^{-2})/\delta)$$

and $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right] \ge 1 - \delta$ for all $t \ge T$. Also, on the same events there exists a T such that

$$T \lesssim n\Delta^{-2}\log(\log(\Delta^{-2})/\delta)$$

and $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right] \ge 1 - \delta$ for all $t \ge T$. Note, neither follows from the other.

Proof. Define the random set $\mathcal{I} = \{i \in \mathcal{H}_1 : \rho_i \geq \delta\}$. Note, this is equivalent to $\mathcal{I} = \{i \in \mathcal{H}_1 : \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \geq \mu_i \ \forall t \in \mathbb{N}\}$ since if $\rho_i \geq \delta$ then $\phi(T_i(t), \delta) > \phi(T_i(t), \rho_i)$ and

$$\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t),\delta) \ge \mu_i + \phi(T_i(t),\delta) - \phi(T_i(t),\rho_i) \ge \mu_i.$$

Our aim is to show that this "well-behaved" set of indices \mathcal{I} will be added to \mathcal{S}_t in the claimed amount of time. This is sufficient for the TPR result because $\mathbb{E}[\mathcal{I}] = \sum_{i \in \mathcal{H}_1} \mathbb{E}[\mathbf{1}\{\rho_i > \delta\}] = \sum_{i \in \mathcal{H}_1} \mathbb{P}(\rho_i > \delta) \ge (1 - \delta)|\mathcal{H}_1|$ since each ρ_i is a sub-uniformly distributed random variable. To be clear, we are not claiming which *particular* indices of \mathcal{H}_1 will be added to \mathcal{S}_t (indeed, \mathcal{I} is a random set), just that their number exceeds $(1 - \delta)|\mathcal{H}_1|$ in expectation.

Note that $S_t \subseteq S_{t+1}$ for all t so define $T = \min\{t \in \mathbb{N} : \mathcal{I} \subseteq S_{t+1}\}$ if such inclusion ever exists, otherwise let $T = \infty$. Then

$$T = \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{I} \not\subseteq \mathcal{S}_t \}$$
$$= \sum_{t=1}^{\infty} \mathbf{1} \{ I_t \in \mathcal{H}_0, \mathcal{I} \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ I_t \in \mathcal{H}_1 \}$$

We will bound each sum separately, starting with the first.

For any $j \in \mathcal{I}$ we have

$$\widehat{\mu}_{j,T_j(t)} + \phi(T_j(t),\delta) \ge \mu_j + \phi(T_j(t),\delta) - \phi(T_j(t),\rho_j) \ge \mu_j.$$

Thus $\{\mathcal{I} \not\subseteq \mathcal{S}_t\}$ implies that

$$\arg\max_{j\in\mathcal{S}_t^c}\widehat{\mu}_{j,T_j(t)} + \phi(T_j(t),\delta) \ge \min_{j\in\mathcal{I}}\mu_j \ge \min_{j\in\mathcal{H}_1}\mu_j.$$

On the other hand, for any $i \in \mathcal{H}_0$ we have

$$\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t),\delta) \le \mu_i + \phi(T_i(t),\delta) + \phi(T_i(t),\rho_i)$$
$$\le \mu_i + 2\phi(T_i(t),\delta\rho_i)$$

so that $\hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \leq \min_{j \in \mathcal{H}_1} \mu_j = \mu_i + \Delta_i$ whenever $T_i(t) \geq \phi^{-1}(\frac{\Delta_i}{2}, \delta\rho_i)$. If $T_i(t)$ were this large, the arm $i \in \mathcal{H}_0$ could not be pulled because its upper confidence bound would be below the upper confidence bound of an arm $j \in \mathcal{I} \cap S_t^c$. Thus,

$$\begin{split} \sum_{t=1}^{\infty} \mathbf{1} \{ I_t \in \mathcal{H}_0, \mathcal{I} \not\subseteq \mathcal{S}_t \} &\leq \sum_{i \in \mathcal{H}_0} \phi^{-1}(\frac{\Delta_i}{2}, \delta \rho_i) \\ &\leq \sum_{i \in \mathcal{H}_0} c \Delta_i^{-2} \log(\log(\Delta_i^{-2})/(\delta \rho_i)) \\ &\leq \sum_{i \in \mathcal{H}_0} c \Delta_i^{-2} \log(\log(\Delta_i^{-2})/\delta) + c \Delta_i^{-2} \log(1/\rho_i) \\ &\stackrel{\mathcal{E}_{4,0}}{\leq} \sum_{i \in \mathcal{H}_0} c \Delta_i^{-2} \log(\log(\Delta_i^{-2})/\delta) + 5c \Delta_i^{-2} \log(1/\delta) \\ &\leq \sum_{i \in \mathcal{H}_0} c' \Delta_i^{-2} \log(\log(\Delta_i^{-2})/\delta) \\ &\leq c'' |\mathcal{H}_0| \Delta^{-2} \log(\log(\Delta^{-2})/\delta) \end{split}$$

which concludes the upper bound on the first term.

For the second term, consider a time that $I_t \in \mathcal{H}_1$. Recall $\delta' = \delta/(9.6 \log(3000/\delta))$. For any $j \in \mathcal{H}_1$ and arbitrary $k \leq n$ by the BH procedure in the algorithm we have

$$\widehat{\mu}_{j,T_j(t)} - \phi(T_j(t), \delta' \frac{k}{n}) \ge \mu_j - \phi(T_j(t), \delta' \frac{k}{n}) - \phi(T_j(t), \rho_j)$$
$$\ge \mu_j - 2\phi(T_j(t), \delta' \rho_j \frac{k}{n})$$

so that $\hat{\mu}_{j,T_j(t)} - \phi(T_j(t), \delta'\rho_j \frac{k}{n}) \ge \mu_0$ whenever $T_j(t) \ge \phi^{-1}(\frac{\mu_j - \mu_0}{2}, \delta'\rho_j \frac{k}{n})$, guaranteeing its spot in s(k). In the worst case, the arms are added one at a time to s(k), instead as a group. Thus, if π is any map $\mathcal{H}_1 \to \{1, \ldots, |\mathcal{H}_1|\}$ then

$$\begin{split} \sum_{t=1}^{\infty} \mathbf{1}\{I_t \in \mathcal{H}_1\} &\leq \max_{\pi} \sum_{i \in \mathcal{H}_1} \phi^{-1} \left(\frac{\mu_i - \mu_0}{2}, \delta' \rho_i \frac{\pi(i)}{n}\right) \\ &\leq \max_{\pi} \sum_{i \in \mathcal{H}_1} \left(c \Delta_i^{-2} \log(\frac{n}{\pi(i)} \log(\Delta_i^{-2}) / \delta') + c \Delta_i^{-2} \log(1 / \rho_i) \right) \\ &\stackrel{\mathcal{E}_{4,1}}{\leq} \max_{\pi} \sum_{i \in \mathcal{H}_1} \left(c \Delta_i^{-2} \log(\frac{n}{\pi(i)} \log(\Delta_i^{-2}) / \delta) + 5 c \Delta_i^{-2} \log(1 / \delta) \right) \\ &\leq \max_{\pi} \sum_{i \in \mathcal{H}_1} c' \Delta_i^{-2} \log(\frac{n}{\pi(i)} \log(\Delta_i^{-2}) / \delta) \\ &= \sum_{i=1}^{|\mathcal{H}_1|} c' \Delta^{-2} \log(\frac{n}{i} \log(\Delta^{-2}) / \delta). \end{split}$$

The first claimed upper bound of T, the diverse means case, is completed by considering the second to last line and noting trivially that $\pi(i) \ge 1$. To obtain the second upper bound of T, we consider the last line and note that

$$\sum_{i=1}^{k} \log(\frac{n}{i}) \le \int_{0}^{k} \log(\frac{n}{x}) dx = k \log(n) - (x \log x - x) \Big|_{x=0}^{k} = k \log(\frac{n}{k}) + k \le n.$$

Combining the previous two displays we get

$$\begin{split} \sum_{t=1}^{\infty} \mathbf{1}\{I_t \in \mathcal{H}_1\} &\leq \sum_{i=1}^{|\mathcal{H}_1|} c\Delta^{-2} \log(\frac{n}{i} \log(\Delta^{-2})/\delta) \\ &\leq c|\mathcal{H}_1|\Delta^{-2} \log(\log(\Delta^{-2})/\delta) + \sum_{i=1}^{|\mathcal{H}_1|} c\Delta^{-2} \log(\frac{n}{i}) \\ &\leq c|\mathcal{H}_1|\Delta^{-2} \log(\log(\Delta^{-2})/\delta) + cn\Delta^{-2} \\ &\leq c'' n\Delta^{-2} \log(\log(\Delta^{-2})/\delta). \end{split}$$

