
Finding All ϵ -Good Arms in Stochastic Bandits

Blake Mason
University of Wisconsin
Madison, WI 53706
bmason3@wisc.edu

Lalit Jain
University of Washington
Seattle, WA 98115
lalitj@uw.edu

Ardhendu Tripathy
University of Wisconsin
Madison, WI 53706
astripathy@wisc.edu

Robert Nowak
University of Wisconsin
Madison, WI 53706
rdnowak@wisc.edu

Abstract

The pure-exploration problem in stochastic multi-armed bandits aims to find one or more arms with the largest (or near largest) means. Examples include finding an ϵ -good arm, best-arm identification, top- k arm identification, and finding all arms with means above a specified threshold. However, the problem of finding *all* ϵ -good arms has been overlooked in past work, although arguably this may be the most natural objective in many applications. For example, a virologist may conduct preliminary laboratory experiments on a large candidate set of treatments and move all ϵ -good treatments into more expensive clinical trials. Since the ultimate clinical efficacy is uncertain, it is important to identify all ϵ -good candidates. Mathematically, the all- ϵ -good arm identification problem presents significant new challenges and surprises that do not arise in the pure-exploration objectives studied in the past. We introduce two algorithms to overcome these and demonstrate their great empirical performance on a large-scale crowd-sourced dataset of 2.2M ratings collected by the New Yorker Caption Contest as well as a dataset testing hundreds of possible cancer drugs.

1 Introduction

We propose a new multi-armed bandit problem where the objective is to return *all* arms that are ϵ -good relative to the best-arm. Concretely, if the arms have means μ_1, \dots, μ_n , with $\mu_1 = \max_{1 \leq i \leq n} \mu_i$, then the goal is to return the set $\{i : \mu_i \geq \mu_1 - \epsilon\}$ in the **additive** case, and $\{i : \mu_i \geq (1 - \epsilon)\mu_1\}$ in the **multiplicative** case. The ALL- ϵ problem is a novel setting in the bandits literature, adjacent to two other methods for finding many good arms: TOP- k where the goal is to return the arms with the k highest means, and threshold bandits where the goal is to identify all arms above a fixed threshold. Building on a metaphor given by [1], if TOP- k is a “contest” and thresholding bandits is an “exam”, ALL- ϵ organically decides which arms are “above the bar” relative to the highest score. We argue that the ALL- ϵ problem formulation is more appropriate in many applications, and we show that it presents some unique challenges that make its solution distinct from TOP- k and threshold bandits.

A Natural and Robust Objective. A motivating example is drug discovery, where pharmacologists want to identify a set of highly-potent drug candidates from potentially millions of compounds using various *in vitro* and *in silico* assays, and only the selected undergo more expansive testing [2]. Since performing the assays can be costly, one would like to use an adaptive, sequential experiment design that requires fewer experiments than a fixed experiment design. In sequential experiment design, it is important to fix the objective at the beginning as that choice affects the experimentation process. Both the objectives of finding the top- k performing drugs, or all drugs above a threshold can result

in failure. In TOP- k , choosing k too small may miss potent compounds, and choosing k too large may yield many ineffective compounds and require an excessively large number of experiments. Setting a threshold suffers from the same issues - with the additional concern that if it is set too high, potentially no drug discoveries are made. In contrast, the ALL- ϵ objective of finding all arms whose potency is within 20% of the best avoids these concerns by giving a robust and natural guarantee: *no* significantly suboptimal arms will be returned and but *every* near-optimal arm will be discovered.

We emphasize that unlike TOP- k or thresholding which require some prior knowledge about the distribution of arms to guarantee a good set of returned arms, choosing the arms relative to the best is a natural, distribution-free metric for finding good arms. As an example, we consider the New Yorker Cartoon Caption Contest (NYCCC). Each week, contestants submit thousands of supposedly funny captions for a cartoon (see Appendix A), which are rated from 1 (unfunny) to 3 (funny) through a crowdsourcing process. The New Yorker editors select final winners from a set with the highest average crowd-ratings (typically over 1 million ratings per contest). The number of truly funny captions varies from week to week, and this makes setting a choice of k or fixed threshold difficult. In Figure 1, we plot the distribution of ratings from 3 different contests. Horizontal lines depict a reasonable threshold of $0.8\mu_1$ in each and vertical lines show the number of arms that exceed this threshold. Both of these quantities vary over weeks and these differences can be stark. In contest 627, only $k = 27$ arms are within 20% of μ_1 , but $k = 748$ are in contest 651. Additionally, a fixed threshold of $\tau = 1.5$, admits captions within 30% of the best in contest 627, but only those within 15% of the best in contest 651. These examples show that it would be imprudent, and indeed, incorrect to choose a value of k or a threshold based on past contests— the far more principled decision is to optimize for the objective of finding the captions that are within a percentage of the best every week.




Figure 1: Mean ratings from contests 627, 651, 690

Though the ALL- ϵ objective is natural and easy to state, it has not been studied in the literature. As we will show, admitting arms relative to the best makes the ALL- ϵ problem inherently more challenging than either TOP- k or thresholding. In particular, it is not easily possible to adapt TOP- k or thresholding algorithms to achieve the instance dependent lower bound for ALL- ϵ . In this work, we provide a careful investigation of the ALL- ϵ problem including theoretical and empirical guarantees.

1.1 Problem Statement and Notation

Fix $\epsilon > 0$ and a failure probability $\delta > 0$. Let $\nu := \{\rho_1, \dots, \rho_n\}$ be an instance of n distributions (or arms) with 1-sub-Gaussian distributions having *unknown* means $\mu_1 \geq \dots \geq \mu_n$. We now formally define our notions of **additive** and **multiplicative** ϵ -good arms.

Definition 1 (**additive** ϵ -good). For a given $\epsilon > 0$, arm i is *additive ϵ -good* if $\mu_i \geq \mu_1 - \epsilon$.

Definition 2 (**multiplicative** ϵ -good). For a given $\epsilon > 0$, arm i is *multiplicative ϵ -good* if $\mu_i \geq (1 - \epsilon)\mu_1$.

Additionally, we define the sets

$$G_\epsilon(\nu) := \{i : \mu_i \geq \mu_1 - \epsilon\} \text{ and } M_\epsilon(\nu) := \{i : \mu_i \geq (1 - \epsilon)\mu_1\} \quad (1)$$

to be the sets of additive and multiplicative ϵ -good arms respectively. Where clear, we take $G_\epsilon = G_\epsilon(\nu)$ and $M_\epsilon = M_\epsilon(\nu)$. Consider an algorithm that at each time s selects an arm $I_s \in [n]$ based on the history $\mathcal{F}_{s-1} = \sigma(I_1, X_1, \dots, I_{s-1}, X_{s-1})$, and observes a reward $X_s \stackrel{\text{iid}}{\sim} \rho_{I_s}$. The objective of the algorithm is to return G_ϵ or M_ϵ using as few total samples as possible.

Definition 3. (ALL- ϵ problem). An algorithm for the ALL- ϵ problem is δ -PAC if (a) the algorithm has a finite stopping time τ with respect to \mathcal{F}_t , (b) at time τ it recommends a set \hat{G} such that with probability at least $1 - \delta$, $\hat{G} = G_\epsilon$ in the **additive** case, or $\hat{G} = M_\epsilon$ in the **multiplicative** case.

Notation: For any arm $i \in [n]$, let $\hat{\mu}_i(t)$ denote the empirical mean after t pulls. For all $i \in [n]$, define the suboptimality gap $\Delta_i := \mu_1 - \mu_i$. Without loss of generality, we denote $k = |G_\epsilon|$ (resp. $k = |M_\epsilon|$). Throughout, we will keep track of the quantity $\alpha_\epsilon := \min_{i \in G_\epsilon} \mu_i - (\mu_1 - \epsilon)$ which is the distance from the smallest additive ϵ -good arm, denoted μ_k , to the threshold $\mu_1 - \epsilon$. Additionally, if

G_ϵ^c is non-empty, we consider $\beta_\epsilon = \min_{i \in G_\epsilon^c} (\mu_1 - \epsilon) - \mu_i$, the distance of the largest arm that is not additive ϵ -good, denoted μ_{k+1} , to the threshold. Equivalently, in the case of returning multiplicative ϵ arms, we define $\tilde{\alpha}_\epsilon := \min_{i \in M_\epsilon} \mu_i - (1 - \epsilon)\mu_1$, $\tilde{\beta}_\epsilon := \min_{i \in M_\epsilon^c} (1 - \epsilon)\mu_1 - \mu_i$, μ_k , and μ_{k+1} to be the smallest differences of arms in M_ϵ and M_ϵ^c to $(1 - \epsilon)\mu_1$ respectively. For our sample complexity results, we also consider a relaxed version of the ALL- ϵ problem, where for a user-given slack $\gamma \geq 0$, we allow our algorithm to return \hat{G} that satisfies $G_\epsilon \subset \hat{G} \subset G_{\epsilon+\gamma}$ in the **additive** case, or $M_\epsilon \subset \hat{G} \subset M_{\epsilon+\gamma}$ in the **multiplicative** case. As we will see, this prevents large or potentially unbounded sample complexities when arms' means are very close to or equal $\mu_1 - \epsilon$.

1.2 Contributions and Summary of Main Results

In this paper we propose the new problem of finding *all* ϵ -good arms and give a precise characterization of its complexity. Our contribution is threefold:

- Information-theoretic lower bounds for the ALL- ϵ problem.
- A novel algorithm, (ST)², that is nearly optimal, is easy to implement, and has excellent empirical performance on real-world data.
- An instance optimal algorithm, FAREAST.

We now summarize our results in the **additive** setting (the **multiplicative** setting is analogous).

Lower Bound and Algorithms. As a preview of our results, we highlight the impact of three key quantities that affect the sample complexity: the user provided ϵ and the instance dependent quantities α_ϵ and β_ϵ , (see Figure 2). In this case, Theorem 2.1 implies that any δ -PAC algorithm requires an expected number of samples exceeding

$$\sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right\} \log \left(\frac{1}{\delta} \right). \quad (2)$$

We provide two algorithms, (ST)² and FAREAST for the ALL- ϵ problem. Our starting point, (ST)² is a novel combination of UCB [3] and LUCB [4] and is easier to implement and has good empirical performance. (ST)² is nearly optimal, however in some instances does not achieve the lower bound. To overcome this gap, we provide an instance optimal algorithm FAREAST which achieves the lower bound, however suffers from larger constants and is not always better in practical applications.

To highlight the difficulty of developing optimal algorithms for the ALL- ϵ problem, we quickly discuss a naive elimination approach that uniformly samples all arms and eliminates arms once they are known to be above or below $\mu_1 - \epsilon$ and not the best arm. Intuitively, such an algorithm would keep pulling arms until $\mu_1 - \epsilon$ is estimated to an accuracy of $O(\min(\alpha_\epsilon, \beta_\epsilon))$ to resolve the arms around the threshold (see Figure 2). An elimination algorithm pays a high cost of exploration - potentially over pulling arms close to μ_1 compared to the lower bound until a time when $\mu_1 - \epsilon$ is estimated sufficiently well. Our algorithm FAREAST provides a novel approach to overcome the issues with this approach. However, as we will show in Section 4, in certain instances a dependence on $\sum_{i=1}^n (\mu_1 + \beta_\epsilon - \mu_i)^{-2}$ is present in *moderate confidence*, i.e., it is not multiplied by $\log(1/\delta)$, unlike the lower bound in equation (2) and becomes negligible compared to other terms as $\delta \rightarrow 0$.

Empirical results. We demonstrate the empirical success of (ST)² on a real world dataset of 9250 captions from the NYCCC. In Fig. 4a, we compare (ST)² to other methods that have been used to run this contest. We show that (ST)² is better able to detect which arms have means within 10% of the best. The plot demonstrates the sub-optimality of using existing sampling schemes such as UCB or LUCB with an incorrect k for the ALL- ϵ problem, providing an additional empirical validation for the study of this paper.

1.3 Connections to prior Bandit art

Our problem is related to several prior pure-exploration settings in the multi-armed bandit literature, including TOP- k bandits, and threshold bandits.




Figure 2: An example instance

TOP- k . In the TOP- k problem, the goal is to identify the set $\{\mu_1, \dots, \mu_k\}$ with probability greater than $1 - \delta$ [4–9]. The ALL- ϵ problem reduces to the setting of the TOP- k problem with $k = |G_\epsilon|$ when $|G_\epsilon|$ is known. In particular, lower bounds for the TOP- k problem apply to our setting. A lower bound (with precise logarithmic factors) given in [9] is $\sum_{i=1}^k (\mu_i - \mu_{k+1})^{-2} \log((n-k)/\delta) + \sum_{i=k+1}^n (\mu_i - \mu_k)^{-2} \log(k/\delta)$. In general, this is smaller than our lower bound in Theorem 2.1 since $\mu_k \geq \mu_1 - \epsilon \geq \mu_{k+1}$. A particular case of this problem is best-arm identification when $k = 1$.

Approximate versions of the TOP- k problem have also been considered where the goal is to return a set of arms \mathcal{S} with $|\mathcal{S}| = k$ and such that with probability greater than $1 - \delta$, each $i \in \mathcal{S}$ satisfies $\mu_i \geq \mu_k - \epsilon$ [4, 10]. In the case where $k = 1$, this is also known as the problem of identifying an (single) ϵ -good arm [4, 7, 9–17] which has received a large amount of interest. If $|G_\epsilon| = k$, [6], demonstrate a lower bound of $O((k\epsilon^{-2} + \sum_{i=k+1}^n (\mu_1 - \mu_i)^{-2}) \log(1/\delta))$ samples in expectation to find such an arm and [10] provide an algorithm that matches this to doubly logarithmic factors, though methods such as [4, 9, 18, 19] achieve better empirical performance. A particular instance of interest is when it is known that one arm is at mean ϵ , and the rest are at mean zero. In this setting, [11] show a lower bound on the sample complexity of $O(n/\epsilon^2 + 1/\epsilon^2 \log(1/\delta))$ highlighting that the dependence on n only occurs in *moderate confidence*, i.e., for a fixed value of δ . They also provide a matching upper bound that motivates our procedure in FAREAST. Finally [15] considers the *unverifiable* regime where there are potentially many ϵ -good arms. In such cases, sample-efficient algorithms exist that return an ϵ -good arm with high probability, but *verifying* it is ϵ -good requires far more samples. Extending these ideas to the setting of ALL- ϵ is a goal of future work.

Threshold Bandits. In the threshold bandit problem, we are given a threshold τ and the goal is to identify the set of arms whose means are greater than the threshold [1, 20]. If the value of μ_1 were known, then ALL- ϵ problem would reduce to a threshold bandit with $\tau = \mu_1 - \epsilon$. A naive sequential sampling scheme that stops sampling an arm when its upper or lower confidence bound clears the threshold has sample complexity $O(\sum_{i=1}^n (\mu_i - \tau)^{-2} \log(n/\delta))$. Up to factors of $\log(n)$, this can be shown to be a lower bound for threshold bandits as well, and as a result is bounded above by the result Theorem 2.1. Hence, ALL- ϵ is intrinsically more difficult than threshold bandits. A naive approach to the ALL- ϵ problem is to first identify the index and mean of the best arm using a best-arm identification algorithm and then utilize it to build an estimate of the threshold $\mu_1 - \epsilon$. In general, this two-step procedure is sub-optimal if there are many arms close to the best-arm in which case identifying the best-arm is both unnecessary and expends unnecessary samples. In the fixed confidence setting, threshold bandits is closely related to that of multiple hypothesis testing, and recent work [21] achieves tight upper and lower bounds for this problem including tighter logarithmic factors similar to those for TOP- k . If μ_1 is known, then the additive ALL- ϵ problem reduces to the FWER (family-wise error rate) and FWER (family-wise probability of detection) setting in [21]. Finally, in the fixed budget setting, [1] proposes an optimal anytime method APT whose sampling strategy we use as a comparison in Section 5.

2 Lower Bound

Theorem 2.1. (*additive and multiplicative lower bounds*) Fix $\delta, \epsilon > 0$. Consider n arms, such that the i^{th} is distributed according to $\mathcal{N}(\mu_i, 1)$. Any δ -PAC algorithm for the *additive* setting satisfies

$$\mathbb{E}[\tau] \geq 2 \sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right\} \log \left(\frac{1}{2.4\delta} \right)$$

and if $\mu_1 > 0$, any δ -PAC algorithm for the *multiplicative* algorithm satisfies,

$$\mathbb{E}[\tau] \geq 2 \sum_{i=1}^n \max \left\{ \frac{1}{((1 - \epsilon)\mu_1 - \mu_i)^2}, \frac{1}{(\mu_1 + \frac{\alpha_\epsilon}{1 - \epsilon} - \mu_i)^2} \right\} \log \left(\frac{1}{2.4\delta} \right).$$

The bounds are different but share a common interpretation. Consider the *additive* case. First, every arm must be sampled inversely proportional to its squared distance to $\mu_1 - \epsilon$. In a manner similar to thresholding [1], even if $\mu_1 - \epsilon$ was known, these number of samples are necessary to decide if an arm's mean is above or below that quantity. This leads to the first term in the $\max\{\cdot, \cdot\}$. The second term in the $\max\{\cdot, \cdot\}$ states that every arm must be sampled inversely proportional to its squared distance to $\mu_1 + \alpha_\epsilon$. Recall that $\alpha_\epsilon = \mu_k - (\mu_1 - \epsilon)$ is the margin by which arm k is good.

Hence, to verify that $k \in G_\epsilon$, it is also necessary to confirm that all means are below $\mu_1 + \alpha_\epsilon$, as $\mu_1 + \alpha_\epsilon - \epsilon \geq \mu_k$ which would imply that k is bad. This represents the necessity of estimating the threshold, and leads to the second term. For arms in G_ϵ^c , comparing against $\mu_1 - \epsilon$ is always more difficult, but for arms in G_ϵ , either constraint may be more challenging to ensure. We state the bound for Gaussian distributions, but the same technique can be used to prove equivalent results for other distributions. Lastly, we note that it is possible to prove bounds with tighter logarithmic terms. For an instance where $O(n^\phi)$ arms have mean 2ϵ for $\phi \in (0, 1)$, and the remaining have mean 0, Theorem 1 of [22] suggests that $\Omega(n/\epsilon^2 \log(n/\delta))$ samples are necessary, exceeding the above bounds by a factor of $\log(n)$.