D Proof of Theorem 3

Our analysis will also make use of the following events, that we will prove each hold with probability at least $1 - \delta$. Let $\beta := \frac{5}{3(1-4\delta)} \log(1/\delta)$ and define:

$$\mathcal{E}_{1} := \left\{ \left| \left\{ i \in \mathcal{H}_{1} : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_{i}| \le \phi(t, \frac{\delta}{\beta}) \} \right\} \right| \ge \frac{1}{2} |\mathcal{H}_{1}| \right\}$$
$$\mathcal{E}_{2,j} := \left\{ \bigcap_{i \in \mathcal{H}_{j}} \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_{i}| \le \phi(t, \frac{\delta}{|\mathcal{H}_{j}|}) \} \right\} \quad j \in \{0, 1\}$$

Also, recall from Lemma 2 the event

$$\begin{aligned} \mathcal{E}_3 &:= \left\{ \forall s \in (0,1] : \sum_{i \in \mathcal{H}_0} \mathbf{1}\{P_{i,*} \le s\} \\ &\le s |\mathcal{H}_0| + (1+4s)\sqrt{2\max\{2s, 2\delta'/n\}|\mathcal{H}_0|\log(\frac{\log_2(n/\delta')}{\delta'})} + \frac{1+4s}{3}\log(\frac{\log_2(n/\delta')}{\delta'}) \right\} \end{aligned}$$

Lemma 2 guarantees that this event holds with probability greater than $1 - \delta$.

We restate the theorem using the above events.

Theorem 3 (FDR, FWPD). For all t we have $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_0|}{|S_t|}\right] \leq \delta$. Moreover, on $\mathcal{E}_1 \cap \mathcal{E}_{2,1} \cap \mathcal{E}_3 \cap \mathcal{E}_{4,0} \cap \mathcal{E}_{4,1}$ (which holds with probability at least $1 - 5\delta$), there exists a T such that

 $T \lesssim (n - |\mathcal{H}_1|)\Delta^{-2}\log(\max\{|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta) + |\mathcal{H}_1|\Delta^{-2}\log(\log(\Delta^{-2})/\delta)$ and $\mathcal{H}_1 \subseteq \mathcal{S}_t$ for all $t \ge T$.

We need a few technical lemmas, specifically, the proof of events \mathcal{E}_1 and $\mathcal{E}_{2,1}$ and their consequences. Lemma 4. $\mathbb{P}(\mathcal{E}_1) \ge 1 - \delta$.

Proof. We break the proof up into two cases based on the cardinality $|\mathcal{H}_1|$. If $|\mathcal{H}_1| \le \beta$ then $\mathbb{P}\left(\bigcup_{i\in\mathcal{H}_1}\bigcup_{t=1}^{\infty}\{|\hat{\mu}_{i,t}-\mu_i|\ge\phi(t,\frac{\delta}{\beta})\}\right) \le \sum_{i\in\mathcal{H}_1}\mathbb{P}\left(\bigcup_{t=1}^{\infty}\{|\hat{\mu}_{i,t}-\mu_i|\ge\phi(t,\frac{\delta}{\beta})\}\right) \le \delta\frac{|\mathcal{H}_1|}{\beta}$

which is less than δ by the case definition. So in what follows, assume that $|\mathcal{H}_1| > \beta$. By definition

$$\mathbb{P}\left(\bigcup_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_i| \ge \phi(t,\delta) \} \right) = \mathbb{P}(\rho_i \le \delta) \le \delta.$$

By Bernstein's inequality, with probability at least $1 - \delta$

$$\sum_{i \in \mathcal{H}_1} \mathbf{1}\{\rho_i \le \delta\} \le \delta |\mathcal{H}_1| + \sqrt{2\delta} |\mathcal{H}_1| \log(1/\delta) + \frac{1}{3} \log(1/\delta)$$
$$\le \delta |\mathcal{H}_1| + 2\sqrt{\frac{1}{2}\delta} |\mathcal{H}_1| \log(1/\delta) + (1 - \frac{1}{3})\frac{1}{2} \log(1/\delta)$$
$$\le 2\delta |\mathcal{H}_1| + \frac{5}{6} \log(1/\delta)$$

where the last line follows from $a + 2\sqrt{ab} + b = (\sqrt{a} + \sqrt{b})^2 \le 2a + 2b$, which implies

$$\sum_{i \in \mathcal{H}_1} \mathbf{1}\{\rho_i > \delta\} \ge (1 - 2\delta)|\mathcal{H}_1| - \frac{5}{6}\log(1/\delta)$$
$$\ge \frac{1}{2}|\mathcal{H}_1|$$

where we use the fact that $|\mathcal{H}_1| > \beta = \frac{5}{3(1-4\delta)} \log(1/\delta)$. Combining these two results, and noting that at most one of the cases $|\mathcal{H}_1| > \beta$ or $|\mathcal{H}_1| \le \beta$ can be true, we obtain that $\mathbb{P}(\mathcal{E}_1) \ge 1 - \delta$. **Lemma 5.** $\min\{\mathbb{P}(\mathcal{E}_{2,0}), \mathbb{P}(\mathcal{E}_{2,1})\} \ge 1 - \delta$.

Proof. The result follows from a union bound:

$$\mathbb{P}(\mathcal{E}_{2,1}^{c}) = \mathbb{P}\left(\left\{\bigcup_{i\in\mathcal{H}_{1}}\bigcup_{t=1}^{\infty}\{|\widehat{\mu}_{i,t}-\mu_{i}|\leq\phi(t,\frac{\delta}{|\mathcal{H}_{1}|})\}\right\}\right)$$
$$\leq \sum_{i\in\mathcal{H}_{1}}\mathbb{P}\left(\left\{\bigcup_{t=1}^{\infty}\{|\widehat{\mu}_{i,t}-\mu_{i}|\leq\phi(t,\frac{\delta}{|\mathcal{H}_{1}|})\}\right\}\right) \leq \sum_{i\in\mathcal{H}_{1}}\frac{\delta}{|\mathcal{H}_{1}|}\leq\delta.$$

The proof that $\mathbb{P}(\mathcal{E}_{2,0}) \geq 1 - \delta$ follows analogously.

The next lemma shows an important consequence of these events holding. Lemma 6. If $\mathcal{E}_1 \cap \mathcal{E}_{2,1}$ then for all t

$$\mathcal{H}_1 \nsubseteq \mathcal{S}_t \implies \max_{i \in \mathcal{H}_1 \cap \mathcal{S}_t^c} \widehat{\mu}_{i, T_i(t)} + \phi(T_i(t), \frac{\delta}{2|\mathcal{S}_t| \lor \beta}) \ge \min_{i \in \mathcal{H}_1} \mu_i.$$

Proof. Define the random set $\mathcal{I}_t = \left\{ i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{2|\mathcal{S}_t|\vee\beta}) \ge \mu_i \right\}$. We will prove that on $\mathcal{E}_1 \cap \mathcal{E}_{2,1}$ we have $\mathcal{I}_t \cap \mathcal{S}_t^c \neq \emptyset$ which implies the result. First we use the fact that $2|\mathcal{S}_t|\vee\beta \ge \beta$ so that

$$\begin{aligned} |\mathcal{I}_t| &= \left| \left\{ i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{2|\mathcal{S}_t|\vee\beta}) \ge \mu_i \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{\beta}) \ge \mu_i \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,T_i(t)} - \mu_i| \le \phi(T_i(t), \frac{\delta}{\beta}) \} \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_i| \le \phi(t, \frac{\delta}{\beta}) \} \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_i| \le \phi(t, \frac{\delta}{\beta}) \} \right\} \right| \end{aligned}$$

Given $|\mathcal{I}_t| \geq \frac{1}{2}|\mathcal{H}_1|$, if $|\mathcal{S}_t| < \frac{1}{2}|\mathcal{H}_1|$ then $|\mathcal{S}_t| < |\mathcal{I}_t|$ which implies $\mathcal{S}_t^c \cap \mathcal{I}_t \neq \emptyset$. On the other hand, if $|\mathcal{S}_t| \geq |\mathcal{H}_1|/2$ then we use the fact that $2|\mathcal{S}_t| \lor \beta \geq 2|\mathcal{S}_t| \geq |\mathcal{H}_1|$ to observe

$$\begin{aligned} |\mathcal{I}_t| &= \left| \left\{ i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{2|\mathcal{S}_t|\vee\beta}) \ge \mu_i \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{|\mathcal{H}_1|}) \ge \mu_i \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,T_i(t)} - \mu_i| \le \phi(T_i(t), \frac{\delta}{|\mathcal{H}_1|}) \} \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_i| \le \phi(t, \frac{\delta}{|\mathcal{H}_1|}) \} \right\} \right| \\ &\geq \left| \left\{ i \in \mathcal{H}_1 : \bigcap_{t=1}^{\infty} \{ |\widehat{\mu}_{i,t} - \mu_i| \le \phi(t, \frac{\delta}{|\mathcal{H}_1|}) \} \right\} \right| \end{aligned}$$

which implies $\mathcal{I}_t = \mathcal{H}_1$, thus $\mathcal{I}_t \cap \mathcal{S}_t^c = \mathcal{H}_1 \cap \mathcal{S}_t^c$ which is non-empty by assumption.

We are now ready to prove Theorem 3.

Proof. We proceed similarly to Theorem 1. Note that $S_t \subseteq S_{t+1}$ for all t so define $T = \min\{t \in \mathbb{N} : H_1 \subseteq S_{t+1}\}$ if such inclusion ever exists, otherwise let $T = \infty$. Then

$$T = \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{H}_1 \not\subseteq \mathcal{S}_t \}$$
$$= \sum_{t=1}^{\infty} \mathbf{1} \{ I_t \in \mathcal{H}_0, \mathcal{H}_1 \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ I_t \in \mathcal{H}_1 \}$$

Note that we are in the TPR setting, like Theorem 2, and so the upperbound $\sum_{t=1}^{\infty} \mathbf{1}\{I_t \in \mathcal{H}_1\} \le cn\Delta^{-2}\log(\log(\Delta^{-2})/\delta)$ applies here as well. Thus, we only need to bound the first sum.

Lemma 6 (which requires $\mathcal{E}_1 \cap \mathcal{E}_{2,1}$) says that if there is at least one arm from \mathcal{H}_1 not in \mathcal{S}_t then the largest upper confidence bound of some arm in $\mathcal{H}_1 \cap \mathcal{S}_t^c$ is *at least* as large as $\min_{j \in \mathcal{H}_1} \mu_j \ge \mu_0 + \Delta$. Thus, $\{\mathcal{H}_1 \not\subseteq \mathcal{S}_t\}$ implies that

$$\arg\max_{i\in\mathcal{S}_t^c}\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t),\delta) \ge \min_{i\in\mathcal{H}_1}\mu_i = \mu_0 + \Delta.$$

On the other hand, for $\kappa = \frac{1}{1-2\delta'(1+4\delta')} |\mathcal{H}_1| + \frac{4(1+4\delta')/3}{1-2\delta'(1+4\delta')} \log(\frac{5\log_2(n/\delta')}{\delta'})$, we have $|\mathcal{S}_t| \le \kappa$ from Lemma 2 (which requires \mathcal{E}_3), and for any $i \in \mathcal{H}_0$ we have

$$\begin{aligned} \widehat{\mu}_{i,T_{i}(t)} + \phi(T_{i}(t), \frac{\delta}{|\mathcal{S}_{t}| \lor \beta}) &\leq \mu_{i} + \phi(T_{i}(t), \frac{\delta}{|\mathcal{S}_{t}| \lor \beta}) + \phi(T_{i}(t), \rho_{i}) \\ &\leq \mu_{i} + \phi(T_{i}(t), \frac{\delta}{\kappa}) + \phi(T_{i}(t), \rho_{i}) \\ &\leq \mu_{i} + 2\phi(T_{i}(t), \frac{\delta\rho_{i}}{\kappa}) \\ &\leq \mu_{0} + 2\phi(T_{i}(t), \frac{\delta\rho_{i}}{\kappa}) \end{aligned}$$

so that $\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{|\mathcal{S}_t| \lor \beta}) \le \mu_0 + \Delta$ whenever $T_i(t) \ge \phi^{-2}(\frac{\Delta}{2}, \frac{\delta \rho_i}{\kappa})$. Thus,

$$\sum_{t=1} \mathbf{1} \{ I_t \in \mathcal{H}_0, \mathcal{H}_1 \not\subseteq \mathcal{S}_t \} \leq \sum_{i \in \mathcal{H}_0} \phi^{-2}(\frac{\Delta}{2}, \frac{\delta \rho_i}{\kappa})$$

$$\leq \sum_{i \in \mathcal{H}_0} c \Delta^{-2} \log(\kappa \log(\Delta^{-2})/\delta) + c \Delta^{-2} \log(1/\rho_i)$$

$$\stackrel{\mathcal{E}_{4,0}}{\leq} \sum_{i \in \mathcal{H}_0} c' \Delta^{-2} \log(\max\{|\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta) + 5c \Delta^{-2} \log(1/\delta)$$

$$\leq |\mathcal{H}_0| c'' \Delta^{-2} \log(\max\{|\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta)$$

We conclude that

and \mathcal{H}_1

$$T = \sum_{t=1}^{T} \mathbf{1} \{ I_t \in \mathcal{H}_0, \mathcal{H}_1 \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ I_t \in \mathcal{H}_1, \mathcal{H}_1 \not\subseteq \mathcal{S}_t \}$$

$$\leq (n - |\mathcal{H}_1|) c \Delta^{-2} \log(\max\{|\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta) + cn\Delta^{-2} \log(\log(\Delta^{-2})/\delta)$$

$$\leq (n - |\mathcal{H}_1|) c'\Delta^{-2} \log(\max\{|\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta) + c'|\mathcal{H}_1|\Delta^{-2} \log(\log(\Delta^{-2})/\delta).$$

E Proof of Theorem 4

We restate the theorem using the above events.

Theorem 4 (FWER, FWPD). For all t we have $\mathbb{E}[\frac{|\mathcal{S}_t \cap \mathcal{H}_0|}{|\mathcal{S}_t|}] \leq \delta$. Moreover, on $\mathcal{E}_1 \cap \mathcal{E}_{2,0} \cap \mathcal{E}_{2,1} \cap \mathcal{E}_3 \cap \mathcal{E}_{4,0} \cap \mathcal{E}_{4,1}$ (which holds with probability at least $1 - 6\delta$), we have $\mathcal{H}_0 \cap \mathcal{R}_t = \emptyset$ for all $t \in \mathbb{N}$ and there exists a T such that

$$T \leq (n - |\mathcal{H}_1|)\Delta^{-2}\log(\max\{|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta) \\ + |\mathcal{H}_1|\Delta^{-2}\log(\max\{n - (1 - 2\delta(1 + 4\delta))|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta) \\ \subseteq \mathcal{R}_t \text{ for all } t \geq T. \text{ Note, together this implies } \mathcal{H}_1 = \mathcal{R}_t \text{ for all } t \geq T.$$

The following lemma shows that a tight control on the size of S_t allows us to conclude a FWER. Lemma 7. If $\mathcal{E}_3 \cap \mathcal{E}_{2,0}$ holds then $\mathcal{R}_t \cap \mathcal{H}_0 = \emptyset$ for all t.

Proof. By Lemma 2 (which requires \mathcal{E}_3) we have $|\mathcal{S}_t| \leq \frac{1}{1-2\delta'(1+4\delta')}|\mathcal{H}_1| + \frac{4(1+4\delta')/3}{1-2\delta'(1+4\delta')}\log(\frac{5\log_2(n/\delta')}{\delta'}) = \frac{|\mathcal{H}_1|+\eta}{1-2\delta'(1+4\delta')}$ for all times t and $\eta = \frac{4(1+4\delta')}{3}\log(5\log_2(n)/\delta')$. This implies

$$n - (1 - 2\delta'(1 + 4\delta'))|\mathcal{S}_t| + \eta \ge n - |\mathcal{H}_1| = |\mathcal{H}_0|$$
(3)

but for $i \in S_t \cap \mathcal{H}_0 \subseteq \mathcal{H}_0$ we have that $\mathcal{E}_{2,0}$ applies so

$$\widehat{\mu}_{i,T_{i}(t)} - \phi(T_{i}(t), \frac{\delta}{n - (1 - 2\delta'(1 + 4\delta'))|\mathcal{S}_{t}| + \eta}) \leq \widehat{\mu}_{i,T_{i}(t)} - \phi(T_{i}(t), \frac{\delta}{|\mathcal{H}_{0}|}) \stackrel{\mathcal{E}_{2,0}}{\leq} \mu_{i} \leq \mu_{0}$$

where the last inequality holds because $\max_{i \in \mathcal{H}_0} \mu_i \leq \mu_0$. Thus, no arms from \mathcal{H}_0 will be added to \mathcal{R}_t .

Now that we have FWER control, we need to show that all the arms in \mathcal{H}_1 are added to \mathcal{R}_t in the claimed amount of time.