3 An Optimism Algorithm for ALL- ϵ

We propose algorithm 1 called $(ST)^2$, (Sample the Threshold, Split the Threshold) to return a set containing all ϵ -good arms and none worse than $(\epsilon + \gamma)$ -good with probability $1 - \delta$. Intuitively, $(ST)^2$ runs UCB and LUCB1 in parallel. At all times, $(ST)^2$ pulls three arms. We pull the arm with the highest upper confidence bound, similarly to the UCB algorithm, [3], to refine an estimate of the threshold using the highest empirical mean (Sample the Threshold). Using the empirical estimate of the threshold, we pull an arm above it and an arm below it whose confidence bounds cross it, similar to LUCB1, [4] (Split the Threshold). Using these bounds, $(ST)^2$ forms upper and lower bounds on the true threshold, i.e. $\mu_1 - \epsilon$ (resp. $(1 - \epsilon)\mu_1$) and terminates when it can declare that all arms are either in $G_{\epsilon+\gamma}$ or G_ϵ^c . To do so, $(ST)^2$ maintains anytime confidence widths, $C_{\delta/n}(t)$ such that for an empirical mean $\hat{\mu}_i(t)$ of t samples, we have $\mathbb{P}(\bigcup_{t=1}^\infty |\hat{\mu}_i(t) - \mu_i| > C_{\delta/n}(t)) \leq \delta/n$. For this work, we take $C_\delta(t) = \sqrt{\frac{c_\phi \log(\log_2(2t)/\delta)}{t}}$ for a constant c_ϕ . It suffices to take $c_\phi = 4$, though tighter bounds are known and should be used in practice, e.g. [6, 23, 24].

Algorithm 1 $(ST)^2$: Sample the Threshold, Split the Threshold

Require: $\epsilon, \delta > 0, \gamma \geq 0$, instance ν

- 1: Pull each arm once, initialize $T_i \leftarrow 1$, update $\hat{\mu}_i$ for each $i \in \{1, 2, \dots, n\}$
- 2: Empirically good arms: $\hat{G} = \{i : \hat{\mu}_i \geq \max_j \hat{\mu}_j - \epsilon\}$, $\hat{G} = \{i : \hat{\mu}_i \geq (1 - \epsilon) \max_j \hat{\mu}_j\}$
- 3: $U_t = \max_j \hat{\mu}_j(T_j) + C_{\delta/n}(T_j) - \epsilon - \gamma$ and $L_t = \max_j \hat{\mu}_j(T_j) - C_{\delta/n}(T_j) - \epsilon$
- 4: $U_t = (1 - \epsilon - \gamma) (\max_j \hat{\mu}_j(t) + C_{\delta/n}(T_j))$ and $L_t = (1 - \epsilon) (\max_j \hat{\mu}_j(t) - C_{\delta/n}(T_j))$
- 5: Known arms: $K = \{i : \hat{\mu}_i(T_i) + C_{\delta/n}(T_i) < L_t \text{ or } \hat{\mu}_i(T_i) - C_{\delta/n}(T_i) > U_t\}$
- 6: **while** $K \neq [n]$ **do**
- 7: Pull arm $i_1(t) = \arg \min_{i \in \hat{G} \setminus K} \hat{\mu}_i(T_i) - C_{\delta/n}(T_i)$, update $T_{i_1}, \hat{\mu}_{i_1}$
- 8: Pull arm $i_2(t) = \arg \max_{i \in \hat{G}^c \setminus K} \hat{\mu}_i(T_i) + C_{\delta/n}(T_i)$, update $T_{i_2}, \hat{\mu}_{i_2}$
- 9: Pull arm $i^*(t) = \arg \max_i \hat{\mu}_i(T_i) + C_{\delta/n}(T_i)$, update $T_{i^*}, \hat{\mu}_{i^*}$
- 10: Update bounds L_t, U_t , sets \hat{G}, K
- return** The set of good arms $\{i : \hat{\mu}_i(T_i) - C_{\delta/n}(T_i) > U_t\}$

3.1 Theoretical guarantees

Next we present a pair of theorems on the sample complexity of $(ST)^2$. For clarity, we omit doubly logarithmic terms and defer such statements to Appendix B. Below we denote $a \wedge b := \min\{a, b\}$.

Theorem 3.1 (Additive Case). *Fix $\epsilon > 0, 0 < \delta \leq 1/2, \gamma \leq 16$ and an instance ν such that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i . With probability at least $1 - \delta$, there is a constant c_1 such that $(ST)^2$ returns a set \hat{G} such that $G_\epsilon \subset \hat{G} \subset G_{(\epsilon+\gamma)}$ in at most the following number of samples.*

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \right\} \wedge \frac{1}{\gamma^2} \quad (3)$$

Given a positive slack γ , we are allowed to return an arm that is $(\epsilon + \gamma)$ -good. Thus a confidence width less than $\Omega(\gamma)$ on any arm is not needed, resulting in the $1/\gamma^2$ term in Theorem 3.1. In particular this prevents unbounded sample complexities if there is an arm at the threshold $\mu_1 - \epsilon$. For $\gamma = 0$, the first two terms inside the max are also present in the lower bound (Theorem 2.1). When α_ϵ is within a constant factor of β_ϵ , the second and third term in the max have the same order, and the upper bound matches the lower bound up to a $\log(n)$ factor.

If $\beta_\epsilon \ll \alpha_\epsilon$, (3) has a different scaling than the lower bound. In such restrictive settings the upper bound above can be significantly larger than the lower bound. In the next section, we provide an algorithm that overcomes these issues and is optimal over all parameter regimes. The **multiplicative** case has different terms but follows the same intuition.

Theorem 3.2 (Multiplicative Case). Fix $\epsilon \in (0, 1/2]$, $\gamma \in [0, \min(16/\mu_1, 1/2)]$ and $0 < \delta \leq 1/2$ and an instance ν such that $\mu_1 \geq 0$ and $\max(\Delta_i, |\epsilon\mu_1 - \Delta_i|) \leq 2$ for all i . With probability at least $1 - \delta$, for a constant c_1 $(ST)^2$ returns a set G such that $M_\epsilon \subset G \subset M_{(\epsilon+\gamma)}$ with sample complexity:

$$c_1 \log\left(\frac{n}{\delta}\right) \sum_{i=1}^n \max \left\{ \frac{1}{((1-\epsilon)\mu_1 - \mu_i)^2}, \frac{1}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2}, \frac{1}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \right\} \wedge \frac{1}{\gamma^2 \mu_1^2}.$$

4 Surprising Complexity of Finding All ϵ -Good arms

When α_ϵ and β_ϵ are not of the same order, $(ST)^2$ is not optimal. In this section we present an algorithm that is optimal for all parameter regimes. We focus on the **additive** case here, and defer the **multiplicative** case to Appendix E. We first state an improved sample complexity lower bound for a family of problem instances that makes explicit the *moderate confidence* terms.

Theorem 4.1. Fix $\delta \leq 1/16$, $n \geq 2/\delta$, and $\epsilon > 0$. Let ν be an instance of n arms such that the i^{th} is distributed as $\mathcal{N}(\mu_i, 1)$, $|G_{2\beta_\epsilon}| = 1$, and $\beta_\epsilon < \epsilon/2$. Select a permutation $\pi : [n] \rightarrow [n]$ uniformly from the set of $n!$ permutations, and consider the permuted instance $\pi(\nu)$. Any algorithm that returns $G_\epsilon(\pi(\nu))$ on $\pi(\nu)$ correctly with probability at least $1 - \delta$ requires at least the following number of samples in expectation over randomness in ν and π for a universal constant c_2 .

$$\left[c_2 \sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right\} \log\left(\frac{1}{2.4\delta}\right) \right] + c_2 \sum_{i=1}^n \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \quad (4)$$

Proof. (Sketch) To give a tight lower bound in the setting where $|G_{2\beta_\epsilon}| = 1$ and $\beta_\epsilon < \epsilon/2$, we break our argument into pieces performing a series of reductions that link the ALL- ϵ problem to a hypothesis test, and then the hypothesis test to the problem of identifying the best-arm. We apply the Simulator technique from [9] to compute precise moderate confidence bounds. Other works that prove strong lower bounds in moderate confidence include [25]. We extend the Simulator technique via a novel reduction to composite hypothesis testing in order to connect to ALL- ϵ . In all cases, we consider sample complexity in expectation with respect to the randomness in the outcomes and a randomly chosen permutation of the means.

Step 1. Finding an isolated best arm: Consider the problem of finding the best arm where $\mu_1 = \beta > 0$ and $\mu_2, \dots, \mu_n \leq -\beta$. This relates to the problem of finding a β -good arm when μ_1 is known, studied by [11]. We use the Simulator technique, [9], to show that any algorithm requires $\Omega(\sum_{i=2}^n \Delta_i^{-2})$ samples in expectation. This can be significantly larger than the asymptotically optimal rate of $O(\beta^{-2} \log(1/\delta))$ (which was proven by [11]) for non-asymptotic δ , e.g. $\delta = 0.05$.

Step 2. Deciding if Any mean is positive: We then consider a composite hypothesis test on n distributions where the null hypothesis, H_0 , is that the mean of each distribution is less than $-\beta$ and the alternate hypothesis, H_1 , is that there exists a *single* distribution i^* with mean β and the remainder have mean less than $-\beta$. Importantly, an algorithm does not need to declare which arm is i^* , otherwise the bound from step 1 applies immediately. Instead, to link this to step 1, we develop a novel extension of the simulator technique and use this to show that if an algorithm can solve this composite hypothesis test in fewer than $o(\sum_{i=2}^n \Delta_i^{-2})$ samples, then one may design a method to solve the problem in step 1 in $o(\sum_{i=2}^n \Delta_i^{-2})$ samples which is a contradiction. Hence any algorithm for this hypothesis test requires $\Omega(\sum_{i=2}^n \Delta_i^{-2})$ samples in expectation.

Step 3: Reducing ALL- ϵ to Step 2: Finally, we show that a generic algorithm for ALL- ϵ can be used to solve the hypothesis test in step 2. Hence the lower bound from step 2 applies to finding all ϵ -good arms as well. In the case of the instances considered in the theorem statement, $O(\sum_{i=2}^n \Delta_i^{-2}) = O(\sum_{i=2}^n (\mu_1 + \beta_\epsilon - \mu_i)^{-2})$. Combining this bound, which is independent of δ with the result from Theorem 2.1 gives the result. \square

Theorem 4.1 states that an additional $\Omega(\sum_{i=1}^n (\mu_1 + \beta_\epsilon - \mu_i)^{-2})$ samples are necessary for instances where no arm is within $2\beta_\epsilon$ of μ_1 compared to the lower bound Theorem 2.1. Somewhat surprisingly, these samples are *necessary in moderate confidence*, independent of δ and negligible as $\delta \rightarrow 0$. For non-asymptotic values of δ , such as the common choice of $\delta = .05$ in scientific applications, this term is present and can even dominate the sample complexity when $\beta_\epsilon \ll \alpha_\epsilon$. As an extreme example, if $\mu_1 = \beta > 0$, $\mu_2 \dots, \mu_{n-1} = -\beta$, $\mu_n = -\epsilon$, the first term in 4 scales like $((n-1)/\epsilon^2 + 1/\beta^2) \log(1/\delta)$ but the second term scales like n/β^2 , which is $O(n)$ larger than the first term for small β and fixed δ . Furthermore, we point out that Theorem 4.1 highlights that $(ST)^2$ is optimal on these instances up to a log factor! The algorithm we present next, FAREAST, improves $(ST)^2$'s dependence on δ and matches the lower bound in Theorem 4.1 for certain instances. Though moderate confidence terms can dominate the sample complexity in practice, few works have focused on understanding their effect.

4.1 FAREAST

We focus on the **additive** case with $\gamma = 0$ in Algorithm 4.1, FAREAST, and defer the more general case (**multiplicative** and $\gamma > 0$) to Algorithm E.1 in the supplementary. FAREAST matches the instance dependent lower bound in Theorem 2.1 as $\delta \rightarrow 0$. At a high level, FAREAST (**F**ast **A**rm **R**emoval **E**limination **A**lgorithm for a **S**ampled **T**hreshold) proceeds in rounds r and maintains sets \hat{G}_r and \hat{B}_r of arms thus far declared to be good or bad. It sorts unknown arms into either set through use of a good filter to detect arms in G_ϵ and a bad filter to detect arms in G_ϵ^c .

Good Filter: The good filter is a simple elimination scheme. It maintains an upper bound U_t and lower bound L_t on $\mu_1 - \epsilon$. If an arm's upper bound drops below L_t (line 20), the good filter eliminates that arm, otherwise, if an arm's lower bound rises above U_t (19), the good filter adds the arm to \hat{G}_r , but only eliminates this arm if its upper bound falls below the highest lower bound. This ensures that μ_1 is never eliminated and U_t and L_t are always valid bounds¹. As the sampling is split across rounds, the good filter always samples the least sampled arm, breaking ties arbitrarily. The number of samples given to the good filter in each round is such that both filters receive identically many samples. This prevents the good filter from over-sampling bad arms and vice versa. In our proof we show that in an unknown round, $\hat{G}_r = G_\epsilon$, ie all good arms have been found, having used fewer than $O(\sum_{i=1}^n \max\{(\mu_1 - \epsilon - \mu_i)^{-2}, (\mu_1 + \alpha_\epsilon - \mu_i)^{-2}\} \log(n/\delta))$ samples, matching the lower bound.

FAREAST cannot yet terminate, however, as it must also verify that any remaining arms are in G_ϵ^c .

Bad Filter: The bad filter removes arms that are not ϵ -good. To show an arm i is in G_ϵ^c , it suffices to find any j such that $\mu_j - \mu_i > \epsilon$. To motivate the idea of lines 9-12, consider the following procedure in the special case where $\beta_i = \mu_1 - \epsilon - \mu_i$ is known. In each round we first run Median-Elimination, [12], with failure probability $1/16$, to find an arm \hat{i} that is $\beta_i/2$ -good in $O(n/\beta_i^2)$ samples². We then pull both i and \hat{i} roughly $O(1/\beta_i^2 \log(1/\delta))$ times and can check whether $\mu_{\hat{i}} - \mu_i > \epsilon$ with probability greater than $1 - \delta$. This procedure relies on Median-Elimination succeeding, which happens with probability $15/16$. In the case that it fails and we declare $\mu_{\hat{i}} - \mu_i < \epsilon$, we merely repeat this process until it succeeds—on average $O(1)$ times. This gives an expected sample complexity of $O(n/\beta_i^2 + 1/\beta_i^2 \log(1/\delta))$ for any $i \in G_\epsilon^c$. Of course, β_i is unknown to the algorithm. Instead, in each round r , the bad filter guesses that $\beta_i \geq 2^{-r}$ for all unknown arms $i \notin \hat{G}_r \cup \hat{B}_r$ and performs the above procedure. The following theorem demonstrates that this algorithm matches our lower bounds asymptotically as $\delta \rightarrow 0$.


Theorem 4.2. Fix $0 < \epsilon$, $0 < \delta < 1/8$, and an instance ν of n arms such that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i . There exists an event E such that $\mathbb{P}(E) \geq 1 - \delta$ and on E , FAREAST terminates and returns G_ϵ . Letting T denote the number of samples taken, for a constant c_3

$$\mathbb{E}[\mathbb{1}_E T] \leq \left[c_3 \sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right\} \log \left(\frac{n}{\delta} \right) \right] + c_3 \sum_{i \in G_\epsilon^c} \frac{c'' n}{(\mu_1 - \epsilon - \mu_i)^2}.$$

Additionally for $\gamma \leq 16$ FAREAST terminates on E and returns a set \hat{G} such that $G_\epsilon \subset \hat{G} \subset G_{\epsilon+\gamma}$ in a number of samples no more than a constant times (3), the complexity of $(ST)^2$.

¹This scheme works as an independent algorithm, we analyze it in Appendix E.5.

²Median-Elimination is used for ease of analysis. One can use LUCB [4] or another method instead.



(a) A typical setting of diverse means (black dots).

(b) A more challenging setting

Figure 3: Comparison of $(ST)^2$ and FAREAST averaged over 250 trials plotted with 3 standard errors.

Algorithm 4.1: additive FAREAST with $\gamma = 0$

Input: ϵ, δ , instance ν

Let $\hat{G}_0 = \emptyset$ be the set of arms declared as good and $\hat{B}_0 = \emptyset$ the set of arms declared as bad.

Let $\mathcal{A} = [n]$ be the active set, $N_i = 0$ track the total number of samples of arm i by the Good Filter.