Proof. Note that $\mathcal{R}_t \subseteq \mathcal{R}_{t+1}$ for all t so define $T = \min\{t \in \mathbb{N} : \mathcal{H}_1 \subseteq \mathcal{R}_{t+1}\}$ if such inclusion ever exists, otherwise let $T = \infty$. Noting that $\mathcal{R}_t \subseteq \mathcal{S}_t$ we have

$$T = \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{H}_1 \not\subseteq \mathcal{R}_t \}$$
$$= \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{H}_1 \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ \mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t \}$$
$$= \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{H}_1 \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ J_t \in \mathcal{H}_0, \mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t \} + \mathbf{1} \{ J_t \in \mathcal{H}_1, \mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t \}$$

First sum. Note that we are in the FWPD setting, so the selection rule for I_t identical to that of the setting of Theorem 3, and $\sum_{t=1}^{\infty} \mathbf{1}\{\mathcal{H}_1 \not\subseteq \mathcal{S}_t\}$ is precisely what is bounded in the proof of Theorem 3 (which requires $\mathcal{E}_1 \cap \mathcal{E}_{2,1} \cap \mathcal{E}_3 \cap \mathcal{E}_{4,0} \cap \mathcal{E}_{4,1}$). Thus,

$$\sum_{t=1}^{T} \mathbf{1}\{\mathcal{H}_{1} \not\subseteq \mathcal{S}_{t}\}$$

$$\leq (n - |\mathcal{H}_{1}|)c'\Delta^{-2}\log(\max\{|\mathcal{H}_{1}|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta) + c'|\mathcal{H}_{1}|\Delta^{-2}\log(\log(\Delta^{-2})/\delta).$$

Second sum. Recall that $J_t = \arg \min_{i \in S_t \setminus \mathcal{R}_t} \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{|\mathcal{S}_t|})$. Now, on the event $\{\mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t\}$ we have that $|\mathcal{H}_1| \leq |\mathcal{S}_t|$ and that there exists a $j \in \mathcal{H}_1 \cap (\mathcal{S}_t \setminus \mathcal{R}_t)$ such that

$$\begin{aligned} \widehat{\mu}_{j,T_{j}(t)} + \phi(T_{j}(t), \frac{\delta}{|\mathcal{S}_{t}|}) &\geq \widehat{\mu}_{j,T_{j}(t)} + \phi(T_{j}(t), \frac{\delta}{|\mathcal{H}_{1}|}) \\ &\stackrel{\mathcal{E}_{2,1}}{&\geq} \mu_{j} \\ &\geq \mu_{0} + \Delta. \end{aligned}$$

On the other hand, for $\kappa = \frac{1}{1-2\delta'(1+4\delta')}|\mathcal{H}_1| + \frac{4(1+4\delta')/3}{1-2\delta'(1+4\delta')}\log(\frac{5\log_2(n/\delta')}{\delta'})$, we have $|\mathcal{S}_t| \le \kappa$ from Lemma 2 (which requires \mathcal{E}_3), and for any $i \in \mathcal{H}_0 \cap (\mathcal{S}_t \setminus \mathcal{R}_t)$ we have

$$\begin{aligned} \widehat{\mu}_{i,T_{i}(t)} + \phi(T_{i}(t), \frac{\delta}{|\mathcal{S}_{i}|}) &\leq \mu_{i} + \phi(T_{i}(t), \frac{\delta}{\kappa}) + \phi(T_{i}(t), \rho_{i}) \\ &\leq \mu_{i} + 2\phi(T_{i}(t), \frac{\delta\rho_{i}}{\kappa}) \\ &\leq \mu_{0} + 2\phi(T_{i}(t), \frac{\delta\rho_{i}}{\kappa}) \end{aligned}$$

so that $\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \frac{\delta}{|\mathcal{S}_t|}) \le \mu_0 + \Delta$ whenever $T_i(t) \ge \phi^{-2}(\frac{\Delta}{2}, \frac{\delta \rho_i}{\kappa})$. By an identical argument to that made in the proof of Theorem 3 we very crudely have the bound

$$\sum_{t=1}^{\infty} \mathbf{1}\{J_t \in \mathcal{H}_0, \mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t\} \leq \sum_{i \in \mathcal{H}_0} \phi^{-2}(\frac{\Delta}{2}, \frac{\delta \rho_i}{\kappa})$$
$$\leq c|\mathcal{H}_0|\Delta^{-2}\log(\max\{|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta).$$

Third sum. An arm $j \in \mathcal{H}_1 \cap (\mathcal{S}_t \setminus \mathcal{R}_t)$ is accepted into \mathcal{R}_t if $\hat{\mu}_{i,T_i(t)} - \phi(T_i(t), \frac{\delta}{\chi_t}) \ge \mu_0$ where $\chi_t = n - (1 - 2\delta'(1 + 4\delta'))|\mathcal{S}_t| + \frac{4(1+4\delta')}{3}\log(5\log_2(n/\delta')/\delta')$. On the event $\{\mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t\}$ we have $\chi_t \le n - (1 - 2\delta'(1 + 4\delta'))|\mathcal{H}_1| + \frac{4(1+4\delta')}{3}\log(5\log_2(n/\delta')/\delta') =: u$. Thus, for $j \in \mathcal{H}_1 \cap (\mathcal{S}_t \setminus \mathcal{R}_t)$

$$\begin{aligned} \widehat{\mu}_{j,T_{j}(t)} - \phi(T_{j}(t), \frac{\delta}{\chi_{t}}) &\geq \widehat{\mu}_{j,T_{j}(t)} - \phi(T_{j}(t), \frac{\delta}{u}) \\ &\geq \mu_{j} - \phi(T_{j}(t), \frac{\delta}{u}) - \phi(T_{j}(t), \rho_{j}) \\ &\geq \mu_{j} - 2\phi(T_{j}(t), \frac{\delta\rho_{j}}{u}) \\ &\geq \mu_{0} + \Delta - 2\phi(T_{j}(t), \frac{\delta\rho_{j}}{u}) \end{aligned}$$

so that $\hat{\mu}_{j,T_j(t)} - \phi(T_j(t), \frac{\delta}{\chi_t}) \ge \mu_0$ whenever $T_j(t) \ge \phi^{-1}(\frac{\Delta}{2}, \frac{\delta\rho_j}{u})$. By the same arguments made throughout these proofs we have

$$\begin{split} \sum_{t=1}^{\infty} \mathbf{1} \{ J_t \in \mathcal{H}_1, \mathcal{H}_1 \not\subseteq \mathcal{R}_t, \mathcal{H}_1 \subseteq \mathcal{S}_t \} &\leq \sum_{j \in \mathcal{H}_1} \phi^{-1}(\frac{\Delta}{2}, \frac{\delta \rho_j}{u}) \\ &\leq c |\mathcal{H}_1| \Delta^{-2} \log(u \log(\Delta^{-2})/\delta) \\ &\leq c' |\mathcal{H}_1| \Delta^{-2} \log(\max\{n - (1 - 2\delta(1 + 4\delta)) |\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta). \end{split}$$

Summing all three sums together yields the result.

F Proof of Theorem 5

We define a new event,

$$\mathcal{E}_5 := \left\{ \sum_{i \in \mathcal{H}_1} \mathbf{1}\{\rho_i \le \delta\} \le \delta |\mathcal{H}_1| + \sqrt{2\delta |\mathcal{H}_1| \log(1/\delta)} + \frac{1}{3} \log(1/\delta) \right\}$$

By a direct application of Bernstein's inequality, this holds with probability greater than $1 - \delta$.

Theorem 5 (FWER, TPR). For all t we have $\mathbb{E}[\frac{|S_t \cap \mathcal{H}_0|}{|S_t|}] \leq \delta$. Moreover, on $\mathcal{E}_1 \cap \mathcal{E}_{2,1} \cap \mathcal{E}_3 \cap \mathcal{E}_{4,0} \cap \mathcal{E}_{4,1} \cap \mathcal{E}_{2,0} \cap \mathcal{E}_5$ (which holds with probability at least $1 - 7\delta$), we have $\mathcal{H}_0 \cap \mathcal{R}_t = \emptyset$ for all $t \in \mathbb{N}$ and there exists a T such that

$$T \leq (n - |\mathcal{H}_1|)\Delta^{-2}\log(\log(\Delta^{-2})/\delta) + |\mathcal{H}_1|\Delta^{-2}\log(\max\{n - (1 - 2\delta)|\mathcal{H}_1|, \log\log(n/\delta)\}\log(\Delta^{-2})/\delta)$$

and $\mathbb{E}\left[\frac{|\mathcal{R}_t \cap \mathcal{H}_1|}{|\mathcal{H}_1|}\right] \ge 1 - \delta$ for all $t \ge T$.

Proof. First, we invoke Lemma 7 (which requires $\mathcal{E}_3 \cap \mathcal{E}_{2,0}$) which controls the FWER. All that is left is to control the sample complexity.

Let $\mathcal{I} = \{i \in \mathcal{H}_1 : \rho_i \ge \delta\}$ be the same random set defined in the proof of Theorem 2. Note that $\mathcal{R}_t \subseteq \mathcal{R}_{t+1}$ for all t so define $T = \min\{t \in \mathbb{N} : \mathcal{I} \subseteq \mathcal{R}_{t+1}\}$ if such inclusion ever exists, otherwise

let $T = \infty$. Noting that $\mathcal{R}_t \subseteq \mathcal{S}_t$ we have

$$T = \sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{I} \not\subseteq \mathcal{R}_t \}$$

= $\sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{I} \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq \mathcal{S}_t \}$
= $\sum_{t=1}^{\infty} \mathbf{1} \{ \mathcal{I} \not\subseteq \mathcal{S}_t \} + \mathbf{1} \{ J_t \in \mathcal{H}_0, \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq \mathcal{S}_t \} + \mathbf{1} \{ J_t \in \mathcal{H}_1, \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq \mathcal{S}_t \}.$

First sum. Since we are in the TPR setting we have $\xi_t = 1$ so the first sum is precisely the quantity bounded in the proof of Theorem 2. Thus,

$$\sum_{t=1}^{\infty} \mathbf{1}\{\mathcal{I} \not\subseteq \mathcal{S}_t\} \le cn\Delta^{-2}\log(\log(\Delta^{-2})/\delta).$$

Second sum. Recall that $J_t = \arg \max_{i \in S_t \setminus \mathcal{R}_t} \widehat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta)$. Now, on the event $\{ \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq S_t \}$ there exists a $j \in \mathcal{I} \cap (S_t \setminus \mathcal{R}_t)$ such that

$$\begin{aligned} \widehat{\mu}_{j,T_j(t)} + \phi(T_j(t),\delta) &\geq \mu_j + \phi(T_j(t),\delta) - \phi(T_j(t),\rho_j) \\ &\geq \mu_j \\ &\geq \mu_0 + \Delta \end{aligned}$$

where the first inequality follows by the definition of \mathcal{I} . On the other hand, for any $i \in \mathcal{H}_0$ we have

$$\begin{aligned} \widehat{\mu}_{i,T_{i}(t)} + \phi(T_{i}(t),\delta) &\leq \mu_{i} + \phi(T_{i}(t),\delta) + \phi(T_{i}(t),\rho_{i}) \\ &\leq \mu_{i} + 2\phi(T_{i}(t),\delta\rho_{i}) \\ &\leq \mu_{0} + 2\phi(T_{i}(t),\delta\rho_{i}) \end{aligned}$$

so that $\hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \le \mu_0 + \Delta$ whenever $T_i(t) \ge \phi^{-1}(\frac{\Delta}{2}, \delta\rho_i)$. Thus, by the same sequence of steps used in Theorem 2 we have

$$\sum_{t=1}^{\infty} \mathbf{1}\{J_t \in \mathcal{H}_0, \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq \mathcal{S}_t\} \le \sum_{i \in \mathcal{H}_0} \phi^{-1}(\frac{\Delta}{2}, \delta\rho_i) \stackrel{\mathcal{E}_{4,0}}{\le} c|\mathcal{H}_0|\Delta^{-2}\log(\log(\Delta^{-2})/\delta).$$

Third sum. In bounding the analogous sum in the proof of Theorem 4 we used the fact that $\mathcal{H}_1 \subseteq \mathcal{S}_t$ to lowerbound $|\mathcal{S}_t|$ to obtain an upperbound on ξ_t . Now that we only have $\mathcal{I} \subseteq \mathcal{S}_t$ we observe that

$$|\mathcal{S}_t| \ge |\mathcal{I}| \ge |\mathcal{H}_1| - \sum_{i \in \mathcal{H}_1} \mathbf{1}\{\rho_i \le \delta\} \stackrel{\mathcal{E}_5}{\ge} |\mathcal{H}_1| (1 - \delta - \sqrt{\frac{2\delta \log(1/\delta)}{|\mathcal{H}_1|}} - \frac{\log(1/\delta)}{3|\mathcal{H}_1|}).$$

Using this approximation the same argument yields

$$\begin{split} &\sum_{t=1}^{\infty} \mathbf{1} \{ J_t \in \mathcal{H}_1, \mathcal{I} \not\subseteq \mathcal{R}_t, \mathcal{I} \subseteq \mathcal{S}_t \} \\ &\leq c' |\mathcal{H}_1| \Delta^{-2} \log(\max\{n - (1 - 3\delta - \sqrt{\frac{2\delta \log(1/\delta)}{|\mathcal{H}_1|}} - \frac{\log(1/\delta)}{3|\mathcal{H}_1|}) |\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta) \\ &\leq c'' |\mathcal{H}_1| \Delta^{-2} \log(\max\{n - (1 - 3\delta - \sqrt{2\delta \log(1/\delta)}/|\mathcal{H}_1|) |\mathcal{H}_1|, \log \log(n/\delta)\} \log(\Delta^{-2})/\delta). \end{split}$$

G Successive Elimination and Uniform Allocation Algorithms

The following gives a sample complexity bound for Successive Elimination and Uniform allocation strategies. Note that for these algorithms, up to n arms are pulled at each time t.

Algorithm 2 Uniform and Successive elimination algorithms for identifying arms with means above μ_0 .

Input: Threshold μ_0 , confidence δ , confidence interval $\phi(\cdot, \cdot)$ Initialize: Set $S_1 = \emptyset$ For t = 1, 2, ...if Successive Elimination Pull each and every arm in $[n] - S_t$. else if Uniform Pull each and every arm in [n]. Apply Benjamini-Hochberg [11] selection to obtain FDR-controlled set S_t : $s(k) = \{i \in [n] \setminus S_t : \hat{\mu}_{i,t} - \phi(t, \delta \frac{k}{n}) \ge \mu_0\}, \forall k \in [n]$ $S_{t+1} = S_t \cup s(\hat{k})$ where $\hat{k} = \max\{k \in [n] : |s(k)| \ge k\}$ (if $\nexists \hat{k}$ set $S_{t+1} = S_t$)

Theorem 6. For both the Successive Elimination and Uniform Allocation algorithms, for all t we have that $\mathbb{E}\left[\frac{|S_t \cap \mathcal{H}_0|}{|S_t|}\right] \leq \delta$. In addition, in the case of successive elimination, if $\eta = \delta + \sqrt{\frac{2\delta \log(1/\delta)}{|\mathcal{H}_1|}} + \frac{\log(1/\delta)}{3|\mathcal{H}_1|}$, then on the event \mathcal{E}_5 (which holds with probability greater than $1 - \delta$) then there exists a T such that

$$T \lesssim \min\left\{ n\Delta^{-2}\log(\frac{n}{(1-\eta)|\mathcal{H}_1|}\log(\Delta^{-2})/\delta), \\ (n-(1-\eta)|\mathcal{H}_1|)\Delta^{-2}\log(\frac{n}{(1-\eta)|\mathcal{H}_1|}\log(\Delta^{-2})/\delta) + \sum_{i\in\mathcal{H}_1}\Delta_i^{-2}\log(n\log(\Delta_i^{-2})/\delta) \right\}$$

and in the case of Uniform allocation,

$$T \lesssim n\Delta^{-2} \log(\frac{n}{(1-\eta)|\mathcal{H}_1|} \log(\Delta^{-2})/\delta)$$

and $\mathbb{E}\left[\frac{|\mathcal{S}_t \cap \mathcal{H}_1|}{|\mathcal{S}_t|}\right] > 1 - \delta$ for all $t \ge T$.

Proof. Define the random set $\mathcal{I} = \{i \in \mathcal{H}_1 : \rho_i \geq \delta\}$. Note, this is equivalent to $\mathcal{I} = \{i \in \mathcal{H}_1 : \hat{\mu}_{i,T_i(t)} + \phi(T_i(t), \delta) \geq \mu_i \ \forall t \in \mathbb{N}\}$ since if $\rho_i \geq \delta$ then $\phi(T_i(t), \delta) > \phi(T_i(t), \rho_i)$ and

$$\widehat{\mu}_{i,T_i(t)} + \phi(T_i(t),\delta) \ge \mu_i + \phi(T_i(t),\delta) - \phi(T_i(t),\rho_i) \ge \mu_i.$$

Our aim is to show that this "well-behaved" set of indices \mathcal{I} will be added to \mathcal{S}_t in the claimed amount of time. This is sufficient for the TPR result because $\mathbb{E}[\mathcal{I}] = \sum_{i \in \mathcal{H}_1} \mathbb{E}[\mathbf{1}\{\rho_i > \delta\}] = \sum_{i \in \mathcal{H}_1} \mathbb{P}(\rho_i > \delta) \ge (1 - \delta)|\mathcal{H}_1|$ since each ρ_i is a sub-uniformly distributed random variable. To be clear, we are not claiming which *particular* indices of \mathcal{H}_1 will be added to \mathcal{S}_t (indeed, \mathcal{I} is a random set), just that their number exceeds $(1 - \delta)|\mathcal{H}_1|$ in expectation.