Let $t = 0$ denote the total number of times that line 16 is true in the Good Filter.

for $r = 1, 2, \dots$

 Let $\delta_r = \delta/2^{r^2}$, $\tau_r = \left\lceil 2^{2r+3} \log\left(\frac{8n}{\delta_r}\right) \right\rceil$, Initialize $\hat{G}_r = \hat{G}_{r-1}$ and $\hat{B}_r = \hat{B}_{r-1}$

// Bad Filter: find bad arms in G_ϵ^c

 Let $i_r = \text{MedianElimination}(\nu, 2^{-r}, 1/16)$, sample i_r τ_r times and compute $\hat{\mu}_{i_r}$

for $i \notin \hat{G}_{r-1} \cup \hat{B}_{r-1}$:

 Sample μ_i τ_r times and compute $\hat{\mu}_i$

If $\hat{\mu}_{i_r} - \hat{\mu}_i \geq \epsilon + 2^{-r+1}$: Add i to \hat{B}_r **// Bad arm detected**

// Good Filter: find good arms in G_ϵ

for $s = 1, \dots, H_{ME}(n, 2^{-r}, 1/16) + (|\hat{G}_{r-1} \cup \hat{B}_{r-1}| + 1)\tau_r$:

 Pull arm $I_s \in \arg \min_{j \in \mathcal{A}} \{N_j\}$ and set $N_{I_s} \leftarrow N_{I_s} + 1$.

if $\min_{j \in \mathcal{A}} \{N_j\} = \max_{j \in \mathcal{A}} \{N_j\}$:

 Update $t = t + 1$. Let $U_t = \max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/2n}(t) - \epsilon$ and $L_t = \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t) - \epsilon$

for $i \in \mathcal{A}$:

if $\hat{\mu}_i(t) - C_{\delta/2n}(t) \geq U_t$: Add i to \hat{G}_r **// Good arm detected**

if $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq L_t$: Remove i from \mathcal{A} and add i to \hat{B}_r **// Bad arms removed**

if $i \in \hat{G}_r$ and $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$: **// Good arms removed**

 Remove i from \mathcal{A}

if $\mathcal{A} \subset \hat{G}_r$ or $\hat{G}_r \cup \hat{B}_r = [n]$: **Return the set \hat{G}_r**

5 Empirical Performance

We begin by comparing $(ST)^2$ and FAREAST on simulated data. FAREAST is asymptotically optimal, but suffers worse constant factors compared to $(ST)^2$ ³. $(ST)^2$ is optimal *except* when $\beta_\epsilon \ll \alpha_\epsilon$. We compare $(ST)^2$ and FAREAST on two instances in the **additive** case, shown in Figure 3. All arms are Gaussian with $\sigma = 1$. In the first example on the left, $\delta = 0.1$, $\alpha_\epsilon = \beta_\epsilon = 0.05$. Both $(ST)^2$ and FAREAST are optimal in this setting; we show the scaling of their sample complexity as the number of arms increases while keeping the threshold, α_ϵ , and β_ϵ constant. In the second example, $\alpha_\epsilon = \epsilon = 0.99$, and $\beta = 0.01$. When $1/\beta_\epsilon^2 \gg n/\epsilon^2$, Theorem 2.1 suggests that $O(1/\beta_\epsilon^2 \log(1/\delta))$ samples are necessary, independent of n . Indeed, in Figure 3, for $\delta = 0.01$, the average complexity of FAREAST is constant, but $(ST)^2$ scales linearly with n as Theorem 3.1 suggests. Finally, a naive uniform sampling strategy performed very poorly - additional experiments including the uniform sampling method and with $\gamma > 0$ are in the Appendix A.

5.1 Finding all ϵ -good arms in real world data – fast

As discussed in the introduction, in many applications such as the New Yorker Cartoon Caption Contest (NYCCC), the ALL- ϵ objective returns a set of good arms which can then be screened further to choose a favorite. We considered Contest 651, which had 9250 captions whose means we estimated from a total of 2.2 million ratings. We set $\epsilon = 0.1$ and focus on the multiplicative setting, i.e., the objective of recovering all captions within 10% of the funniest one. In this experiment, we contrast $(ST)^2$ with several other methods including two *oracle* methods (marked with \blacktriangle): LUCB1 [4] with

³Implementations of all algorithms and baselines used in this paper are available on [GitHub](#).

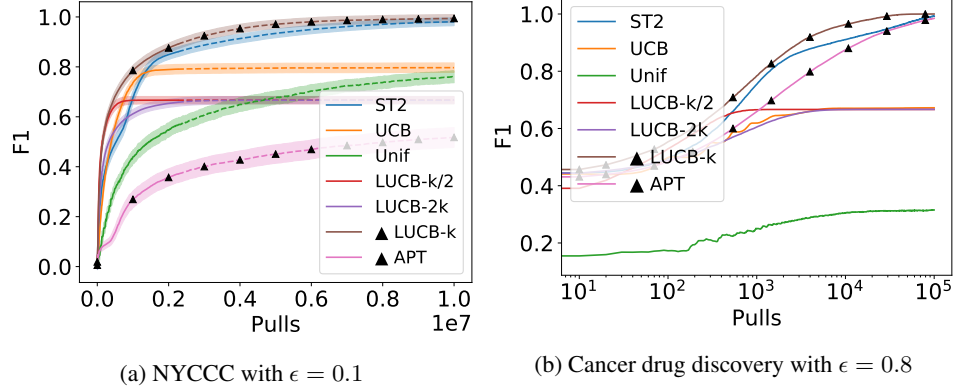


Figure 4: F1 scores averaged over 600 trials with 95% confidence widths for each dataset.




Figure 5: Precision and recall averaged over 600 trials with 95% confidence widths on NYCCC data.

k set to the number of ϵ -good arms (here it was 46), and a threshold-bandit, APT [1] given the value of $0.9\mu_1$. We focus on a common practical requirement, each algorithm’s ability to balance precision and recall as it samples. With every new sample, each method recommends an empirical set of ϵ -good arms based on the empirical means, and we consider the F1 score of this set⁴. We focus on the F1 score as it is practically relevant and provides a continuous measure of performance of each method. $F1 = 1$ indicates that an algorithm has found all ϵ -good arms. As can be seen in Figure 4a, $(ST)^2$ outperforms all baselines including the oracle APT, and almost matches the performance of the TOP- k oracle! We transition from a solid line to a dashed one at 2.2M pulls to mark the number of samples drawn in the real contest from which we gather the data. To illustrate the importance of knowing the correct value of k , we also plot LUCB1 given $k = 46/2 = 23$ and $k = 46 \times 2 = 92$, settings where the experimenter under or over estimates the number of ϵ good arms by as little as a factor of 2. Both cases result in a poor performance. We have also included UCB, currently being used for the contest [26]; the plot shows that UCB is not able to estimate the ϵ -good set. In Figure 5, we show precision and recall curves for each method on the NYCCC data. $(ST)^2$ achieves near-perfect precision quickly, matched only by UCB. APT’s poor performance is a consequence of having low-precision, shown in Figure 5a. $(ST)^2$ achieves high recall more slowly, but is still competitive with other methods. In practical experiments, high precision early on may be more important than high recall, as it guarantees that practitioners can trust the declarations that the algorithm has made, even if some arms are yet to be found. In the Supplementary we show plots for more values of ϵ . Additionally, motivated by drug discovery, we performed an experiment on a dataset [27] of 189 inhibitors whose activities were tested against ACVRL1, a kinase associated with cancer [28]. In this experiment, we use the multiplicative case of ALL- ϵ with $\epsilon = 0.8$ and $\delta = 0.001$, to promote high precision. In this experiment as well, $(ST)^2$ performs best (Figure 4b), with only the oracle methods are competitive with it. We plot on a log-scale to emphasize the early regime.

⁴F1 is the harmonic mean of precision (fraction of captions returned that are actually good) and recall (fraction of all good captions that are actually returned).

6 Broader Impacts and Funding Transparency Statement

6.1 Broader Impacts

The application of machine learning (ML) in domains such as advertising, biology, or medicine brings the possibility of utilizing large computational power and large datasets to solve new problems. It is tempting to use powerful, if not fully understood, ML tools to maximize scientific discovery. However, at times the gap between a tool’s theoretical guarantees and its practical performance can lead to sub-optimal behavior. This is especially true in *adaptive data collection* where misspecifying the model or desired output (e.g., “return the top k performing compounds” vs. “return all compounds with a potency about a given threshold”) may bias data collection and hinder post-hoc consideration of different objectives. In this paper we highlight several such instances in real-life data collection using multi-armed bandits where such a phenomenon occurs. We believe that the objective studied in this work, that of returning all arms whose mean is quantifiably near-best, more naturally aligns with practical objectives as diverse as finding funny captions to performing medical tests. We point out that methods from adaptive data collection and multi-armed bandits can also be used on content-recommendation platforms such as social media or news aggregator sites. In these scenarios, time and again, we have seen that recommendation systems can be greedy, attempting purely to maximize clickthrough with a long term effect of a less informed public. Adjacent to one of the main themes of this paper, we recommend that practitioners not just focus on the objective of recommendation for immediate profit maximization but rather keep track of a more holistic set of metrics. We are excited to see our work used in practical applications and believe it can have a major impact on driving the process of scientific discovery.

6.2 Funding Transparency Statement

The work presented in this paper was supported by ARO grant W911NF-15-1-0479. Additionally, this work was partially supported by the MADLab AF Center of Excellence FA9550-18-1-0166.

References


- [1] Andrea Locatelli, Maurilio Gutzeit, and Alexandra Carpentier. An optimal algorithm for the thresholding bandit problem. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning-Volume 48*, pages 1690–1698. JMLR. org, 2016.
- [2] Serge Christmann-Franck, Gerard JP van Westen, George Papadatos, Fanny Beltran Escudie, Alexander Roberts, John P Overington, and Daniel Domine. Unprecedentedly large-scale kinase inhibitor set enabling the accurate prediction of compound–kinase activities: A way toward selective promiscuity by design? *Journal of chemical information and modeling*, 56(9):1654–1675, 2016.
- [3] Peter Auer, Nicolo Cesa-Bianchi, and Paul Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- [4] 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.
- [5] Sébastien Bubeck, Tengyao Wang, and Nitin Viswanathan. Multiple identifications in multi-armed bandits. In *International Conference on Machine Learning*, pages 258–265, 2013.
- [6] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.
- [7] Victor Gabillon, Mohammad Ghavamzadeh, and Alessandro Lazaric. Best arm identification: A unified approach to fixed budget and fixed confidence. In *Advances in Neural Information Processing Systems*, pages 3212–3220, 2012.
- [8] Wenbo Ren, Jia Liu, and Ness B Shroff. Exploring k out of top p fraction of arms in stochastic bandits. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 2820–2828, 2019.

- [9] 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.
- [10] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pages 1238–1246, 2013.
- [11] Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun):623–648, 2004.
- [12] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pages 255–270. Springer, 2002.
- [13] 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.
- [14] Shivaram Kalyanakrishnan and Peter Stone. Efficient selection of multiple bandit arms: Theory and practice. In *ICML*, volume 10, pages 511–518, 2010.
- [15] Julian Katz-Samuels and Kevin Jamieson. The true sample complexity of identifying good arms. *arXiv preprint arXiv:1906.06594*, 2019.
- [16] Rémy Degenne and Wouter M Koolen. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems*, pages 14564–14573, 2019.
- [17] Emilie Kaufmann and Shivaram Kalyanakrishnan. Information complexity in bandit subset selection. In *Conference on Learning Theory*, pages 228–251, 2013.
- [18] Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan. Pac identification of a bandit arm relative to a reward quantile. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [19] Arghya Roy Chaudhuri and Shivaram Kalyanakrishnan. Pac identification of many good arms in stochastic multi-armed bandits. In *International Conference on Machine Learning*, pages 991–1000, 2019.
- [20] Hideaki Kano, Junya Honda, Kentaro Sakamaki, Kentaro Matsuura, Atsuyoshi Nakamura, and Masashi Sugiyama. Good arm identification via bandit feedback. *Machine Learning*, 108(5):721–745, 2019.
- [21] Kevin Jamieson and Lalit Jain. A bandit approach to multiple testing with false discovery control. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS’18, page 3664–3674, Red Hook, NY, USA, 2018. Curran Associates Inc.
- [22] Matthew L Malloy and Robert D Nowak. Sequential testing for sparse recovery. *IEEE Transactions on Information Theory*, 60(12):7862–7873, 2014.
- [23] 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.
- [24] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Uniform, nonparametric, non-asymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.
- [25] 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.
- [26] Ervin Tanczos, Robert Nowak, and Bob Mankoff. A kl-lucb algorithm for large-scale crowd-sourcing. In *Advances in Neural Information Processing Systems*, pages 5894–5903, 2017.
- [27] David H Drewry, Carrow I Wells, David M Andrews, Richard Angell, Hassan Al-Ali, Alison D Axtman, Stephen J Capuzzi, Jonathan M Elkins, Peter Ettmayer, Mathias Frederiksen, et al. Progress towards a public chemogenomic set for protein kinases and a call for contributions. *PloS one*, 12(8), 2017.
- [28] Matteo Bocci, Jonas Sjölund, Ewa Kurzejamska, David Lindgren, Michael Bartoschek, Mattias Höglund, Kristian Pietras, et al. Activin receptor-like kinase 1 is associated with immune cell infiltration and regulates clec14a transcription in cancer. *Angiogenesis*, 22(1):117–131, 2019.

- [29] Patricia Dranchak, Ryan MacArthur, Rajarshi Guha, William J Zuercher, David H Drewry, Douglas S Auld, and James Inglese. Profile of the gsk published protein kinase inhibitor set across atp-dependent and-independent luciferases: implications for reporter-gene assays. *PloS one*, 8(3), 2013.

Contents

1	Introduction	1
1.1	Problem Statement and Notation	2
1.2	Contributions and Summary of Main Results	3
1.3	Connections to prior Bandit art	3
2	Lower Bound	4
3	An Optimism Algorithm for ALL-ϵ	5
3.1	Theoretical guarantees	5
4	Surprising Complexity of Finding All ϵ-Good arms	6
4.1	FAREAST	7
5	Empirical Performance	8
5.1	Finding all ϵ -good arms in real world data – <i>fast</i>	8
6	Broader Impacts and Funding Transparency Statement	10
6.1	Broader Impacts	10
6.2	Funding Transparency Statement	10
A	Additional Experimental Results	15
B	(ST)², An optimism based algorithm for all-ϵ	20
B.1	Optimism with additive γ	20
B.2	Optimism with multiplicative γ	30
C	Proof of instance dependent lower bounds, Theorem 2.1	41
D	Theorem 4.1: Lower bounds in the moderate confidence regime	42
D.1	Step 1: Finding an Isolated Arm	43
D.2	Step 2. Deciding if an instance is isolated	46
D.3	Step 3: Reducing ALL- ϵ to isolated instance detection	50
E	An optimal method for finding all additive and multiplicative ϵ-good arms	51
E.1	The FAREAST Algorithm	51
E.2	Key ideas of the proof	54
E.3	Proof of Theorem 4.2, FAREAST in the additive regime	55
E.4	Proof of Theorem E.2, FAREAST in the multiplicative regime	68
E.5	An elimination algorithm for all ϵ	82
E.6	Proof of Theorem E.3 EAST in the additive regime	83
E.7	Proof of Theorem E.4, EAST in the multiplicative regime	88



(a) Plot in Figure 3a with uniform sampling included. (b) Plot in Figure 3b with uniform sampling included.

Figure 6: Simulation results with uniform sampling included.




Figure 7: $(ST)^2$ and FAREAST with different values of γ

A Additional Experimental Results

Practical change made to FAREAST for simulations: We make one change to FAREAST that we recommend for practitioners wishing to use FAREAST that improve its empirical performance. In particular, Median-Elimination may instead be replaced by another method, such as LUCB1, [4], to find ϵ -good arms. LUCB1, for instance, has better constant factors and enjoys improved empirical performance versus Median-Elimination. The use of Median-Elimination in this algorithm serves to ease both notation and analysis since its sample complexity is deterministic. To modify the algorithm, simply track the number of samples given to the bad filter in total, which can be a random variable, and give the good filter the same number in that round. The proof then follows identically, with only the moderate confidence term changing in the result.

Additional Simulations Results As mentioned in the Experiments, Section 5, we omitted curves comparing against uniform sampling as they make the plots hard to read with uniform performing much more poorly. For completeness, we include them in Figure 6. Clearly, uniform sampling performs much more poorly than either active method, as expected.

Additionally, we include experiments with $\gamma > 0$ here. For small γ , the only valid solution is G_ϵ (resp. M_ϵ) itself. However, for larger γ , there are many valid solutions. Indeed, any G such that $G_\epsilon \subset G \subset G_{\epsilon+\gamma}$ is valid. To analyze the effect of γ on both $(ST)^2$ and FAREAST, we consider the same type of instances studied in Figure 3b. Here, $n - 1$ arms have means equal to μ_1 , and a single arm is in G_ϵ^c . Again, we take $\epsilon = 0.99$ and $\beta_\epsilon = 0.01$, and additionally, set $n = 150$ arms. Recall that in this setting, FAREAST outperforms $(ST)^2$, as shown in Figure 3b. As we increase γ , the problem becomes easier. We increase γ on an exponential scale, beginning with $\gamma \approx \epsilon/100$ and ending with $\gamma \approx \epsilon/2$. Indeed, for smaller values of γ , FAREAST is superior as it finds the exact solution fastest. For larger γ , $(ST)^2$ is able to terminate more quickly. In Figure 7 we plot these results.

Metrics we consider for real data experiments: For all methods, we track their precision, recall and F1 score with respect to the true set of ϵ -good arms. To compute these metrics, at each time, the algorithm outputs a set that it guesses are the ϵ -good arms based on the data it has gathered thus far. For UCB, Uniform, and (ST)², this is based directly on empirical means, i.e., $\hat{G} = \{i : \hat{\mu}_i \geq \max_j \hat{\mu}_j - \epsilon\}$ or $\hat{G} = \{i : \hat{\mu}_i \geq \max_j (1 - \epsilon)\hat{\mu}_j\}$ in the multiplicative case. Oracle methods may use their additional information to return the set. In particular, APT returns all arms whose empirical means exceed $(1 - \epsilon)\mu_1$ (using knowledge of μ_1) and LUCB1 returns the k largest empirical means (using knowledge that $|M_\epsilon| = k$). Let TP (true positives) denote the number of arms that an algorithm declares as ϵ -good that truly are. Let FN (false negatives) denote the number of arms that an algorithm declares as *not* ϵ -good when in fact they are. Recall, $r \in [0, 1]$, is computed as $r = \frac{TP}{TP + FN}$. Intuitively, recall is the total number of ϵ -good arms that the algorithm detects. Precision, $p \in [0, 1]$, by contrast is the fraction of the arms that an algorithm predicts as ϵ -good that truly are. It is computed as $p = \max(TP/|\hat{G}|, 1)$ where the $\max()$ is necessary to avoid the trivial case that $\hat{G} = \emptyset$. Finally, the F1 is the harmonic mean of precision and recall: $F1 = \frac{2pr}{p+r}$. It balances how precise an algorithm is with how many discoveries it makes. In many cases, F1 may be a more relevant metric than the others, as it avoids trivial edge cases. For instance, an algorithm that always declares every arm as ϵ -good independent of the data, achieves perfect recall because it has 0 false negatives. Similarly, an algorithm that never declares any arms as ϵ -good, again independent of data, achieves perfect precision. Both methods, despite seemingly good performance with respect to their individual metrics, are undesirable in practice. In particular, both would achieve low F1 scores.

The New Yorker Caption Contest: In this section we provide additional experimental results adjoining those in Section 5. The data can be downloaded at <https://github.com/nextml/caption-contest-data>. We chose contest 651 for our experiments, but hundreds of others are available. Captions are rated on a scale of 1 to 3 (“unfunny”, “somewhat funny”, or “funny”). It is desirable to find all captions that are nearly as good as the best. However, setting a fixed number of captions or fraction of captions to accept is undesirable as the number of truly funny captions varies from week to week and represents a small fraction of the submissions. For instance, in the contest that ran the week of 3/14/16, only 8 captions were rated within 20% of the funniest caption. In the following week, by contrast, 187 captions were. Similarly, choosing a fixed threshold of what it means for a caption to be funny is unrealistic. In the same two contests, first week saw 3% of captions be rated at least 1.5 out of 3 whereas the second saw $< 0.1\%$. For this reason, finding all ϵ -good arms is more natural. We consider finding all **multiplicative** ϵ -good arms with $\epsilon = 0.1, 0.15, 0.2$. To keep the comparison fair, all methods use the same confidence widths from [24]. In Figure 9b we plot the average rating of each caption in sorted order with horizontal lines corresponding to $(1 - 0.2)\mu_1$, $(1 - 0.15)\mu_1$, and $(1 - 0.1)\mu_1$. The arms with means above this line are 0.2, 0.15, and 0.1 ϵ -good. The oracle methods tend to achieve high recall, but low precision, and this is especially true for the threshold oracle, APT. In Figures 10, 11, 12 we plot F1, Precision, and Recall curves for all methods tested on $\epsilon = 0.2, 0.15, 0.1$ respectively. As before, all curves are averaged over 600 independent repetitions and plotted with 95% confidence intervals. It is evident from these curves, that (ST)² performs especially well with regard to precision, though it achieves lower recall than some other baselines.

Protein Kinase Inhibitors for Cancer Drug Discovery

Additionally, we consider a second, medically focused experiment. In 2013, researchers at Glaxo-SmithKline published a dataset of protein kinase inhibitors different kinases (PKIS1), primarily from humans [29]. Kinases are a family of enzymes present in many cells and researchers are interested in developing targeted kinase inhibitors to as a new way to treat cancer [2]. The dataset contains numerous measures of how strongly each inhibitor reacts with each kinase. A second, larger dataset (PKIS2) was expanded on by [27]⁵. For the purpose of our experiment, we selected a single Kinase in the dataset, ACVRL1, which researchers have linked to numerous types of cancer, most prominently bladder and prostate cancers [28]. PKIS2 contains 641 different compounds that were tested as being potential kinase inhibitors, though not every compound was tested against every kinase. In particular, 189 were tested against ACVRL1. For each compound, there is an associated average “percent inhibition” that is reported. All numbers are between 0 and 1 and averaged across multiple trials in a single assay. We subtract each number from 1 to compute the percent control, representing how effective any method is relative to a control, an important metric for estimating how effective

⁵The dataset can be downloaded at the following link: <https://doi.org/10.1371/journal.pone.0181585.s004>.




Figure 8: The user interface for the caption contest with the caption for contest 651. “Unfunny” = 1, “Somewhat funny” = 2, “Funny” = 3

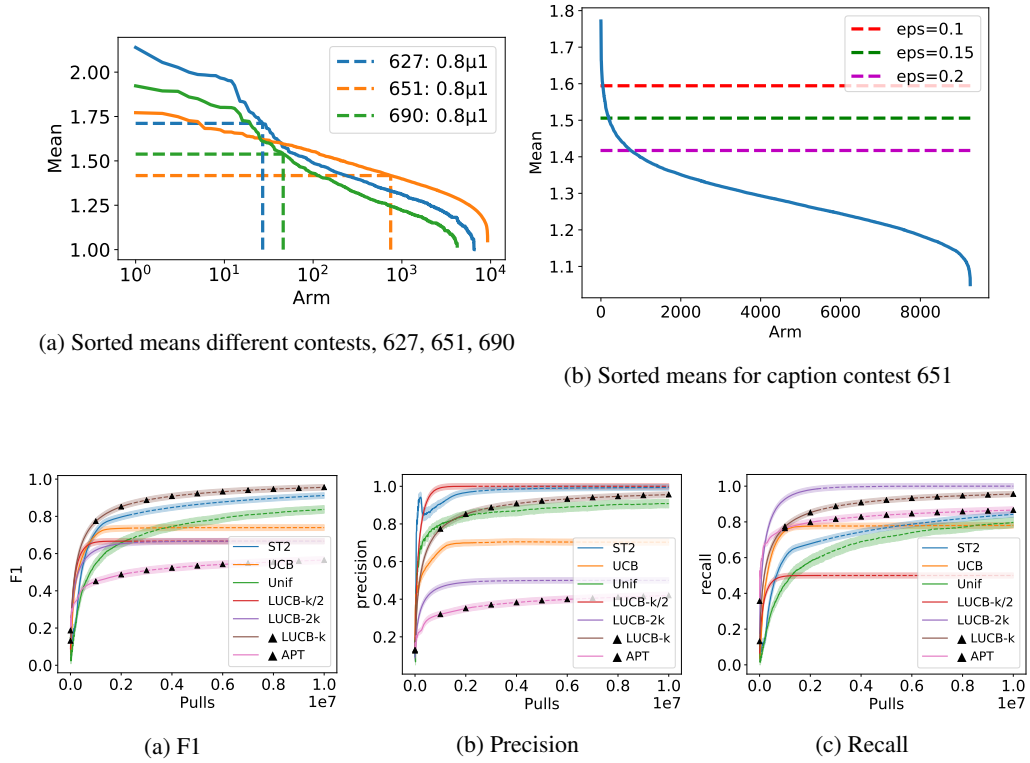


Figure 10: F1, Precision, and Recall scores on the New Yorker Caption Contest with $\epsilon = 0.2$




Figure 11: F1, Precision, and Recall scores on the New Yorker Caption Contest with $\epsilon = 0.15$




Figure 12: F1, Precision, and Recall scores on the New Yorker Caption Contest with $\epsilon = 0.1$

that compound is against the target, ACRVL1. A meta-analysis, done by [2], reported that these values have log-normal distributions with variance less than 1. Therefore, we compute the log of each percent control and may sample from a normal distribution with that mean and variance 1. As before, we plot F1, precision, and recall for all methods. To simulate being in a medical research regime where a higher level of precision is often desired, we take $\delta = 0.001$. We test each method on returning all **multiplicative** ϵ -good arms with $\epsilon = 0.8$ and plot the results in Figure 13. Note that these curves are plotted on a log-scale to emphasize the early regime of this experiment. It is likewise true here that the oracle baselines perform better on recall than they do on precision. $(ST)^2$ again performs well with respect to precision, and is more competitive with respect to recall in this experiment. Finally, $(ST)^2$ is competitive versus oracle methods on F1 score and greatly outperforms UCB and uniform sampling.

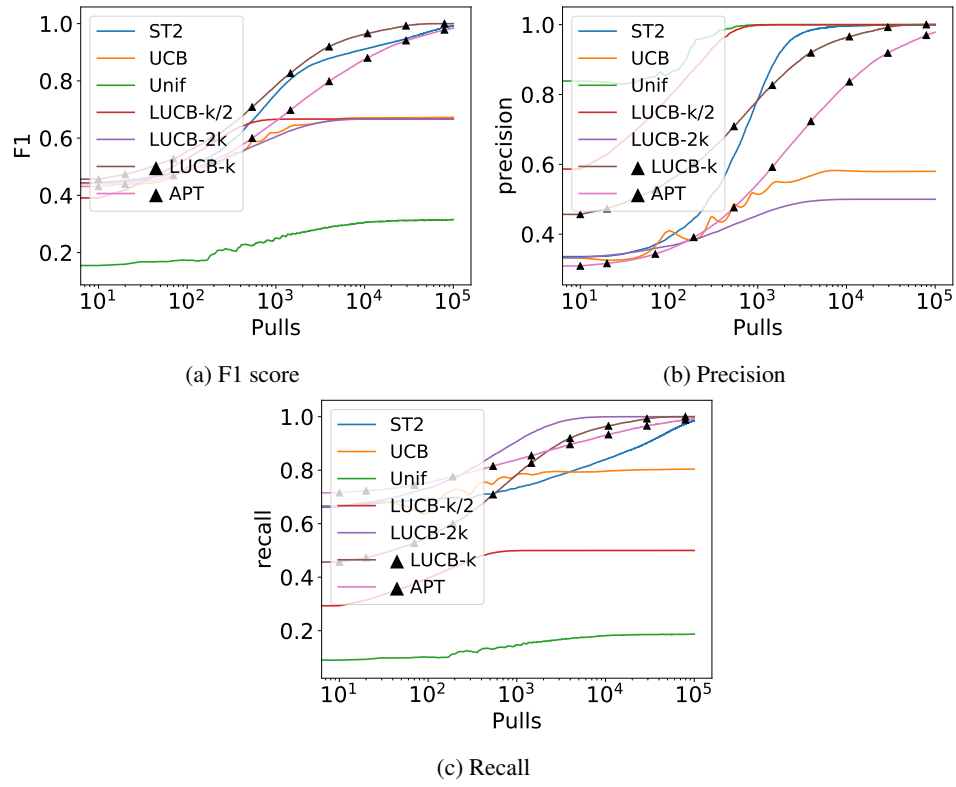


Figure 13: Precision and Recall curves for the PKIS2 cancer drug discovery experiment with $\epsilon = 0.8$

B (ST)², An optimism based algorithm for all- ϵ

Algorithm 2 The (ST)² Algorithm

Require: Instance $\nu, \epsilon > 0, \delta \in (0, 1/2], \gamma \geq 0$ ($\epsilon \in (0, 1/2]$, and $\gamma \in [0, \min(16/\mu_1, 1/2)]$)

- 1: Pull each arm once, initialize $T_i \leftarrow 1$, update $\hat{\mu}_i$ for each $i \in \{1, 2, \dots, n\}$
 - 2: Empirically good arms: $\hat{G} = \{i : \hat{\mu}_i \geq \max_j \hat{\mu}_j - \epsilon\}$ or $\hat{G} = \{i : \hat{\mu}_i \geq (1 - \epsilon) \max_j \hat{\mu}_j\}$
 - 3: $U_t = \max_j \hat{\mu}_j(T_j) + C_{\delta/n}(T_j) - \epsilon - \gamma$ or $U_t = (1 - \epsilon - \gamma) (\max_j \hat{\mu}_j(t) + C_{\delta/n}(T_j))$
 - 4: $L_t = \max_j \hat{\mu}_j(T_j) - C_{\delta/n}(T_j) - \epsilon$ or $L_t = (1 - \epsilon) (\max_j \hat{\mu}_j(t) - C_{\delta/n}(T_j))$
 - 5: Known arms: $K = \{i : \hat{\mu}_i(T_i) + C_{\delta/n}(T_i) < L_t \text{ or } \hat{\mu}_i(T_i) - C_{\delta/n}(T_i) > U_t\}$
 - 6: **while** $K \neq [n]$ **do**
 - 7: Pull arm $i_1(t) = \arg \min_{i \in \hat{G} \setminus K} \hat{\mu}_i(T_i) - C_{\delta/n}(T_i)$, update $T_{i_1}, \hat{\mu}_{i_1}$
 - 8: Pull arm $i_2(t) = \arg \max_{i \in \hat{G} \setminus K} \hat{\mu}_i(T_i) + C_{\delta/n}(T_i)$, update $T_{i_2}, \hat{\mu}_{i_2}$
 - 9: Pull arm $i^*(t) = \arg \max_i \hat{\mu}_i(T_i) + C_{\delta/n}(T_i)$, update $T_{i^*}, \hat{\mu}_{i^*}$
 - 10: Update bounds L_t, U_t , sets \hat{G}, K
 - return** The set of good arms $\{i : \hat{\mu}_i(T_i) - C_{\delta/n}(T_i) > U_t\}$
-

B.1 Optimism with additive γ

Theorem B.1. Fix $\epsilon \geq 0, 0 < \delta \leq 1/2, \gamma \in [0, 16]$ and an instance ν such that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i . In the case that $G_\epsilon = [n]$, let $\alpha_\epsilon = \min(\alpha_\epsilon, \beta_\epsilon)$. With probability at least $1 - \delta$, (ST)² correctly returns a set G such that $G_\epsilon \subset G \subset G_{\epsilon+\gamma}$ in at most

$$12 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\ \frac{4096}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right), \\ \left. \frac{4096}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \right\}, \\ \left. \frac{1}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta\gamma^2} \right) \right) \right\}$$

samples.

Proof. Throughout the proof, recall that $\Delta_i = \mu_1 - \mu_i$ for all i , $\alpha_\epsilon = \min_{i \in G_\epsilon} \mu_i - (\mu_1 - \epsilon)$, and $\beta_\epsilon = \min_{i \in G_\epsilon^c} (\mu_1 - \epsilon) - \mu_i$. Additionally, at any time t , we will take $T_j(t)$ to denote the number of samples of arm j up to time t .

Define the event

$$\mathcal{E} = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &= \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \\ &\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{n} = \delta \end{aligned}$$

Hence, $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$. Throughout, we will make use of a function $h(x, \delta)$ such that if $t \geq h(x, \delta)$, then $C_\delta(t) \leq |x|$. We bound $h(\cdot, \cdot)$ in Lemma F.2. $h(\cdot, \cdot)$ is assumed to decrease monotonically in both arguments and is symmetric in its first argument.

B.1.1 Step 0: Correctness

We begin by showing that on \mathcal{E} , if $(\text{ST})^2$ terminates, it returns a set G such that $G_\epsilon \subset G \subset G_{\epsilon+\gamma}$. Since $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$, this implies that $(\text{ST})^2$ is correct with high probability.

Claim 0: On Event \mathcal{E} , at all times t , $U_t \geq \mu_1 - \epsilon - \gamma$.

Proof.

$$\begin{aligned} U_t &= \max_j \hat{\mu}_j(T_j(t)) + C_{\delta/n}(T_j(t)) - \epsilon - \gamma \geq \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t)) - \epsilon - \gamma \\ &\stackrel{\mathcal{E}}{\geq} \mu_1 - \epsilon - \gamma \end{aligned}$$

□

Claim 1: On Event \mathcal{E} , at all times t , $L_t \leq \mu_1 - \epsilon$.

Proof.

$$L_t = \max_j \hat{\mu}_j(T_j(t)) - C_{\delta/n}(T_j(t)) - \epsilon \stackrel{\mathcal{E}}{\leq} \max_j \mu_j - \epsilon = \mu_1 - \epsilon$$

□

Claim 2: On event \mathcal{E} , if there is a time t such that $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t$, then $i \in G_{\epsilon+\gamma}$.

Proof. Assume for some t , $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t$. Then

$$\mu_i \stackrel{\mathcal{E}}{\geq} \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \geq U_t \stackrel{\text{Claim 0}}{\geq} \mu_1 - \epsilon - \gamma$$

which implies $i \in G_{\epsilon+\gamma}$

□

Claim 3: On event \mathcal{E} , if there is a time t such that $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t$, then $i \in G_\epsilon^c$.

Proof. Assume that is a t for which $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t$. Then

$$\mu_i \stackrel{\mathcal{E}}{\leq} \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \leq L_t \stackrel{\text{Claim 1}}{\leq} \mu_1 - \epsilon$$

which implies $i \in G_\epsilon^c$.

□

$(\text{ST})^2$ terminates at any time t such that simultaneously for all arms i , either $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t$ or $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) < L_t$. On \mathcal{E} , by Claim 3, $G_\epsilon \subset \{i : \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t\}$. On \mathcal{E} , by Claim 2, $\{i : \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t\} \subset G_{\epsilon+\gamma}$. Hence, on the event \mathcal{E} , $(\text{ST})^2$ returns a set G such that $G_\epsilon \subset G \subset G_{\epsilon+\gamma}$.

B.1.2 Step 1: Complexity of estimating the threshold, $\mu_1 - \epsilon$

Let STOP denote the termination event that for all arms i , either $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t$ or $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) < L_t$. Let ω denote the quantity

$$\omega := \max\{\gamma, \min(\alpha_\epsilon, \beta_\epsilon)\}.$$

Let T denote the random variable of the total number of rounds before $(\text{ST})^2$ terminates. At most 3 samples are drawn in any round. Hence, the total sample complexity is bounded by $3T$. We may write T as

$$T := |\{t : \neg \text{STOP}\}| = |\{t : \neg \text{STOP and } i^* \notin G_\omega\}| + |\{t : \neg \text{STOP and } i^* \in G_\omega\}|$$

Next, we bound the first event in this decomposition.

Claim 0: On \mathcal{E} ,

$$|\{t : \neg \text{STOP and } i^* \notin G_\omega\}| \leq \sum_{i \in G_\omega^c} \min \left\{ h\left(\frac{\gamma}{2}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{2}, \frac{\delta}{n}\right), h\left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{2}, \frac{\delta}{n}\right) \right] \right\}.$$

Proof. If for each $i \in G_\omega^c$, $\mu_i + 2C_{\delta/n}(T_i(t)) < \mu_1$ is true, which is ensured when $T_i(t) > h(\Delta_i/2, \frac{\delta}{n})$ for all $i \in G_\omega^c$, then

$$\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(T_i(t)) < \mu_1 \stackrel{\mathcal{E}}{\leq} \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t))$$

which implies that $i \neq i^*$. Additionally, since $i \in G_\omega^c$ by assumption, we have that $\mu_1 - \omega - \mu_i \geq 0$, which reduces to $\Delta_i \geq \omega$. Since $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$, it is likewise true that

$$h\left(\frac{\Delta_i}{2}, \frac{\delta}{n}\right) = \min\left[h\left(\frac{\gamma}{2}, \frac{\delta}{n}\right), \min\left\{h\left(\frac{\Delta_i}{2}, \frac{\delta}{n}\right), h\left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{2}, \frac{\delta}{n}\right)\right\}\right].$$

Summing over all $i \in G_\omega^c$ achieves the result. \square

We may decompose the set $\{t : \neg\text{STOP and } i^* \in G_\omega\}$ as

$$\begin{aligned} & \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16}\right\} \\ & \cup \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16}\right\} \end{aligned}$$

$$\begin{aligned} \text{Claim 1: } & \left| \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16}\right\} \right| \leq \\ & \sum_{i \in G_\omega} \min\left\{h\left(\frac{\gamma}{16}, \frac{\delta}{n}\right), \min\left[h\left(\frac{\Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n}\right)\right]\right\} \end{aligned}$$

Proof. $C_{\delta/n}(T_i(t)) \leq \frac{\omega}{16}$ is true when $T_i(t) \geq h\left(\frac{\omega}{16}, \frac{\delta}{n}\right)$. Since $i^* \in G_\omega$, $\mu_i - (\mu_1 - \omega) \geq 0$, which implies $\Delta_i \leq \omega$. By definition, $\omega = \min(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$. Hence, by monotonicity of $h(\cdot, \cdot)$,

$$\begin{aligned} h\left(\frac{\omega}{16}, \frac{\delta}{n}\right) &= \min\left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\omega}{16}, \frac{\delta}{n}\right)\right] \\ &= \min\left\{h\left(\frac{\gamma}{16}, \frac{\delta}{n}\right), \min\left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n}\right)\right]\right\} \end{aligned}$$

Summing over all $i \in G_\omega$ achieves the desired result. \square

B.1.3 Step 2: Controlling “crossing” events

Recall that we sample $i_1(t) \in \widehat{G}$ and $i_2(t) \in \widehat{G}^c$. In this section, we control the number of times that $i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}^c$ and $i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}$.