First we consider the case of Successive Elimination. Let $T_i = \min\{t \in \mathbb{N} : i \notin S_t\}$ be the random number of times arm *i* is pulled until the last time when it is added to S_t (may possibly be infinite). At round *t* all arms in $[n] - S_t$ are pulled and once an arm is added to S_t it will never be pulled again. Note that $i \in s(k)$ implies $\hat{\mu}_{i,t} - \phi(t, \delta \frac{k}{n}) \ge \mu_0$. Since for all $i \in \mathcal{I}$

$$\begin{aligned} \widehat{\mu}_{i,t} - \phi(t, \delta \frac{k}{n}) &\geq \mu_i - \phi(t, \delta \frac{k}{n}) - \phi(t, \delta) \\ &\geq \mu_i - 2\phi(t, \delta \frac{k}{n}) \\ &= \mu_0 + \Delta_i - 2\phi(t, \delta \frac{k}{n}) \end{aligned}$$

we have that $i \in s(k)$ whenever $t \ge \phi^{-1}(\frac{\Delta_i}{2}, \delta_n^k)$. In particular, because $i \in s(1)$ implies $i \in S_t$ we have that $T_i \le \phi^{-1}(\frac{\Delta_i}{2}, \delta_n^1)$ for all $i \in \mathcal{I}$. But if $\Delta = \min_{i \in \mathcal{H}_1} \Delta_i$ then we also have necessarily that $t \le \phi^{-1}(\frac{\Delta}{2}, \delta_n^{|\mathcal{I}|})$ since at this time, all arms in \mathcal{I} will be in $s(|\mathcal{I}|)$, which means $\hat{k} \ge |\mathcal{I}|$ and so $\mathcal{I} \subseteq S_t$. This, of course, implies $T_i \le \phi^{-1}(\frac{\Delta}{2}, \delta_n^{|\mathcal{I}|})$ for all $i \in [n] \setminus \mathcal{I}$. Putting these pieces together,

we have that

$$\sum_{i=1}^{n} T_{i} \leq \min\left\{n\phi^{-1}(\frac{\Delta}{2},\delta\frac{|\mathcal{I}|}{n}), (n-|\mathcal{I}|)\phi^{-1}(\frac{\Delta}{2},\delta\frac{|\mathcal{I}|}{n}) + \sum_{i\in\mathcal{I}}\phi^{-1}(\frac{\Delta_{i}}{2},\delta\frac{1}{n})\right\}$$
$$\leq \min\left\{cn\Delta^{-2}\log(\frac{n}{|\mathcal{I}|}\log(\Delta^{-2})/\delta), \\ c(n-|\mathcal{I}|)\Delta^{-2}\log(\frac{n}{|\mathcal{I}|}\log(\Delta^{-2})/\delta) + \sum_{i\in\mathcal{I}}c\Delta_{i}^{-2}\log(n\log(\Delta_{i}^{-2})/\delta)\right\}.$$

Now on event \mathcal{E}_5 (which holds with probability at least $1 - \delta$) we have

$$|\mathcal{I}| \ge (1 - \delta - \sqrt{\frac{2\delta \log(1/\delta)}{|\mathcal{H}_1|}} - \frac{\log(1/\delta)}{3|\mathcal{H}_1|})|\mathcal{H}_1|$$

which implies that for $\eta = \delta + \sqrt{\frac{2\delta \log(1/\delta)}{|\mathcal{H}_1|}} + \frac{\log(1/\delta)}{3|\mathcal{H}_1|}$ we have

$$\sum_{i=1}^{n} T_{i} \leq \min \Big\{ cn\Delta^{-2} \log(\frac{n}{(1-\eta)|\mathcal{H}_{1}|} \log(\Delta^{-2})/\delta), \\ c(n-(1-\eta)|\mathcal{H}_{1}|)\Delta^{-2} \log(\frac{n}{(1-\eta)|\mathcal{H}_{1}|} \log(\Delta^{-2})/\delta) + \sum_{i \in \mathcal{H}_{1}} c\Delta_{i}^{-2} \log(n\log(\Delta_{i}^{-2})/\delta) \Big\}.$$

Now, in the case of Uniform allocation, we never stop pulling arms once they enter S_t . Hence, we need to consider the number of samples needed before all the arms in $\mathcal{I} \subset S_t$. By the same reasoning as above, this is bounded by $cn\Delta^{-2}\log(\frac{n}{(1-\eta)|\mathcal{H}_1|}\log(\Delta^{-2})/\delta)$.

The following two theorems provide lower bounds for Uniform allocation.

Theorem 7 (FDR, TPR). Fix $\delta < 1/40$, $\Delta > 0$. and k < n/2. For any $\mathcal{H}_1 \subseteq [n]$ such that $|\mathcal{H}_1| = k$ define an instance $(\{\nu_i\}_{i=1}^n, \mu_0)$ with $\mu_0 = 0$, $\nu_i = \mathcal{N}(\Delta, 1)$ for $i \in \mathcal{H}_1$, $\nu_i = \mathcal{N}(0, 1)$ for $i \in \mathcal{H}_0$. Any algorithm that samples each arm an equal number of times before outputting a set $S \subseteq [n]$ after τ total samples, and is FDR- δ and TPR- δ , τ on $(\{\nu_i\}_{i=1}^n, \mu_0)$ for all $\mathcal{H}_1 \subseteq [n]$ such that $|\mathcal{H}_1| = k$ simultaneously, must satisfy $t \gtrsim n\Delta^{-2} \log (n/k)$.

Proof. The proof is based on the construction of [29] which states that for any n > 2k, there exists a collection $\mathcal{M}_{n,k}$ of subsets of [n] where a) each $\pi \in \mathcal{M}_{n,k}$ has weight $|\pi| = k$, b) $2k \ge |\pi \bigtriangleup \pi'| > k$ for all $\pi \ne \pi' \in \mathcal{M}_{n,k}$, and c) $|\mathcal{M}_{n,k}| \ge (\frac{n}{6k})^{k/4}$. Each π gives rise to an instance of the problem $(\{\nu_i\}_{i=1}^n, \mu_0)$ where $\mu_0 = 0$, $\mathcal{H}_1 = \pi$ and so $\nu_i = \mathcal{N}(\Delta, 1)$ if $i \in \pi$, otherwise and $\nu_i = \mathcal{N}(\mu_0, 1)$ otherwise. In particular, for each instance $|\mathcal{H}_1| = k$.

Note that if (i),
$$\frac{|\mathcal{H}_{1} \cap \mathcal{S}_{t}^{c}|}{|\mathcal{H}_{1}|} = 1 - \frac{|\mathcal{H}_{1} \cap \mathcal{S}_{t}|}{|\mathcal{H}_{1}|} \leq \eta \text{ and } (ii) \frac{|\mathcal{H}_{0} \cap \mathcal{S}_{t}|}{|\mathcal{H}_{0} \cap \mathcal{S}_{t}| + |\mathcal{H}_{1} \cap \mathcal{S}_{t}|} = \frac{|\mathcal{H}_{0} \cap \mathcal{S}_{t}|}{|\mathcal{S}_{t}|} \leq \eta \text{ then}$$
$$|\mathcal{H}_{1} \triangle \mathcal{S}_{t}| = |\mathcal{H}_{1} \cap \mathcal{S}_{t}^{c}| + |\mathcal{H}_{1}^{c} \cap \mathcal{S}_{t}|$$
$$= |\mathcal{H}_{1} \cap \mathcal{S}_{t}^{c}| + |\mathcal{H}_{0} \cap \mathcal{S}_{t}|$$
$$\stackrel{(ii)}{\leq} \eta |\mathcal{H}_{1}| + |\mathcal{H}_{0} \cap \mathcal{S}_{t}|$$
$$\stackrel{(iii)}{\leq} \eta |\mathcal{H}_{1}| + \frac{\eta}{1-\eta} |\mathcal{H}_{1} \cap \mathcal{S}_{t}|$$
$$\leq \frac{2\eta}{1-\eta} |\mathcal{H}_{1}|.$$

Thus, if (i) and (ii) hold and $\eta < 1/5$ then $|\mathcal{H}_1 \triangle \mathcal{S}_t| \le k/2$ and so $\mathcal{H}_1 \triangle \mathcal{S}_t$ can therefore be used as an estimator for any $\mathcal{H}_1 = \pi \in \mathcal{M}_{n,k}$ since $\min_{\pi,\pi' \in \mathcal{M}_{n,k}} |\pi \triangle \pi'| > k$.