To do so, we first decompose the set $\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16}\}$ as

$$\begin{aligned} & \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}^c\right\} \\ & \cup \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}\right\} \end{aligned}$$

$$\begin{aligned} \text{Claim 0: } & \left| \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}^c\right\} \right| \leq \\ & \sum_{i \in G_{\epsilon+\frac{\gamma}{2}}^c} \min\left[h\left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right]. \end{aligned}$$

Proof. Recall that \widehat{G} is the set of all arms whose empirical means exceed $\max_i \hat{\mu}_i(T_i(t)) - \epsilon$, and $i_1(t) \in \widehat{G}$ by definition. Note that $\max_i \hat{\mu}_i(T_i(t)) - \epsilon > \max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) - \epsilon = L_t$. Hence, if an arm’s upper bound is below L_t , then the arm cannot be in \widehat{G} and thus not be $i_1(t)$. By the above event, $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16}$. Hence,

$$\mu_i^* + \frac{\omega}{8} \geq \mu_i^* + 2C_{\delta/n}(T_{i^*}(t)) \stackrel{\mathcal{E}}{\geq} \hat{\mu}_i^*(T_{i^*}(t)) + C_{\delta/n}(T_{i^*}(t)) \geq \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t)) \stackrel{\mathcal{E}}{\geq} \mu_1.$$

Therefore, $\mu_{i^*} \geq \mu_1 - \frac{\omega}{8}$ or equivalently, $i^* \in G_{\omega/8}$. Using this,

$$\begin{aligned} L_t = \max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) - \epsilon &\geq \hat{\mu}_{i^*}(T_{i^*}(t)) - C_{\delta/n}(T_{i^*}(t)) - \epsilon \\ &\stackrel{\mathcal{E}}{\geq} \mu_{i^*} - 2C_{\delta/n}(T_{i^*}(t)) - \epsilon \\ &\stackrel{\mathcal{E}}{\geq} \mu_{i^*} - \frac{\omega}{8} - \epsilon \\ &\geq \mu_1 - \frac{\omega}{4} - \epsilon \end{aligned}$$

Next, we bound the number of times an arm $i \in G_{\epsilon+\frac{\gamma}{2}}^c$ is sampled before its upper bound is below $\mu_1 - \frac{\omega}{4} - \epsilon$. Note that $C_{\delta/n}(T_i(t)) < \frac{1}{2}(\mu_1 - \frac{\omega}{4} - \epsilon - \mu_i)$, true when $T_i(t) > h(\frac{1}{2}(\mu_1 - \frac{\omega}{4} - \epsilon - \mu_i), \frac{\delta}{n})$ implies that

$$\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \stackrel{\epsilon}{\leq} \mu_i + 2C_{\delta/n}(T_i(t)) < \mu_1 - \frac{\omega}{4} - \epsilon \leq L_t.$$

Finally, we turn our attention to the difference $\mu_1 - \frac{\omega}{4} - \epsilon - \mu_i$. Recall that $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$.

$$\begin{aligned} \mu_1 - \frac{\omega}{4} - \epsilon - \mu_i &= (\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\omega \\ &= (\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon)). \end{aligned}$$

By definition, $\beta_\epsilon = \min_{i \in G_\epsilon^c}(\mu_1 - \epsilon) - \mu_i$. Hence, $\min(\alpha_\epsilon, \beta_\epsilon) \leq (\mu_1 - \epsilon) - \mu_i$ for all $i \in G_{\epsilon+\frac{\gamma}{2}}^c$. Similarly, since $i \in G_{\epsilon+\frac{\gamma}{2}}^c$ by assumption, $(\mu_1 - \epsilon - \frac{\gamma}{2}) - \mu_i \geq 0$, which rearranges to $\frac{\gamma}{2} \leq (\mu_1 - \epsilon) - \mu_i$. Therefore,

$$(\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon)) \geq \frac{1}{2}((\mu_1 - \epsilon) - \mu_i) = \frac{\Delta_i - \epsilon}{2}.$$

Hence, by monotonicity of $h(\cdot, \cdot)$,

$$h\left(\frac{1}{2}\left(\mu_1 - \frac{\omega}{4} - \epsilon - \mu_i\right), \frac{\delta}{n}\right) \leq h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right).$$

Lastly, as above, since $i \in G_{\epsilon+\frac{\gamma}{2}}^c$, we have that $\Delta_i - \epsilon = (\mu_1 - \epsilon) - \mu_i \geq \frac{1}{2}\gamma$. Hence,

$$h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right].$$

Putting this together, if $T_i(t) \geq \min\left[h\left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right]$, then $i \neq i_1(t)$ for all $i \in G_{\epsilon+\frac{\gamma}{2}}^c$. Summing over all such i bounds the size of set stated in the claim. \square

We decompose the remaining event

$$\left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}\right\}$$

as

$$\begin{aligned} &\left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}\right\} \\ &\cup \left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}^c\right\}. \end{aligned}$$

We proceed by bounding the size of the first set.

Claim 1:

$$\begin{aligned} &\left|\left\{t : \neg\text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}\right\}\right| \\ &\leq \sum_{i \in G_{\epsilon+\frac{\gamma}{2}}^c} \min\left[h\left(\frac{\epsilon\Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right] \end{aligned}$$

Proof. Recall that $K = \{i : \hat{\mu}(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t \text{ or } \hat{\mu}(T_i(t)) - C_{\delta/n}(T_i(t)) > L_t\}$ and i_2 is sampled from the set $\widehat{G}^c \setminus K$, ie all arms in \widehat{G}^c who have not been declared as above U_t or below L_t . Hence, if an arm's lower bound exceeds $U_t = \max_i \hat{\mu}(T_i(t)) + C_{\delta/n}(T_i(t)) - \epsilon - \gamma$, it must be in K and thus cannot be i_2 . Recall that $i^*(t) = \arg \max \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t))$. By the above event, $i^*(t) \in G_\omega$ and $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16}$. Hence,

$$U_t = \max_i \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) - \epsilon - \gamma = \hat{\mu}_{i^*(t)}(T_{i^*(t)}(t)) + C_{\delta/n}(T_{i^*(t)}(t)) - \epsilon - \gamma$$

$$\begin{aligned}
&\stackrel{\mathcal{E}}{\leq} \mu_{i^*(t)} + 2C_{\delta/n}(T_{i^*(t)}(t)) - \epsilon - \gamma \\
&\leq \mu_{i^*(t)} + \frac{\omega}{8} - \epsilon - \gamma \\
&\leq \mu_1 + \frac{\omega}{8} - \epsilon - \gamma
\end{aligned}$$

Next, we bound the number of times an arm $i \in G_{\epsilon+\frac{\gamma}{2}}$ is sampled before its lower bound is above $\mu_1 + \frac{\omega}{8} - \epsilon - \gamma$. Note that $C_{\delta/n}(T_i(t)) < \frac{1}{2}(\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma))$, true when $T_i(t) > h\left(\frac{1}{2}(\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma)), \frac{\delta}{n}\right)$ implies that

$$\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \stackrel{\mathcal{E}}{\geq} \mu_i - 2C_{\delta/n}(T_i(t)) > \mu_1 + \frac{\omega}{8} - \epsilon - \gamma.$$

Finally, we turn our attention to the difference $\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma)$. Recall that $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$.

$$\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right) = \mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8}\omega$$

Case 1a, $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$ and $i \in G_\epsilon$:

By definition, $\alpha_\epsilon = \min_{i \in G_\epsilon} \mu_i - (\mu_1 - \epsilon)$. Hence, $\min(\alpha_\epsilon, \beta_\epsilon) \leq \mu_i - (\mu_1 - \epsilon)$ for all $i \in G_\epsilon$. Therefore,

$$\begin{aligned}
\mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8}\omega &= \mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8}\min(\alpha_\epsilon, \beta_\epsilon) \\
&\geq \max\left(\mu_i - (\mu_1 - \epsilon) - \frac{1}{8}\min(\alpha_\epsilon, \beta_\epsilon), \gamma\right) \\
&\geq \max\left(\frac{7}{8}(\mu_i - (\mu_1 - \epsilon)), \gamma\right)
\end{aligned}$$

Case 1b, $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$ and $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$

Since $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$, if $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$, then $\frac{1}{2}\gamma < \min(\alpha_\epsilon, \beta_\epsilon)$. Since $\min(\alpha_\epsilon, \beta_\epsilon) = \min |\mu_i - (\mu_1 - \epsilon)|$, the set $G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$ is empty and there is nothing to prove.

Case 2a, $\omega = \gamma$ and $i \in G_\epsilon$:

$$\mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8}\omega = \mu_i - (\mu_1 - \epsilon) + \frac{7}{8}\gamma \geq \max\left(\mu_i - (\mu_1 - \epsilon), \frac{7}{8}\gamma\right)$$

Case 2b, $\omega = \gamma$ and $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$:

For $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$, we have that $\mu_i - (\mu_1 - \epsilon - \gamma/2) \geq 0$. Hence $\mu_i - (\mu_1 - \epsilon) \geq \frac{-\gamma}{2}$. Therefore,

$$\mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8}\omega \geq \frac{3}{8}\gamma = \max\left(\frac{3}{8}((\mu_1 - \epsilon) - \mu_i), \frac{3}{8}\gamma\right).$$

Applying the above cases and using monotonicity of $h(\cdot, \cdot)$, we see that for $i \in G_{\epsilon+\frac{\gamma}{2}}$,

$$h\left(\frac{1}{2}\left(\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon\right)\right), \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right].$$

Hence, if any $i \in G_{\epsilon+\frac{\gamma}{2}}$ has received this many samples, then its lower bound exceeds U_t and thus the arm must be in \hat{G} . Putting this together, if $T_i(t) \geq \min\left[h\left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right]$, then $i \neq i_2(t)$ for all $i \in G_{\epsilon+\frac{\gamma}{2}}$. Summing over all such i bounds the size of set stated in the claim. \square

B.1.4 Step 3: Controlling the complexity until stopping occurs

In this step, we turn our attention to the final event to control:

$$\mathcal{S} := \left\{ t : \neg \text{STOP and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon+\frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}^c \right\}.$$

For brevity, we will refer to this set as \mathcal{S} for this step. The objective will be to bound the time before each arms lower bound either clears U_t or its upper bound clears L_t which implies the stopping condition. To do so, we introduce, two events:

$$E_1(t) := \{\hat{\mu}_{i_1(t)}(T_{i_1(t)}(t)) - C_{\delta/n}(T_{i_1(t)}(t)) > U_t\} \quad (5)$$

and

$$E_2(t) := \{\hat{\mu}_{i_2(t)}(T_{i_2(t)}(t)) + C_{\delta/n}(T_{i_2(t)}(t)) < L_t\}. \quad (6)$$

If $E_1(t)$ is true, then $\hat{\mu}_i(T_i) - C_{\delta/n}(T_i(t)) > L_t$ for all $i \in \hat{G}$. If $E_2(t)$ is true, then $\hat{\mu}_i(T_i) + C_{\delta/n}(T_i(t)) < U_t$ for all $i \in \hat{G}^c$. Hence, by line 6 of (ST)², if both $E_1(t)$ and $E_2(t)$ are true, then (ST)² terminates.

Claim 0: $|\mathcal{S} \cap \{t : \neg E_1(t)\}| \leq \sum_{i \in G_{\epsilon+\frac{\gamma}{2}}} \min \left[h\left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right) \right].$

Proof. Recall that by the set \mathcal{S} , we have that $i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}$. Furthermore, by the set \mathcal{S} , we have that $i^*(t) \in G_\omega$ and $C_{\delta/n}(T_{i^*}(t)) \leq \omega/16$. Hence,

$$\begin{aligned} U_t &= \max_i \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) - \epsilon - \gamma \\ &= \hat{\mu}_{i^*(t)}(T_{i^*}(t)) + C_{\delta/n}(T_{i^*}(t)) - \epsilon - \gamma \\ &\stackrel{\mathcal{E}}{\leq} \mu_{i^*(t)} + 2C_{\delta/n}(T_{i^*}(t)) - \epsilon - \gamma \\ &\leq \mu_{i^*(t)} + \frac{\omega}{8} - \epsilon - \gamma \\ &\leq \mu_1 + \frac{\omega}{8} - \epsilon - \gamma \end{aligned}$$

If $C_{\delta/n}(T_i) \leq \frac{1}{2}(\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma))$ which is true when $T_i \geq h\left(\frac{1}{2}(\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma)), \frac{\delta}{n}\right)$, then

$$\hat{\mu}_i(T_i) - C_{\delta/n}(T_i) \geq \mu_i - 2C_{\delta/n}(T_i) \geq \mu_1 + \frac{\omega}{8} - \epsilon - \gamma \geq U_t.$$

The remainder of the proof of this claim focuses on controlling the difference: $\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma)$ in the case that $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$ and $\omega = \gamma$. Recall that $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$. Hence, if any possible $i \in G_{\epsilon+\frac{\gamma}{2}}$ has received sufficiently many samples, since $i_1(t) \in G_{\epsilon+\frac{\gamma}{2}}$, this implies $E_1(t)$.

Case 1a, $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$ and $i \in G_\epsilon$

We focus on the difference $\mu_i - (\mu_1 + \frac{\omega}{8} - \epsilon - \gamma)$.

$$\begin{aligned} \mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right) &= \mu_i - \left(\mu_1 + \frac{\min(\alpha_\epsilon, \beta_\epsilon)}{8} - \epsilon - \gamma\right) \\ &= \mu_i - (\mu_1 - \epsilon) + \gamma - \frac{1}{8} \min(\alpha_\epsilon, \beta_\epsilon) \\ &\stackrel{(\gamma \geq 0)}{\geq} \frac{1}{2}(\mu_i - (\mu_1 - \epsilon)) = \frac{\epsilon - \Delta_i}{2} \end{aligned}$$

where the final step follows since $\min(\alpha_\epsilon, \beta_\epsilon) \leq \alpha_\epsilon \leq \mu_i - (\mu_1 - \epsilon)$ by definition for all $i \in G_\epsilon$. Then by monotonicity of $h(\cdot, \cdot)$,

$$h\left(\frac{1}{2}(\mu_i - (\mu_1 - \epsilon)), \frac{\delta}{n}\right) \leq h\left(\frac{\epsilon - \Delta_i}{4}, \frac{\delta}{n}\right).$$

Lastly, in this setting, $\gamma \leq \min(\alpha_\epsilon, \beta_\epsilon) \leq \epsilon - \Delta_i$ since $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$. Hence, it is trivially true that

$$h\left(\frac{\epsilon - \Delta_i}{4}, \frac{\delta}{n}\right) = \min \left[h\left(\frac{\epsilon - \Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right]$$

Case 1b, $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$ and $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$

Since $\omega = \max(\gamma, \min(\alpha_\epsilon, \beta_\epsilon))$, if $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$, then $\frac{1}{2}\gamma < \min(\alpha_\epsilon, \beta_\epsilon)$. Since $\min(\alpha_\epsilon, \beta_\epsilon) = \min|\mu_i - (\mu_1 - \epsilon)|$, the set $G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$ is empty and there is nothing to prove.