By assumption, $\mathbb{E}[\frac{|\mathcal{H}_0 \cap S_t|}{|S_t|}] \leq \delta$ so by Markov's inequality we have $\mathbb{P}(\frac{|\mathcal{H}_0 \cap S_t|}{|S_t|} \geq 8\delta) \leq 1/8$. Likewise, by assumption $\mathbb{E}[\frac{|\mathcal{H}_1 \cap S_t^c|}{|\mathcal{H}_1|}] = 1 - \mathbb{E}[\frac{|\mathcal{H}_1 \cap S_t|}{|\mathcal{H}_1|}] \leq \delta$, so again $\mathbb{P}(\frac{|\mathcal{H}_1 \cap S_t^c|}{|\mathcal{H}_1|} \geq 8\delta) \leq 1/8$. Thus, with probability at least 3/4, $\frac{|\mathcal{H}_1 \cap S_t^c|}{|\mathcal{H}_1|} < 8\delta$ and $\frac{|\mathcal{H}_0 \cap S_t|}{|S_t|} \leq 8\delta$. To apply the above argument, we just need $8\delta < 1/5$ which holds when $\delta < 1/40$, then S_t could predict the correct $\pi \in \mathcal{M}_{n,k}$ with probability at least 1/4. We will now use an information theoretic inequality to lower bound the t that would make such an estimator possible. Let P_{π} be the probability law of sampling τ samples from each arm under π . Then $KL(P_{\pi}, P_{\pi'}) \leq \Delta^2 \tau |\pi \Delta \pi'|/2 \leq \Delta^2 \tau k$. Directly applying Theorem 2.5 of [30] to our $|\mathcal{M}_{n,k}| \geq (\frac{n}{6k})^{k/4}$ hypotheses, we have that any estimator has a probability of misidentification of at

$$\frac{1}{2} \left(1 - \frac{2\Delta^2 \tau k}{\log |\mathcal{M}_{n,k}|} - \sqrt{\frac{2\Delta^2 \tau k}{\log^2 |\mathcal{M}_{n,k}|}} \right) \ge \frac{1}{2} \left(1 - \frac{8\Delta^2 \tau}{\log(n/6k)} - \sqrt{\frac{32\Delta^2 \tau}{k \log^2(n/k)}} \right)$$

is at least 1/4 unless $\tau \gtrsim \Delta^{-2} \log(n/k)$.

which is at least 1/4 unless $\tau \gtrsim \Delta^{-2} \log(n/k)$.

Theorem 8 (FWER, FWPD). Fix $\delta < 3/8$ and $\Delta > 0$. For any $\mathcal{H}_1 \subseteq [n]$ consider an instance $(\{\nu_i\}_{i=1}^n, \mu_0)$ such that $\mu_0 = -\Delta/2$ and $\nu_i = \mathcal{N}(\mu_i, 1)$ where $\mu_i = \Delta/2$ if $i \in \mathcal{H}_1$ and $\mu_i = -\Delta/2$ if $i \in \mathcal{H}_0$. Any algorithm that samples each arm an equal number of times before outputting a set $S \subseteq [n]$ after τ total samples, and is FWER- δ and FWPD- δ , τ on $(\{\nu_i\}_{i=1}^n, \mu_0)$ for all $\mathcal{H}_1 \subseteq [n]$ simultaneously, must satisfy $\tau \gtrsim n\Delta^{-2}\log(n)$.

Proof. Fix $t \in \mathbb{N}$. Because the empirical mean is a sufficient statistic for each arm and they are independent, the joint probability distribution of the from the n arms is given by $P_0 := \prod_{i=1}^n \mathcal{N}(\mu_i, 1)$ and the distribution after t-pulls on each arm is $P_0^t := \prod_{i=1}^n \mathcal{N}(\mu_i, 1/t)$

For any $j \in [n]$ define P_j as the joint distribution if the *j*th arm's identity was flipped: if $j \in \mathcal{H}_0$ replace its mean with 1, if $j \in \mathcal{H}_1$ replace its means with \mathcal{H}_0 . Note that $P_j := P_0 \frac{\mathcal{N}(-\mu_i, 1)}{\mathcal{N}(\mu_i, 1)}$ and that $KL(P_j|P_0) = KL(\mathcal{N}(-\mu_i, 1)|\mathcal{N}(\mu_i, 1)) = 2\mu_i^2 = \Delta^2/2$ and $KL(P_j^t|P_0^t) = t\Delta^2/2$.

Now, because the algorithm was assumed FWER- δ and FWPD- δ , t on all instances indexed by $\mathcal{H}_1 \subseteq [n]$, it will be able to distinguish between $\{P_k\}_{k=0}^n$ using just t per arm with probability at least $1 - 2\delta \ge 1/4$. The multiple hypothesis testing lower bound of Tsybakov [30, Theorem 2.5], implies that the probability of misclassification of any estimator will be at least

$$\frac{1}{2}(1-\frac{t\Delta^2}{\log(n)}-\sqrt{\frac{t\Delta^2}{\log^2(n)}})\geq \frac{1}{4}$$

unless $t \gtrsim \Delta^{-2} \log n$.

Technical Lemmas Η

Lemma 8. Fix $a \in \mathbb{R}^n_+$ and for i = 1, ..., n let Z_i be independent random variables satisfying $\mathbb{P}(Z_i \geq t) \leq \exp(-t/a_i)$. Then

$$\mathbb{P}\left(\sum_{i=1}^{n} (Z_i - a_i) \ge t\right) \le \exp\left(-\min\{\frac{t}{4||a||_{\infty}}, \frac{t^2}{4||a||_2^2}\}\right)$$

and moreover, with probability at least $1 - \delta$, $\sum_{i=1}^{n} Z_i \leq 5 \log(1/\delta) \sum_{i=1}^{n} a_i$.

Proof. For i = 1, ..., n we have that $Z_i = a_i \log(1/\rho_i)$ are independent random variables satisfying $\mathbb{P}(Z_i \ge t) \le \exp(-t/a_i)$, because the ρ_i are independent *sub-uniformly* distributed random variables. It is straightforward to verify that $\log(\mathbb{E}[\exp(\lambda(Z_i - a_i))]) \leq -a_i\lambda - \log(1 - \lambda a_i) \leq \frac{a_i^2\lambda^2}{2(1 - \lambda a_i)}$ for $\lambda \leq 1/||a||_{\infty}$. Using the standard Chernoff-bound technique, we have for $\lambda = \min\{\frac{t}{2||a||_{\infty}^2}, \frac{1}{2||a||_{\infty}}\}$ that

$$\mathbb{P}\left(\sum_{i=1}^{n} (Z_i - a_i) \ge t\right) \le \exp\left(-\lambda t + \sum_{i=1}^{n} \frac{a_i^2 \lambda^2}{2(1 - \lambda a_i)}\right)$$
$$\le \exp\left(-\lambda t + \lambda^2 ||a||_2^2\right)$$
$$\le \exp\left(-\min\left\{\frac{t}{4||a||_{\infty}}, \frac{t^2}{4||a||_2^2}\right\}\right)$$
$$\le \exp\left(-\frac{1}{4}\min\left\{\frac{t}{||a||_1}, \frac{t^2}{||a||_1^2}\right\}\right)$$

where the last inequality holds by $||a||_{\infty} \leq ||a||_{2} \leq ||a||_{1} = \sum_{i=1}^{n} a_{i}$. The result follows from setting the right hand side equal to δ and solving for t.

Lemma 9. Fix $\delta \in (0, 1/2]$. Let $X_i \in [0, 1]$ for i = 1, ..., m be independent random variables, each satisfying $\mathbb{P}(X_i \leq s) \leq s$. Define $A(s) = \sum_{i=1}^m \mathbf{1}\{X_i \leq s\}$, then for any $c \in (0, 1)$

$$\mathbb{P}\left(\exists s \in (0,1] : A(s) > sm + (1+4s)\sqrt{2\max\{2s,c\}m\log(\frac{\log_2(2/c)}{\delta})} + \frac{1+4s}{3}\log(\frac{\log_2(2/c)}{\delta})\right) \le \delta.$$

Moreover,

$$\mathbb{P}\left(\exists s \in (0,1] : A(s) > sm + (1+2s)\sqrt{4sm\log(\frac{2\log_2(2/s)^2}{\delta})} + \frac{1+2s}{3}\log(\frac{2\log_2(2/s)^2}{\delta})\right) \le \delta.$$

Also, recall that by the Dvoretzky-Kiefer-Wolfowitz inequality [31] we have $\mathbb{P}\left(\exists s \in (0,1] : A(s) > sm + \sqrt{m\log(1/\delta)/2}\right) \leq \delta.$

Proof. First note that $M(s) = \frac{A(s) - \mathbb{E}[A(s)]}{1-s}$ is a martingale with respect to the filtration $\mathcal{F}_t = \{A(s) : s \leq t\}$. Thus, for $\lambda > 0$ we have that $\exp(\lambda M(s))$ is a non-negative sub-martingale and we can apply Doob's maximal inequality to obtain

$$\mathbb{P}\left(\sup_{s \le t} M(s) \ge \epsilon/(1-t)\right) = \mathbb{P}\left(\sup_{s \le t} \exp(\lambda M(s)) \ge \exp(\lambda \epsilon/(1-t))\right)$$
$$\le \exp(-\lambda \epsilon/(1-t))\mathbb{E}\left[\exp(\lambda M(t))\right]$$
$$= \exp(-\frac{\lambda}{1-t}\epsilon)\mathbb{E}\left[\exp(\frac{\lambda}{1-t}(A(t) - \mathbb{E}[A(s)]))\right]$$

Observe that for all t < 1 we have

$$\min_{\lambda} \exp(-\frac{\lambda}{1-t}\epsilon) \mathbb{E}\left[\exp(\frac{\lambda}{1-t}(A(t) - \mathbb{E}[A(s)]))\right] = \min_{\lambda} \exp(-\lambda\epsilon) \mathbb{E}\left[\exp(\lambda(A(t) - \mathbb{E}[A(s)]))\right]$$

Noting that A(t) is a sum of independent random variables with each in [0, 1] and expectation less than t so that $\mathbb{E}[A(t)] \leq mt$, we apply Bernstein's inequality to obtain $\log \mathbb{E} \left[\exp(\lambda(A(t) - \mathbb{E}[A(t)])) \right] \leq \frac{mt\lambda^2}{2(1-\lambda/3)}$ for $\lambda \in (0, 3)$. Optimizing over $\lambda \in (0, 3)$ we have

$$\mathbb{P}\left(\exists s \le t : A(s) > \mathbb{E}[A(s)] + \frac{1-s}{1-t}\sqrt{2mt\log(1/\delta)} + \frac{1-s}{1-t}\log(1/\delta)/3\right) \le \delta$$

For $k \in \mathbb{N}$ define $T_k = \{s \in [0,1] : 2^{-k-1} < s \le 2^{-k}\}$. Note that $2^{-\lfloor \log_2(2/c) \rfloor} \le c$. So for any $k = 1, 2, \ldots, \lfloor \log_2(2/c) \rfloor$, with probability at least $1 - \frac{\delta}{\lfloor \log_2(2/c) \rfloor}$ we have that for any $s \in T_k$

$$A(s) \leq \mathbb{E}[A(s)] + \frac{1-s}{1-2^{-k}} \sqrt{2\max\{c, 2^{-k}\}m\log(\log_2(2/c)/\delta) + \frac{1-s}{1-2^{-k}}\log(\log_2(2/c)/\delta)/3}.$$

1Note that cases $s \in T_k$ for $k > \lfloor \log_2(2/c) \rfloor$ are handled by $k = \lfloor \log_2(2/c) \rfloor$. For any $k \ge 1$ and $s \in T_k$ we have $s \ge 2^{-k-1}$ so that $1 + 4s \ge \frac{1}{1 - \min\{2s, 2^{-1}\}} \ge \frac{1-s}{1-2^{-k}}$ and $2s \ge 2^{-k}$. Thus,

$$\mathbb{P}\left(\exists s \in T_k : A(s) > \mathbb{E}[A(s)] + (1+4s)\sqrt{2\max\{c, 2s\}m\log(\log_2(2/c)/\delta)} + (1+4s)\log(\log_2(2/c)/\delta)/3\right)$$

$$\leq \mathbb{P}\left(\exists s \in T_k : A(s) > \mathbb{E}[A(s)] + \frac{1-s}{1-2^{-k}}\sqrt{2\max\{c, 2^{-k}\}m\log(\log_2(2/c)/\delta)} + \frac{1-s}{1-2^{-k}}\log(\log_2(2/c)/\delta)/3\right)$$

$$\leq \delta \frac{1}{\log_2(2/c)} \leq \delta \frac{1}{\lfloor \log_2(2/c) \rfloor}.$$
Using haunding curve $h = 2$ and $\lfloor \log_2(2/c) \rfloor$ handles $\downarrow \otimes T_k = (0, 1/4)$. To handle $s \in (1/4, 1]$

Union bounding over $k = 2, ..., \lfloor \log_2(2/c) \rfloor$ handles $\bigcup_{k=2}^{\infty} T_k = (0, 1/4]$. To handle $s \in (1/4, 1]$, we note that

$$\mathbb{P}\left(\exists s \in (1/4, 1] : A(s) > \mathbb{E}[A(s)] + (1 + 4s)\sqrt{2\max\{c, 2s\}m\log(\log_2(2/c)/\delta)} + (1 + 4s)\log(\log_2(2/c)/\delta)/3\right) \\
\leq \mathbb{P}\left(\exists s \in T_k : A(s) > \mathbb{E}[A(s)] + \sqrt{m\log(\log_2(2/c)/\delta)}\right) \\
\leq \left(\delta \frac{1}{\log_2(2/c)}\right)^2 \leq \delta \frac{1}{\lfloor \log_2(2/c) \rfloor}$$

where the second to last inequality holds by the DKW inequality [31].

On the other hand, for any
$$k \ge 1$$
 and $s \in T_k$ we have $s \ge 2^{-k-1}$ so that $1 + 2s \ge \frac{1-s}{1-\min\{2s,2^{-1}\}} \ge \frac{1-s}{1-\min\{2s,2^{-1}\}} \ge \frac{1-s}{1-2^{-k}}$ and $2^{-k-1} \le s \le 2^{-k}$.

$$\mathbb{P}\left(\exists s \in T_k : A(s) > \mathbb{E}[A(s)] + (1+2s)\sqrt{4sm\log(2\log_2(\frac{2}{s})^2/\delta)} + (1+2s)\log(2\log_2(\frac{2}{s})^2/\delta)/3\right)$$

$$\le \mathbb{P}\left(\exists s \in T_k : A(s) > \mathbb{E}[A(s)] + \frac{1-s}{1-2^{-k}}\sqrt{2 \cdot 2^{-k}m\log(2(k+1)^2/\delta)} + \frac{1-s}{1-2^{-k}}\log(2(k+1)^2/\delta)/3\right)$$

$$\le \delta \frac{1}{2(k+1)^2}.$$

Union bounding over all $k \ge 0$ and noting that $\sum_{k=0}^{\infty} \frac{1}{2(k+1)^2} \le 1$ completes the proof since $\bigcup_{k\ge 0} T_k = (0,1]$

uniform distribution on [0, 1] is given by,

$$\frac{d\mathbb{P}(p_{\ell}^{0} \le x)}{dx} = \ell \binom{m}{\ell} (1-x)^{m-\ell} x^{\ell-1}.$$

First we consider the case when $\ell > 1$.

$$\begin{split} \sum_{k=1}^{\infty} \frac{1}{k} \mathbb{P}(\frac{(k-1)\alpha}{n} \le p_{\ell}^{0} \le \frac{k\alpha}{n}) \le \mathbb{P}(p_{\ell}^{0} \le \frac{\alpha}{n}) + \int_{x=\frac{\alpha}{n}}^{1} \frac{\alpha}{nx} d\mathbb{P}(p_{\ell}^{0} \le x) \\ &= \mathbb{P}(p_{\ell}^{0} \le \frac{\alpha}{n}) + \frac{\alpha}{n} \ell\binom{m}{\ell} \int_{x=\frac{\alpha}{n}}^{1} (1-x)^{m-\ell} x^{\ell-2} dx \\ &\le \frac{\alpha}{n} + \frac{\alpha}{n} \ell\binom{m}{\ell} (\ell-1)^{-1} \binom{m-1}{\ell-1}^{-1} \int_{\alpha/n}^{1} (\ell-1) \binom{m-1}{\ell-1} (1-x)^{m-\ell} x^{\ell-2} dx \\ &\le \frac{\alpha}{n} + \frac{\alpha}{n} \frac{m!}{(\ell-1)!(m-\ell)!} \frac{(\ell-2)!(m-\ell)!}{(m-1)!} \\ &\le \frac{\alpha}{n} + \frac{\alpha}{n} \frac{m}{\ell-1} \end{split}$$

When $\ell = 1$, the argument is slightly different. Firstly note that

$$\mathbb{P}(\frac{(k-1)\alpha}{n} \le p_{\ell}^{0} \le \frac{k\alpha}{n}) \le$$
$$\sum_{k=1}^{\infty} \frac{1}{k} \mathbb{P}(\frac{(k-1)\alpha}{n} \le p_{\ell}^{0} \le \frac{k\alpha}{n}) \le$$