Case 2a, $\omega = \gamma$ and $i \in G_\epsilon$

Again, we bound the difference $\mu_i - (\mu_1 + \frac{\omega}{4} - \epsilon - \gamma)$.

$$\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right) = \mu_i - (\mu_1 - \epsilon) + \frac{7}{8}\gamma$$

Since $i \in G_\epsilon$, $\mu_i - (\mu_1 - \epsilon) \geq 0$. Hence,

$$\begin{aligned} \mu_i - (\mu_1 - \epsilon) + \frac{7}{8}\gamma &\geq \max\left(\mu_i - (\mu_1 - \epsilon), \frac{7}{8}\gamma\right) \\ &\geq \frac{1}{2} \max(\epsilon - \Delta_i, \gamma) \end{aligned}$$

Therefore, we have that

$$h\left(\frac{1}{2}\left(\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right)\right), \frac{\delta}{n}\right) \leq h\left(\frac{\epsilon - \Delta_i}{4}, \frac{\delta}{n}\right)$$

and

$$h\left(\frac{1}{2}\left(\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right)\right), \frac{\delta}{n}\right) \leq h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right).$$

Hence,

$$h\left(\frac{1}{2}\left(\mu_i - \left(\mu_1 + \frac{\omega}{4} - \epsilon - \gamma\right)\right), \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\epsilon - \Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right)\right].$$

Case 2b, $\omega = \gamma$ and $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$

As before,

$$\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right) = \mu_i - (\mu_1 - \epsilon) + \frac{7}{8}\gamma$$

Since $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$, we have that $\mu_i - (\mu_1 - \epsilon - \frac{\gamma}{2}) \geq 0$. Rearranging implies that $\mu_i - (\mu_1 - \epsilon) \geq \frac{-1}{2}\gamma$. Hence,

$$\mu_i - (\mu_1 - \epsilon) + \frac{7}{8}\gamma \geq \frac{3}{8}\gamma.$$

Hence,

$$h\left(\frac{1}{2}\left(\mu_i - \left(\mu_1 + \frac{\omega}{8} - \epsilon - \gamma\right)\right), \frac{\delta}{n}\right) \leq h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right).$$

Additionally, as above, if $i \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$, we have that $\mu_i - (\mu_1 - \epsilon - \frac{\gamma}{2}) \geq 0$ which implies that $(\mu_1 - \epsilon) - \mu_i \leq \gamma$. Hence

$$h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right) = \min\left[h\left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right].$$

Therefore, if T_i exceeds the above, then $E_1(t)$ is true for an $i_1 \in G_\epsilon^c \cap G_{\epsilon+\frac{\gamma}{2}}$. Combining all cases, and noting that $h(x, \delta) \geq h(x/2, \delta) \forall x$, we see that for $i_1 \in G_{\epsilon+\frac{\gamma}{2}}$, if

$$T_{i_1(t)}(t) > \min\left[h\left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right],$$

Then $E_1(t)$ is true. Summing over all possible $i_1 \in G_{\epsilon+\frac{\gamma}{2}}$ proves the claim. \square

Claim 1: $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| \leq \sum_{i \in G_{\epsilon+\frac{\gamma}{2}}^c} \min\left[h\left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right].$

Proof. By the events in set \mathcal{S} , $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16}$. Hence,

$$\mu_i^* + \frac{\omega}{8} \geq \mu_i^* + 2C_{\delta/n}(T_{i^*}(t)) \stackrel{\mathcal{E}}{\geq} \hat{\mu}_i^*(T_{i^*}(t)) + C_{\delta/n}(T_{i^*}(t)) \geq \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t)) \stackrel{\mathcal{E}}{\geq} \mu_1.$$

Therefore, $\mu_{i^*} \geq \mu_1 - \frac{\omega}{8}$ or equivalently, $i^* \in G_{\omega/8}$. Using this,

$$\begin{aligned} L_t = \max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) - \epsilon &\geq \hat{\mu}_{i^*}(T_{i^*}(t)) - C_{\delta/n}(T_{i^*}(t)) - \epsilon \\ &\stackrel{\mathcal{E}}{\geq} \mu_{i^*} - 2C_{\delta/n}(T_{i^*}(t)) - \epsilon \\ &\stackrel{\mathcal{E}}{\geq} \mu_{i^*} - \frac{\omega}{8} - \epsilon \\ &\geq \mu_1 - \frac{\omega}{4} - \epsilon \end{aligned}$$

For $i \in G_{\epsilon+\frac{\gamma}{2}}^c$, if $C_{\delta/n}(T_i) \leq \frac{1}{2}((\mu_1 - \frac{\omega}{4} - \epsilon) - \mu_i)$, true when $T_i \geq h(\frac{1}{2}((\mu_1 - \frac{\omega}{4} - \epsilon) - \mu_i), \frac{\delta}{n})$, then

$$\hat{\mu}_i(T_i) + C_{\delta/n}(T_i) \leq \mu_i + 2C_{\delta/n}(T_i) \leq \mu_1 - \frac{\omega}{4} - \epsilon \leq L_t.$$

As before, we seek a lower bound for the difference $(\mu_1 - \frac{\omega}{4} - \epsilon) - \mu_i$.

Case 1: $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$

$$\begin{aligned} \left(\mu_1 - \frac{\omega}{4} - \epsilon\right) - \mu_i &= (\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\min(\alpha_\epsilon, \beta_\epsilon) \\ &\geq \frac{1}{2}((\mu_1 - \epsilon) - \mu_i) \end{aligned}$$

since $(\mu_1 - \epsilon) - \mu_i \geq \min(\alpha_\epsilon, \beta_\epsilon)$. Therefore, we have that

$$h\left(\frac{1}{2}\left((\mu_1 - \frac{\omega}{4} - \epsilon) - \mu_i\right), \frac{\delta}{n}\right) \leq h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right).$$

Lastly, in this setting, $\gamma \leq \min(\alpha_\epsilon, \beta_\epsilon) \leq \epsilon - \Delta_i$ since $\omega = \min(\alpha_\epsilon, \beta_\epsilon)$. Hence, it is trivially true that

$$h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right) = \min\left[h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right)\right].$$

Case 2: $\omega = \gamma$

Assume that $\gamma > \min(\alpha_\epsilon, \beta_\epsilon)$, as equality is covered by the previous case. Hence,

$$\left(\mu_1 - \frac{\omega}{4} - \epsilon\right) - \mu_i = (\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\gamma$$

Recall that we seek to control $i_2 \in G_{\epsilon+\frac{\gamma}{2}}^c$. For any $i \in G_{\epsilon+\frac{\gamma}{2}}^c$, we have that $\mu_1 - \epsilon - \frac{\gamma}{2} - \mu_i \geq 0$.

Rearranging, we see that $(\mu_1 - \epsilon) - \mu_i \geq \frac{1}{2}\gamma$ which implies that

$$(\mu_1 - \epsilon) - \mu_i - \frac{1}{4}\gamma \geq \frac{1}{2}((\mu_1 - \epsilon) - \mu_i).$$

Therefore, we have that

$$h\left(\frac{1}{2}\left((\mu_1 - \frac{\omega}{4} - \epsilon) - \mu_i\right), \frac{\delta}{n}\right) \leq h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right)$$

is this setting as well. Similarly, since $\Delta_i - \epsilon \geq \frac{1}{2}\gamma$, we likewise have that

$$h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma}{8}, \frac{\delta}{n}\right)\right].$$

Hence, if T_i exceeds the right-hand side of the preceding inequality, then for any $i \in G_{\epsilon+\frac{\gamma}{2}}^c$, its upper bound is below L_t . Hence for $i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}^c$, this implies event $E_2(t)$. Summing over all possible values of $i_2(t) \in G_{\epsilon+\frac{\gamma}{2}}^c$ proves the claim. \square

Claim 2: The cardinality of \mathcal{S} is bounded as $|\mathcal{S}| \leq \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right]$.

Proof. First, \mathcal{S} may be decomposed as

$$|\mathcal{S}| = |\mathcal{S} \cap \{t : \neg E_1(t)\}| + |\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| + |\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : E_2(t)\}|$$

Note that $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : E_2(t)\}| = 0$ because we have assumed in set \mathcal{S} that $(\text{ST})^2$ has not stopped, and $\{t : E_1(t)\} \cap \{t : E_2(t)\}$ implies termination. By Claim 0, $|\mathcal{S} \cap \{t : \neg E_1(t)\}| \leq \sum_{i \in G_{\epsilon + \frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right]$. By Claim 1, $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| \leq \sum_{i \in G_{\epsilon + \frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right]$. Recalling that h is assumed to be symmetric in its first argument proves the claim. \square

B.1.5 Step 4: Putting it all together

Recall that the total number of rounds T that $(\text{ST})^2$ runs for is given by $T = |\{t : \neg \text{STOP}\}|$. To bound this quantity, we have decomposed the set $\{t : \neg \text{STOP}\}$ into many subsets. Below, we show this decomposition.

$$\begin{aligned} \{t : \neg \text{STOP}\} = & \\ \{t : \neg \text{STOP} \text{ and } i^* \notin G_\omega\} & \\ \cup \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16} \right\} & \\ \cup \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}}^c \right\} & \\ \cup \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon + \frac{\gamma}{2}} \right\} & \\ \cup \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon + \frac{\gamma}{2}}^c \right\}. & \end{aligned}$$

Hence, by a union bound and plugging in the results of the above steps,

$$\begin{aligned} |\{t : \neg \text{STOP}\}| \leq & \\ |\{t : \neg \text{STOP} \text{ and } i^* \notin G_\omega\}| & \\ + \left| \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } \exists i \in G_\omega : C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16} \right\} \right| & \\ + \left| \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}}^c \right\} \right| & \\ + \left| \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon + \frac{\gamma}{2}} \right\} \right| & \\ + \left| \left\{ t : \neg \text{STOP} \text{ and } i^* \in G_\omega \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16} \text{ and } i_1(t) \in G_{\epsilon + \frac{\gamma}{2}} \text{ and } i_2(t) \in G_{\epsilon + \frac{\gamma}{2}}^c \right\} \right| & \\ \leq \sum_{i \in G_\omega^c} \min \left\{ h \left(\frac{\gamma}{2}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{2}, \frac{\delta}{n} \right), h \left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{2}, \frac{\delta}{n} \right) \right] \right\} & \\ + \sum_{i \in G_\omega} \min \left\{ h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n} \right) \right] \right\} & \\ + \sum_{i \in G_{\epsilon + \frac{\gamma}{2}}^c} \min \left[h \left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right] & \\ + \sum_{i \in G_{\epsilon + \frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right] & \\ + \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right] & \\ \stackrel{(\epsilon \leq 1/2)}{\leq} \sum_{i=1}^n \min \left\{ h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n} \right) \right] \right\} & \end{aligned}$$

$$\begin{aligned}
& + 2 \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma}{8}, \frac{\delta}{n} \right) \right] \\
& \leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n} \right) \right] \right\}, \right. \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\}
\end{aligned}$$

Next, by Lemma F.3, we may bound the minimum of $h(\cdot, \cdot)$ functions.

$$\begin{aligned}
& 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{16}, \frac{\delta}{n} \right) \right] \right\}, \right. \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\} \\
& = 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), \right. \right. \\
& \quad \left. \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), \max \left[h \left(\frac{\alpha_\epsilon}{16}, \frac{\delta}{n} \right), h \left(\frac{\beta_\epsilon}{16}, \frac{\delta}{n} \right) \right] \right] \right\}, \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\} \\
& \leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), \right. \right. \\
& \quad \left. \max \left[h \left(\frac{\Delta_i + \alpha_\epsilon}{32}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \beta_\epsilon}{32}, \frac{\delta}{n} \right) \right] \right\}, \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\} \\
& = 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \alpha_\epsilon}{32}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \beta_\epsilon}{32}, \frac{\delta}{n} \right) \right\}, \right. \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\}
\end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $|\epsilon - \Delta_i| \leq 8$ for all i and $\epsilon \leq 1/2$ implies that $\frac{1}{8}|\epsilon - \Delta_i| \leq 2$ and $\frac{1}{8}\min(\alpha_\epsilon, \beta_\epsilon) \leq 2$. $\Delta_i \leq 16$ for all i , gives $0.125\Delta_i \leq 2$. Lastly, $\gamma \leq 16$ implies that $\frac{\gamma}{8} \leq 2$. Therefore,

$$\begin{aligned}
& 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon}{16}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \alpha_\epsilon}{32}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \beta_\epsilon}{32}, \frac{\delta}{n} \right) \right\}, \right. \\
& \quad \left. h \left(\frac{\gamma}{16}, \frac{\delta}{n} \right) \right\} \\
& \leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{(\epsilon - \Delta_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta(\epsilon - \Delta_i)^2} \right) \right), \right. \right. \\
& \quad \frac{4096}{(\Delta_i + \alpha_\epsilon)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\Delta_i + \alpha_\epsilon)^2} \right) \right), \\
& \quad \frac{4096}{(\Delta_i + \beta_\epsilon)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\Delta_i + \beta_\epsilon)^2} \right) \right) \left. \right\}, \\
& \quad \frac{1}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta\gamma^2} \right) \right) \left. \right\}
\end{aligned}$$

$$\begin{aligned}
&= 4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\
&\quad \frac{4096}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right), \\
&\quad \left. \frac{4096}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \right\}, \\
&\quad \left. \frac{1}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta\gamma^2} \right) \right) \right\}.
\end{aligned}$$

The above bounds the number of rounds T . Therefore, the total number of samples is at most $3T$. \square

B.2 Optimism with multiplicative γ

Theorem B.2. Fix $\epsilon \in (0, 1/2]$, $0 < \delta \leq 1/2$, $\gamma \in [0, \min(16/\mu_1, 1/2)]$ and an instance ν such that $\max(\Delta_i, |\epsilon\mu_1 - \Delta_i|) \leq 8$ for all i . In the case that $M_\epsilon = [n]$, let $\tilde{\alpha}_\epsilon = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. With probability at least $1 - \delta$, $(\text{ST})^2$ correctly returns a set G such that $M_\epsilon \subset G \subset M_{\epsilon+\gamma}$ in at most

$$\begin{aligned}
&12 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right. \\
&\quad \frac{4096}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right), \\
&\quad \left. \frac{4096}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon})^2} \right) \right) \right\}, \\
&\quad \left. \frac{1024}{\gamma^2 \mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta\gamma^2 \mu_1^2} \right) \right) \right\}
\end{aligned}$$

samples.

Proof. Throughout the proof, recall that $\Delta_i = \mu_1 - \mu_i$ for all i , $\tilde{\alpha}_\epsilon = \min_{i \in M_\epsilon} \mu_i - (1-\epsilon)\mu_1$, and $\tilde{\beta}_\epsilon = \min_{i \in M_\epsilon^c} (1-\epsilon)\mu_1 - \mu_i$. Additionally, at any time t , we will take $T_j(t)$ to denote the number of samples of arm j up to time t .

Define the event

$$\mathcal{E} = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\begin{aligned}
\mathbb{P}(\mathcal{E}^c) &= \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \\
&\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{n} = \delta
\end{aligned}$$

Hence, $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$. Throughout, we will make use of a function $h(x, \delta)$ such that if $t \geq h(x, \delta)$, then $C_\delta(t) \leq |x|$. We bound $h(\cdot, \cdot)$ in Lemma F.2. $h(\cdot, \cdot)$ is assumed to decrease monotonically in both arguments and is symmetric in its first argument.

B.2.1 Step 0: Correctness

We begin by showing that on \mathcal{E} , if $(\text{ST})^2$ terminates, it returns a set G such that $M_\epsilon \subset G \subset M_{(\epsilon+\gamma)}$. Since $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$, this implies that $(\text{ST})^2$ is correct with high probability.

Claim 0: On Event \mathcal{E} , at all times t , $U_t \geq (1 - \epsilon - \gamma)\mu_1$.

Proof.

$$\begin{aligned} U_t &= (1 - \epsilon - \gamma)(\max_j \hat{\mu}_j(T_j(t)) + C_{\delta/n}(T_j(t))) \geq (1 - \epsilon - \gamma)(\hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t))) \\ &\stackrel{\mathcal{E}}{\geq} (1 - \epsilon - \gamma)\mu_1 \end{aligned}$$

□

Claim 1: On Event \mathcal{E} , at all times t , $L_t \leq (1 - \epsilon)\mu_1$.

Proof.

$$L_t = (1 - \epsilon) \left(\max_j \hat{\mu}_j(T_j(t)) - C_{\delta/n}(T_j(t)) \right) \stackrel{\mathcal{E}}{\leq} (1 - \epsilon) \max_j \mu_j = (1 - \epsilon)\mu_1$$

□

Claim 2: On event \mathcal{E} , if there is a time t such that $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t$, then $i \in M_{\epsilon+\gamma}$.

Proof. Assume for some t , $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t$. Then

$$\mu_i \stackrel{\mathcal{E}}{\geq} \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \geq U_t \stackrel{\text{Claim 0}}{\geq} (1 - \epsilon - \gamma)\mu_1$$

which implies $i \in M_{\epsilon+\gamma}$

□

Claim 3: On event \mathcal{E} , if there is a time t such that $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t$, then $i \in M_\epsilon^c$.

Proof. Assume that is a t for which $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t$. Then

$$\mu_i \stackrel{\mathcal{E}}{\leq} \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \leq L_t \stackrel{\text{Claim 1}}{\leq} (1 - \epsilon)\mu_1$$

which implies $i \in M_\epsilon^c$.

□

(ST)² terminates at any time t such that simultaneously for all arms i , either $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t$ or $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) < L_t$. On \mathcal{E} , by Claim 3, $M_\epsilon \subset \{i : \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t\}$. On \mathcal{E} , by Claim 2, $\{i : \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t\} \subset M_{\epsilon+\gamma}$. Hence, on the event \mathcal{E} , (ST)² returns a set G such that $M_\epsilon \subset G \subset M_{\epsilon+\gamma}$.

B.2.2 Step 1: Complexity of estimating the threshold, $(1 - \epsilon)\mu_1$

Let STOP denote the termination event that for all arms i , either $\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) > U_t$ or $\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) < L_t$. Let ω denote the quantity

$$\omega := \max\{\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)\}.$$

Let T denote the random variable of the total number of rounds before (ST)² terminates. At most 3 samples are drawn in any round. Hence, the total sample complexity is bounded by $3T$. We may write T as

$$T \equiv |\{t : \neg \text{STOP}\}| = |\{t : \neg \text{STOP and } i^* \notin M_{\omega/\mu_1}\}| + |\{t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1}\}|$$

Next, we bound the first event in this decomposition.

Claim 0: On \mathcal{E} , $|\{t : \neg \text{STOP and } i^* \notin M_{\omega/\mu_1}\}| \leq \sum_{i \in M_{\omega/\mu_1}^c} \min \left\{ h \left(\frac{\gamma\mu_1}{2}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{2}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{2}, \frac{\delta}{n} \right) \right] \right\}.$

Proof. For each $i \in M_{\omega/\mu_1}^c$, $\mu_i + 2C_{\delta/n}(T_i(t)) < \mu_1$, true when $T_i(t) > h(\Delta_i/2, \frac{\delta}{n})$ implies that

$$\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(T_i(t)) < \mu_1 \stackrel{\mathcal{E}}{\leq} \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t))$$

which implies that $i \neq i^*$. Additionally, since $i \in M_{\omega/\mu_1}^c$ by assumption, we have that $(1 - \omega/\mu_1)\mu_1 - \mu_i \geq 0$, which reduces to $\Delta_i \geq \omega$. Since $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$, it is likewise true that

$$h \left(\frac{\Delta_i}{2}, \frac{\delta}{n} \right) = \min \left[h \left(\frac{\gamma\mu_1}{2}, \frac{\delta}{n} \right), \min \left\{ h \left(\frac{\Delta_i}{2}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{2}, \frac{\delta}{n} \right) \right\} \right].$$

Summing over all $i \in M_{\omega/\mu_1}^c$ achieves the result. \square

We may decompose the event $\{t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1}\}$ as

$$\begin{aligned} & \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } \exists i \in M_{\omega/\mu_1} : C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16(1-\epsilon)} \right\} \\ & \cup \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \right\} \end{aligned}$$

Claim 1: $\left| \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \geq \frac{\omega}{16(1-\epsilon)} \right\} \right| \leq \sum_{i \in M_{\omega/\mu_1}} \min \left\{ h\left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \right\}$

Proof. $C_{\delta/n}(T_i(t)) \leq \frac{\omega}{16(1-\epsilon)}$ is true when $T_i(t) \geq h\left(\frac{\omega}{16(1-\epsilon)}, \frac{\delta}{n}\right)$. Since $i^* \in M_{\omega/\mu_1}$, $\mu_i - (1 - \omega/\mu_1)\mu_1 \geq 0$, which implies $\Delta_i \leq \omega$. By definition, $\omega = \min(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$. Hence, by monotonicity of $h(\cdot, \cdot)$,

$$\begin{aligned} h\left(\frac{\omega}{16(1-\epsilon)}, \frac{\delta}{n}\right) &= \min \left[h\left(\frac{\Delta_i}{16(1-\epsilon)}, \frac{\delta}{n}\right), h\left(\frac{\omega}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \\ &= \min \left\{ h\left(\frac{\gamma\mu_1}{16(1-\epsilon)}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{16(1-\epsilon)}, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \right\} \\ &\leq \min \left\{ h\left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \right\}. \end{aligned}$$

Summing over all $i \in M_{\omega/\mu_1}$ achieves the desired result. \square

B.2.3 Step 2: Controlling “crossing” events

Recall that we sample $i_1(t) \in \hat{G}$ and $i_2(t) \in \hat{G}^c$. In this section, we control the number of times that $i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}^c$ and $i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}$.

To do so, we first decompose the set $\{t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)}\}$ as

$$\begin{aligned} & \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\} \\ & \cup \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\} \end{aligned}$$

Claim 0: $\left| \left\{ t : \neg\text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\} \right| \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}^c} \min \left[h\left(\frac{\Delta_i - \epsilon\mu_1}{16}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n}\right) \right]$.

Proof. Recall that \hat{G} is the set of all arms whose empirical means exceed $(1 - \epsilon) \max_i \hat{\mu}_i(T_i(t))$, and $i_1(t) \in \hat{G}$ by definition. Note that $(1 - \epsilon) \max_i \hat{\mu}_i(T_i(t)) > (1 - \epsilon) (\max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t))) = L_t$. Hence, if an arm’s upper bound is below L_t , then the arm cannot be in \hat{G} and thus not be $i_1(t)$. By the above event, $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)}$. Therefore,

$$\begin{aligned} \mu_{i^*} + \frac{\omega}{8(1-\epsilon)} &\geq \mu_{i^*} + 2C_{\delta/n}(T_{i^*}(t)) \stackrel{\mathcal{E}}{\geq} \hat{\mu}_{i^*}(T_{i^*}(t)) + C_{\delta/n}(T_{i^*}(t)) \geq \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t)) \\ &\stackrel{\mathcal{E}}{\geq} \mu_1. \end{aligned}$$

Hence, $\mu_{i^*} \geq \mu_1 - \frac{\omega}{8(1-\epsilon)}$. Rearranging this, we see that $\mu_{i^*} - \left(1 - \frac{\omega}{8\mu_1(1-\epsilon)}\right) \mu_1 \geq 0$ which implies that $i^* \in M_{\frac{\omega}{8\mu_1(1-\epsilon)}}$. Hence,

$$L_t = (1 - \epsilon) \left(\max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \right) (1 - \epsilon) (\hat{\mu}_{i^*}(T_{i^*}(t)) - C_{\delta/n}(T_{i^*}(t)))$$

$$\begin{aligned}
&\stackrel{\mathcal{E}}{\geq} (1 - \epsilon) (\mu_{i^*} - 2C_{\delta/n}(T_{i^*}(t))) \\
&\geq (1 - \epsilon) \left(\mu_{i^*} - \frac{\omega}{8(1 - \epsilon)} \right) \\
&\geq (1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right)
\end{aligned}$$

Next, we bound the number of times an arm $i \in M_{\epsilon+\frac{\gamma}{2}}^c$ is sampled before its upper bound is below $(1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right)$. Note that $C_{\delta/n}(T_i(t)) < \frac{1}{2} \left((1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) - \mu_i \right)$, true when $T_i(t) > h \left(\frac{1}{2} \left((1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) - \mu_i \right), \frac{\delta}{n} \right)$ implies that

$$\hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(T_i(t)) < (1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) \leq L_t.$$

Finally, we turn our attention to the difference $(1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) - \mu_i$. Recall that $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$.

$$\begin{aligned}
(1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) - \mu_i &= (1 - \epsilon)\mu_1 - \mu_i - \frac{1}{4}\omega \\
&= (1 - \epsilon)\mu_1 - \mu_i - \frac{1}{4} \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)).
\end{aligned}$$

By definition, $\tilde{\beta}_\epsilon = \min_{i \in M_\epsilon^c} (1 - \epsilon)\mu_1 - \mu_i$. Hence, $\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq (1 - \epsilon)\mu_1 - \mu_i$ for all $i \in M_{\epsilon+\frac{\gamma}{2}}^c$. Similarly, since $i \in M_{\epsilon+\frac{\gamma}{2}}^c$ by assumption, $(1 - \epsilon - \frac{\gamma}{2})\mu_1 - \mu_i \geq 0$, which rearranges to $\frac{\gamma\mu_1}{2} \leq (1 - \epsilon)\mu_1 - \mu_i$. Therefore,

$$(1 - \epsilon)\mu_1 - \mu_i - \frac{1}{4} \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)) \geq \frac{1}{2} ((1 - \epsilon)\mu_1 - \mu_i) = \frac{\Delta_i - \epsilon\mu_1}{2}.$$

Hence, by monotonicity of $h(\cdot, \cdot)$,

$$h \left(\frac{1}{2} \left((1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1 - \epsilon)} \right) - \mu_i \right), \frac{\delta}{n} \right) \leq h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right).$$

Lastly, as above, since $i \in M_{\epsilon+\frac{\gamma}{2}}^c$, we have that $\Delta_i - \epsilon\mu_1 = (1 - \epsilon)\mu_1 - \mu_i \geq \frac{1}{2}\gamma\mu_1$. Hence,

$$h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right) \leq \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right].$$

Putting this together, if $T_i(t) \geq \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right]$, then $i \neq i_1(t)$ for all $i \in M_{\epsilon+\frac{\gamma}{2}}^c$. Summing over all such i bounds the size of set stated in the claim. \square

We decompose the remaining event

$$\left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1 - \epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\}$$

as

$$\begin{aligned}
&\left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1 - \epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \\
&\quad \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\} \\
&\cup \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1 - \epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \\
&\quad \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\}.
\end{aligned}$$

We proceed by bounding the cardinality of the first set.

Claim 1:

$$\left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \\ \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\} \\ \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right]$$

Proof. Recall that $K = \{i : \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) < L_t \text{ or } \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) > U_t\}$ is the set of known arms and i_2 is sampled from $\hat{G}^c \setminus K$. Hence, if an arm's lower bound exceeds U_t , it must be in K and therefore cannot be i_2 . Recall that $i^*(t) = \arg \max \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t))$. By the above event, $i^*(t) \in M_{\omega/\mu_1}$ and $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)}$. Hence,

$$\begin{aligned} U_t &= (1 - \epsilon - \gamma) \left(\max_i \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \right) \\ &= (1 - \epsilon - \gamma) \left(\hat{\mu}_{i^*(t)}(T_{i^*(t)}(t)) + C_{\delta/n}(T_{i^*(t)}(t)) \right) \\ &\stackrel{\mathcal{E}}{\leq} (1 - \epsilon - \gamma) \left(\mu_{i^*(t)} + 2C_{\delta/n}(T_{i^*(t)}(t)) \right) \\ &\leq (1 - \epsilon - \gamma) \left(\mu_{i^*(t)} + \frac{\omega}{8(1-\epsilon)} \right) \\ &\leq (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \end{aligned}$$

Next, we bound the number of times an arm $i \in M_{\epsilon+\frac{\gamma}{2}}$ is sampled before its lower bound is above $(1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right)$. Note that $C_{\delta/n}(T_i(t)) < \frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \right)$, true when $T_i(t) > h \left(\frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \right), \frac{\delta}{n} \right)$ implies that

$$\hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \stackrel{\mathcal{E}}{\geq} \mu_i - 2C_{\delta/n}(T_i(t)) > (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \geq U_t.$$

Finally, we turn our attention to the difference $\mu_i - (1 - \epsilon) \left(\mu_1 + \frac{\omega}{8} \right)$. Recall that $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$. Additionally, recall $\epsilon + \gamma \leq 1$.

$$\begin{aligned} \mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) &= \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \left(\frac{1 - \epsilon - \gamma}{1 - \epsilon} \right) \omega \\ &\geq \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \omega \end{aligned}$$

Case 1a, $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$ and $i \in M_\epsilon$:

By definition, $\tilde{\alpha}_\epsilon = \min_{i \in M_\epsilon} \mu_i - (1 - \epsilon)\mu_1$. Hence, $\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq \mu_i - (1 - \epsilon)\mu_1$ for all $i \in M_\epsilon$. Therefore,

$$\begin{aligned} \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \omega &= \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \\ &\geq \max \left(\mu_i - (1 - \epsilon)\mu_1 - \frac{1}{8} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon), \gamma\mu_1 \right) \\ &\geq \max \left(\frac{7}{8} (\mu_i - (1 - \epsilon)\mu_1), \gamma\mu_1 \right) \end{aligned}$$

Case 1b, $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$ and $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$

Since $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$, if $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$, then $\frac{1}{2}\gamma\mu_1 < \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. Since $\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) = \min |\mu_i - (1 - \epsilon)\mu_1|$, the set $M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$ is empty and there is nothing to prove.

Case 2a, $\omega = \gamma\mu_1$ and $i \in M_\epsilon$

$$\mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8}\omega = \mu_i - (1 - \epsilon)\mu_1 + \frac{7}{8}\gamma\mu_1 \geq \max\left(\mu_i - (1 - \epsilon)\mu_1, \frac{7}{8}\gamma\mu_1\right).$$

Case 2b, $\omega = \gamma\mu_1$ and $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$

For $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$, $\mu_i - (1 - \epsilon - \frac{\gamma}{2})\mu_1 \geq 0$. Hence, $\mu_i - (1 - \epsilon)\mu_1 \geq \frac{-\gamma\mu_1}{2}$. Therefore,

$$\mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8}\omega = \mu_i - (1 - \epsilon)\mu_1 + \frac{7}{8}\gamma\mu_1 \geq \frac{3}{8}\gamma\mu_1 \geq \max\left(\frac{1}{4}\gamma\mu_1, \frac{(1 - \epsilon)\mu_1 - \mu_i}{4}\right).$$

Combining all cases, by monotonicity of $h(\cdot, \cdot)$ and symmetry in its first argument, we see that

$$h\left(\frac{1}{2}\left(\mu_i - (1 - \epsilon - \gamma)\left(\mu_1 + \frac{\omega}{8(1 - \epsilon)}\right)\right), \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n}\right), h\left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n}\right)\right].$$

Putting this together, if $T_i(t) \geq \min\left[h\left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n}\right)\right]$, then $i \neq i_2(t)$ for all $i \in M_{\epsilon+\frac{\gamma}{2}}$. Summing over all such i bounds the size of set stated in the claim. \square

B.2.4 Step 3: Controlling the complexity until stopping occurs

In this step, we turn our attention to the final event to control:

$$\mathcal{S} := \left\{t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1 - \epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \text{ and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c\right\}. \quad (7)$$

For brevity, we will refer to this set as \mathcal{S} for this step. The objective will be to bound the time before each arms lower bound either clears U_t or its upper bound clears L_t which implies the stopping condition. To do so, we introduce, two events:

$$E_1(t) := \{\hat{\mu}_{i_1(t)}(T_{i_1(t)}(t)) - C_{\delta/n}(T_{i_1(t)}(t)) > U_t\} \quad (8)$$

and

$$E_2(t) := \{\hat{\mu}_{i_2(t)}(T_{i_2(t)}(t)) + C_{\delta/n}(T_{i_2(t)}(t)) < L_t\}. \quad (9)$$

If $E_1(t)$ is true, then $\hat{\mu}_i(T_i) - C_{\delta/n}(T_i(t)) > L_t$ for all $i \in \hat{G}$. If $E_2(t)$ is true, then $\hat{\mu}_i(T_i) + C_{\delta/n}(T_i(t)) < U_t$ for all $i \in \hat{G}^c$. Hence, by line 6 of (ST)², if both $E_1(t)$ and $E_2(t)$ are true, then (ST)² terminates.

Claim 0: $|\mathcal{S} \cap \{t : \neg E_1(t)\}| \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}} \min\left[h\left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n}\right)\right]$.

Proof. Recall that by the set \mathcal{S} , we have that $i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}$. Furthermore, by the set \mathcal{S} , we have that $i^*(t) \in M_{\omega/\mu_1}$ and $C_{\delta/n}(T_{i^*}(t)) \leq \omega/16(1 - \epsilon)$. Hence,

$$\begin{aligned} U_t &= (1 - \epsilon - \gamma) \left(\max_i \hat{\mu}_i(T_i(t)) + C_{\delta/n}(T_i(t)) \right) \\ &= (1 - \epsilon - \gamma) \left(\hat{\mu}_{i^*(t)}(T_{i^*(t)}(t)) + C_{\delta/n}(T_{i^*(t)}(t)) \right) \\ &\stackrel{\mathcal{E}}{\leq} (1 - \epsilon - \gamma) \left(\mu_{i^*(t)} + 2C_{\delta/n}(T_{i^*(t)}(t)) \right) \\ &\leq (1 - \epsilon - \gamma) \left(\mu_{i^*(t)} + \frac{\omega}{8(1 - \epsilon)} \right) \\ &\leq (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1 - \epsilon)} \right) \end{aligned}$$

If $C_{\delta/n}(T_i) \leq \frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1 - \epsilon)} \right) \right)$, true when $T_i \geq h\left(\frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1 - \epsilon)} \right) \right), \frac{\delta}{n}\right)$, then

$$\hat{\mu}_i(T_i) - C_{\delta/n}(T_i) \geq \mu_i - 2C_{\delta/n}(T_i) \geq (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1 - \epsilon)} \right) \geq U_t.$$

The remainder of the proof of this claim focuses on controlling the difference: $\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right)$ in the case that $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$ and $\omega = \gamma\mu_1$. Recall that $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$. Hence, if any possible $i \in M_{\epsilon+\frac{\gamma}{2}}$ has received sufficiently many samples, since $i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}$, this implies $E_1(t)$.

Case 1a, $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$ and $i \in M_\epsilon$

We focus on the difference $\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right)$. Recall that $\epsilon + \gamma \leq 1$.

$$\begin{aligned} \mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) &= \mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{8(1-\epsilon)} \right) \\ &= \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \left(\frac{1 - \epsilon - \gamma}{1 - \epsilon} \right) \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \\ &\stackrel{\gamma \geq 0 \text{ and } \epsilon + \gamma \leq 1}{\geq} \mu_i - (1 - \epsilon)\mu_1 - \frac{1}{8} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \\ &\geq \frac{1}{2}(\mu_i - (1 - \epsilon)\mu_1) = \frac{\epsilon\mu_1 - \Delta_i}{2} \end{aligned}$$

where the final step follows since $\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq \tilde{\alpha}_\epsilon \leq \mu_i - (1 - \epsilon)\mu_1$ by definition for all $i \in M_\epsilon$. Then by monotonicity of $h(\cdot, \cdot)$,

$$h \left(\frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \right), \frac{\delta}{n} \right) \leq h \left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n} \right).$$

Lastly, in this setting, $\gamma\mu_1 \leq \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq \epsilon\mu_1 - \Delta_i$ since $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. Hence, it is trivially true that

$$h \left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n} \right) = \min \left[h \left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n} \right) \right]$$

Case 1b, $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$ and $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$

Since $\omega = \max(\gamma\mu_1, \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon))$, if $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$, then $\frac{1}{2}\gamma\mu_1 < \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. Since $\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) = \min |\mu_i - (1 - \epsilon)\mu_1|$, the set $M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$ is empty and there is nothing to prove.

Case 2a, $\omega = \gamma\mu_1$ and $i \in M_\epsilon$

Next, we bound the difference $\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{4(1-\epsilon)} \right)$.

$$\begin{aligned} \mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) &= \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8} \left(\frac{1 - \epsilon - \gamma}{1 - \epsilon} \right) \gamma\mu_1 \\ &\geq \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 \left(1 - \frac{1}{8} \left(\frac{1 - \epsilon - \gamma}{1 - \epsilon} \right) \right) \end{aligned}$$

Since $i \in M_\epsilon$, $\mu_i - (1 - \epsilon)\mu_1 \geq 0$. Using this and the fact that $\epsilon, \gamma \geq 0$ and $\epsilon + \gamma \leq 1$,

$$\begin{aligned} \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 \left(1 - \frac{1}{8} \left(\frac{1 - \epsilon - \gamma}{1 - \epsilon} \right) \right) &\geq \mu_i - (1 - \epsilon)\mu_1 + \frac{7}{8}\gamma\mu_1 \\ &\geq \max \left(\mu_i - (1 - \epsilon)\mu_1, \frac{7}{8}\gamma\mu_1 \right) \\ &\geq \frac{1}{2} \max(\epsilon\mu_1 - \Delta_i, \gamma\mu_1) \end{aligned}$$

Therefore, we have that

$$h \left(\frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \right), \frac{\delta}{n} \right) \leq h \left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n} \right)$$

and

$$h \left(\frac{1}{2} \left(\mu_i - (1 - \epsilon - \gamma) \left(\mu_1 + \frac{\omega}{8(1-\epsilon)} \right) \right), \frac{\delta}{n} \right) \leq h \left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n} \right).$$

Hence,

$$h\left(\frac{1}{2}\left(\mu_i - (1 - \epsilon - \gamma)\left(\mu_1 + \frac{\omega}{8(1 - \epsilon)}\right)\right), \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n}\right)\right].$$

Case 2b, $\omega = \gamma\mu_1$ and $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$

As before,

$$\mu_i - (1 - \epsilon - \gamma)\left(\mu_1 + \frac{\omega}{8(1 - \epsilon)}\right) = \mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8}\left(\frac{1 - \epsilon - \gamma}{1 - \epsilon}\right)\gamma\mu_1$$

Since $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$, we have that $\mu_i - (1 - \epsilon - \frac{\gamma}{2})\mu_1 \geq 0$. Rearranging implies that $\mu_i - (1 - \epsilon)\mu_1 \geq \frac{-1}{2}\gamma\mu_1$. Hence,

$$\mu_i - (1 - \epsilon)\mu_1 + \gamma\mu_1 - \frac{1}{8}\left(\frac{1 - \epsilon - \gamma}{1 - \epsilon}\right)\gamma\mu_1 \geq \frac{1}{2}\gamma\mu_1 - \frac{1}{8}\left(\frac{1 - \epsilon - \gamma}{1 - \epsilon}\right)\gamma\mu_1 \geq \frac{3}{8}\gamma\mu_1.$$

Hence,

$$h\left(\frac{1}{2}\left(\mu_i - (1 - \epsilon - \gamma)\left(\mu_1 + \frac{\omega}{8(1 - \epsilon)}\right)\right), \frac{\delta}{n}\right) \leq h\left(\frac{3\gamma\mu_1}{8}, \frac{\delta}{n}\right).$$

Additionally, as above, if $i \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$, we have that $\mu_i - (1 - \epsilon - \frac{\gamma}{2})\mu_1 \geq 0$ which implies that $(1 - \epsilon)\mu_1 - \mu_i \leq \frac{1}{2}\gamma\mu_1$. Hence

$$h\left(\frac{3\gamma\mu_1}{8}, \frac{\delta}{n}\right) \leq \min\left[h\left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n}\right)\right].$$

Therefore, if T_i exceeds the above, then $E_1(t)$ is true for an $i_1 \in M_\epsilon^c \cap M_{\epsilon+\frac{\gamma}{2}}$. Combining all cases, we see that for $i_1 \in M_{\epsilon+\frac{\gamma}{2}}$, if

$$T_{i_1(t)}(t) > \min\left[h\left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n}\right)\right],$$

Then $E_1(t)$ is true. Summing over all possible $i_1 \in M_{\epsilon+\frac{\gamma}{2}}$ proves the claim. \square

Claim 1: $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}^c} \min\left[h\left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n}\right)\right].$

Proof. By the events in set \mathcal{S} , $C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)}$. Therefore,

$$\begin{aligned} \mu_{i^*} + \frac{\omega}{8(1-\epsilon)} &\geq \mu_{i^*} + 2C_{\delta/n}(T_{i^*}(t)) \stackrel{\mathcal{E}}{\geq} \hat{\mu}_{i^*}(T_{i^*}(t)) + C_{\delta/n}(T_{i^*}(t)) \geq \hat{\mu}_1(T_1(t)) + C_{\delta/n}(T_1(t)) \\ &\stackrel{\mathcal{E}}{\geq} \mu_1. \end{aligned}$$

Hence, $\mu_{i^*} \geq \mu_1 - \frac{\omega}{8(1-\epsilon)}$. Rearranging this, we see that $\mu_{i^*} - \left(1 - \frac{\omega}{8\mu_1(1-\epsilon)}\right)\mu_1 \geq 0$ which implies that $i^* \in M_{\frac{\omega}{8\mu_1(1-\epsilon)}}$. Hence,

$$\begin{aligned} L_t &= (1 - \epsilon) \left(\max_i \hat{\mu}_i(T_i(t)) - C_{\delta/n}(T_i(t)) \right) (1 - \epsilon) (\hat{\mu}_{i^*}(T_{i^*}(t)) - C_{\delta/n}(T_{i^*}(t))) \\ &\stackrel{\mathcal{E}}{\geq} (1 - \epsilon) (\mu_{i^*} - 2C_{\delta/n}(T_{i^*}(t))) \\ &\geq (1 - \epsilon) \left(\mu_{i^*} - \frac{\omega}{8(1-\epsilon)} \right) \\ &\geq (1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) \end{aligned}$$

As before, we seek a lower bound for the difference $(1 - \epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) - \mu_i$.

Case 1: $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$

$$\begin{aligned} (1-\epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) - \mu_i &= (1-\epsilon)\mu_1 - \mu_i - \frac{1}{4} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \\ &\geq \frac{1}{2} ((1-\epsilon)\mu_1 - \mu_i) \end{aligned}$$

since $(1-\epsilon)\mu_1 - \mu_i \geq \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. Therefore, we have that

$$h \left(\frac{1}{2} \left((1-\epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) - \mu_i \right), \frac{\delta}{n} \right) \leq h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right).$$

Lastly, in this setting, $\gamma\mu_1 \leq \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq \epsilon\mu_1 - \Delta_i$ since $\omega = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. Hence, it is trivially true that

$$h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right) = \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n} \right) \right].$$

Case 2: $\omega = \gamma\mu_1$

Assume that $\gamma\mu_1 > \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$, as equality is covered by the previous case. Hence,

$$(1-\epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) - \mu_i = (1-\epsilon)\mu_1 - \mu_i - \frac{1}{4}\gamma\mu_1$$

Recall that we seek to control $i_2 \in M_{\epsilon+\frac{\gamma}{2}}^c$. For any $i \in M_{\epsilon+\frac{\gamma}{2}}^c$, we have that $(1-\epsilon-\frac{\gamma}{2})\mu_1 - \mu_i \geq 0$.

Rearranging, we see that $(1-\epsilon)\mu_1 - \mu_i \geq \frac{1}{2}\gamma\mu_1$ which implies that

$$(1-\epsilon)\mu_1 - \mu_i - \frac{1}{4}\gamma\mu_1 \geq \frac{1}{2}((1-\epsilon)\mu_1 - \mu_i).$$

Therefore, we have that

$$h \left(\frac{1}{2} \left((1-\epsilon) \left(\mu_1 - \frac{\omega}{4(1-\epsilon)} \right) - \mu_i \right), \frac{\delta}{n} \right) \leq h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right)$$

is this setting as well. Similarly, since $\Delta_i - \epsilon\mu_1 \geq \frac{1}{2}\gamma\mu_1$, we likewise have that

$$h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n} \right) \leq \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right].$$

Hence, if T_i exceeds the right-hand side of the preceding inequality, then for any $i \in M_{\epsilon+\frac{\gamma}{2}}^c$, its upper bound is below L_t . Hence, for $i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c$, this implies event $E_2(t)$. Summing over all possible values of $i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c$ proves the claim. \square

Claim 2: The cardinality of \mathcal{S} is bounded as $|\mathcal{S}| \leq \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right]$.

Proof. First, \mathcal{S} may be decomposed as

$$|\mathcal{S}| = |\mathcal{S} \cap \{t : \neg E_1(t)\}| + |\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| + |\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : E_2(t)\}|$$

Note that $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : E_2(t)\}| = 0$ because we have assumed in set \mathcal{S} that $(\text{ST})^2$ has not stopped, and $\{t : E_1(t)\} \cap \{t : E_2(t)\}$ implies termination. By Claim 0, $|\mathcal{S} \cap \{t : \neg E_1(t)\}| \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{4}, \frac{\delta}{n} \right) \right]$. By Claim 1, $|\mathcal{S} \cap \{t : E_1(t)\} \cap \{t : \neg E_2(t)\}| \leq \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}^c} \min \left[h \left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right]$. Recalling that h is assumed to be symmetric in its first argument and summing the two terms proves the claim. \square

B.2.5 Step 4: Putting it all together

Recall that the total number of rounds T that $(\text{ST})^2$ runs for is given by $T = |\{t : \neg \text{STOP}\}|$. To bound this quantity, we have decomposed the set $\{t : \neg \text{STOP}\}$ into many subsets. Below, we show this decomposition.

$$\{t : \neg \text{STOP}\} =$$

$$\begin{aligned}
& \{t : \neg \text{STOP and } i^* \notin M_{\omega/\mu_1}\} \\
& \cup \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) > \frac{\omega}{16(1-\epsilon)} \right\} \\
& \cup \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\} \\
& \cup \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \\
& \quad \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\} \\
& \cup \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \\
& \quad \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\}.
\end{aligned}$$

Hence, by a union bound and plugging in the results of the above steps,

$$\begin{aligned}
& |\{t : \neg \text{STOP}\}| \leq \\
& |\{t : \neg \text{STOP and } i^* \notin M_{\omega/\mu_1}\}| \\
& + \left| \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } \exists i \in M_{\omega/\mu_1} : C_{\delta/n}(T_i(t)) > \frac{\omega}{8(1-\epsilon)} \right\} \right| \\
& + \left| \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\} \right| \\
& + \left| \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \right. \\
& \quad \left. \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}} \right\} \right| \\
& + \left| \left\{ t : \neg \text{STOP and } i^* \in M_{\omega/\mu_1} \text{ and } C_{\delta/n}(T_{i^*}(t)) \leq \frac{\omega}{16(1-\epsilon)} \text{ and } i_1(t) \in M_{\epsilon+\frac{\gamma}{2}} \right. \right. \\
& \quad \left. \left. \text{and } i_2(t) \in M_{\epsilon+\frac{\gamma}{2}}^c \right\} \right| \\
& \leq \sum_{i \in M_{\omega/\mu_1}^c} \min \left\{ h \left(\frac{\gamma\mu_1}{2}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{2}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{2}, \frac{\delta}{n} \right) \right] \right\} \\
& \quad + \sum_{i \in M_{\omega/\mu_1}} \min \left\{ h \left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n} \right) \right] \right\} \\
& \quad + \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}^c} \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right] \\
& \quad + \sum_{i \in M_{\epsilon+\frac{\gamma}{2}}} \min \left[h \left(\frac{\epsilon\mu_1 - \Delta_i}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right] \\
& \quad + \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right] \\
& \stackrel{(\epsilon \leq 1/2)}{\leq} \sum_{i=1}^n \min \left\{ h \left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n} \right) \right] \right\} \\
& \quad + 2 \sum_{i=1}^n \min \left[h \left(\frac{\Delta_i - \epsilon\mu_1}{8}, \frac{\delta}{n} \right), h \left(\frac{\gamma\mu_1}{8}, \frac{\delta}{n} \right) \right] \\
& \leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{16}, \frac{\delta}{n} \right), \min \left[h \left(\frac{\Delta_i}{16}, \frac{\delta}{n} \right), h \left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n} \right) \right] \right\} \right\},
\end{aligned}$$

$$h\left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n}\right)\Big\}$$

Next, by Lemma F.3, we may bound the minimum of $h(\cdot, \cdot)$ functions.

$$\begin{aligned} & 4 \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon\mu_1}{16}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \right\}, \right. \\ & \quad \left. h\left(\frac{\gamma\mu_1}{16}, \frac{\delta}{n}\right) \right\} \\ &= 4 \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon\mu_i}{16}, \frac{\delta}{n}\right), \right. \right. \\ & \quad \left. \min \left[h\left(\frac{\Delta_i}{16}, \frac{\delta}{n}\right), \max \left[h\left(\frac{\tilde{\alpha}_\epsilon}{16(1-\epsilon)}, \frac{\delta}{n}\right), h\left(\frac{\tilde{\beta}_\epsilon}{16(1-\epsilon)}, \frac{\delta}{n}\right) \right] \right] \right\}, \\ & \quad \left. h\left(\frac{\gamma\mu_i}{16}, \frac{\delta}{n}\right) \right\} \\ &\leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon\mu_i}{16}, \frac{\delta}{n}\right), \right. \right. \\ & \quad \left. \max \left[h\left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right) \right] \right\}, \\ & \quad \left. h\left(\frac{\gamma\mu_i}{16}, \frac{\delta}{n}\right) \right\} \\ &= 4 \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon\mu_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right) \right\}, \right. \\ & \quad \left. h\left(\frac{\gamma\mu_i}{16}, \frac{\delta}{n}\right) \right\} \end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $|\epsilon\mu_1 - \Delta_i| \leq 8$ for all i and $\epsilon \leq 1/2$ implies that $\frac{1}{8(1-\epsilon)}|\epsilon\mu_1 - \Delta_i| \leq 2$ and $\frac{1}{8(1-\epsilon)}\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq 2$. $\Delta_i \leq 16$ for all i , gives $0.125\Delta_i \leq 2$. Lastly, $\gamma \leq 16/\mu_1$ implies that $\frac{\gamma\mu_1}{8} \leq 2$. Therefore,

$$\begin{aligned} & 4 \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon\mu_i}{16}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{32}, \frac{\delta}{n}\right) \right\}, \right. \\ & \quad \left. h\left(\frac{\gamma\mu_i}{16}, \frac{\delta}{n}\right) \right\} \\ &\leq 4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{(\epsilon\mu_1 - \Delta_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta(\epsilon\mu_1 - \Delta_i)^2} \right) \right), \right. \right. \\ & \quad \frac{4096}{(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right), \\ & \quad \left. \frac{4096}{(\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon})^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon})^2} \right) \right) \right\} \\ & \quad \frac{1024}{\gamma^2\mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta\gamma^2\mu_1^2} \right) \right) \Big\} \\ &= 4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1024}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right. \end{aligned}$$

$$\begin{aligned} & \frac{4096}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right), \\ & \frac{4096}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{12288n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \right) \right) \Bigg\}, \\ & \frac{1024}{\gamma^2 \mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{3072n}{\delta \gamma^2 \mu_1^2} \right) \right) \Bigg\}. \end{aligned}$$

The above bounds the number of rounds T . Therefore, the total number of samples is at most $3T$. \square

C Proof of instance dependent lower bounds, Theorem 2.1

First we restate and prove the lower bound.

Theorem C.1. (*additive and multiplicative lower bound*) Fix $\delta, \epsilon > 0$. Consider n arms, such that the i^{th} is distributed according to $\mathcal{N}(\mu_i, 1)$. Any δ -PAC algorithm for the *additive* setting satisfies

$$\mathbb{E}[\tau] \geq 2 \sum_{i=1}^n \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right\} \log \left(\frac{1}{2.4\delta} \right)$$

and if $\mu_1 > 0$ any δ -PAC algorithm for the *multiplicative* algorithm satisfies,

$$E[\tau] \geq 2 \sum_{i=1}^n \max \left\{ \frac{1}{((1-\epsilon)\mu_1 - \mu_i)^2}, \frac{1}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \right\} \log \left(\frac{1}{2.4\delta} \right)$$

Proof of Theorem 2.1 in the additive case. Recall that ν denotes the given instance, and without loss of generality we have assumed that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. Then $G_\epsilon(\nu) = \{1, \dots, k\}$. Consider the event E that an algorithm returns $\{1, \dots, k\}$. For any δ -PAC algorithm, E occurs with probability at least $1 - \delta$. For each arm $i \in [n]$ we consider two alternative instances

$$\nu'_i = \{\mu_1, \dots, \mu'_i, \dots, \mu_n\}$$

and

$$\nu''_i = \{\mu_1, \dots, \mu''_i, \dots, \mu_n\}$$

such that only the mean of arm i differs compared to ν but $G_\epsilon(\nu) \neq G_\epsilon(\nu'_i)$ and $G_\epsilon(\nu) \neq G_\epsilon(\nu''_i)$. Therefore, on these alternate instances, E occurs with probability at most δ .

For ν'_i , if $i \leq k$, let $\mu'_i = \mu_1 - \epsilon - \eta$. Then $i \in G_\epsilon(\nu)$ but $i \notin G_\epsilon(\nu'_i)$. If $k < n$ and $i \geq k+1$, let $\mu'_i = \mu_1 - \epsilon + \eta$. Then $i \notin G_\epsilon(\nu)$ but $i \in G_\epsilon(\nu'_i)$.

More subtly, for ν''_i , for any $i \in [n] \setminus \{k\}$, let $\mu''_i = \mu_k + \epsilon + \eta$. In particular, arm i is now the best arm. Under this definition, $\mu''_i - \epsilon > \mu_k$. Therefore, $k \notin G_\epsilon(\nu''_i)$ but $k \in G_\epsilon(\nu)$.

The above holds for all $\eta > 0$. Let N_i denote the random variable of the number of samples of arm i and \mathbb{E}_ν denote expectation with respect to instance ν . Using the fact that we have assumed the distributions are Gaussian, considering ν'_i , by Lemma 1 of [6], taking $\eta \rightarrow 0$ we have that for any δ -PAC algorithm,

$$\mathbb{E}_\nu[N_i] \geq \frac{2 \log(1/2.4\delta)}{(\mu_i - (\mu_1 - \epsilon))^2}.$$

Furthermore, considering ν''_i , and again taking $\eta \rightarrow 0$, we have by the same lemma that for $i \neq k$

$$\mathbb{E}_\nu[N_i] \geq \frac{2 \log(1/2.4\delta)}{(\mu_k + \epsilon - \mu_i)^2} = \frac{2 \log(1/2.4\delta)}{(\mu_1 + \alpha_\epsilon - \mu_i)^2},$$

where the later equality holds since $\mu_k + \epsilon = \mu_1 + \alpha_\epsilon$ by definition of α_ϵ . For $i = k$, note that $\frac{1}{(\mu_k - (\mu_1 - \epsilon))} = \frac{1}{\alpha_\epsilon^2} \geq \frac{1}{\epsilon^2} = \frac{1}{(\mu_k - \mu_k - \epsilon)^2}$ since $\alpha_\epsilon = \min_{i \in G_\epsilon} \mu_i - (\mu_1 - \epsilon) = \min_{i \in G_\epsilon} \epsilon - \Delta_i$. Putting these pieces together, we see that for any i ,

$$\mathbb{E}_\nu[N_i] \geq \max \left(\frac{1}{(\mu_i - (\mu_1 - \epsilon))^2}, \frac{1}{(\mu_k + \epsilon - \mu_i)^2} \right) 2 \log(1/2.4\delta).$$

Summing over all i establishes a lower bound in the *additive* case. \square

Proof of Theorem 2.1 in the multiplicative case. Recall that ν denotes the given instance, and without loss of generality we have assumed that $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$. Let $M_\epsilon(\nu) = \{1, \dots, k\}$. Consider the event E that an algorithm returns $\{1, \dots, k\}$. For any δ -PAC algorithm, E occurs with probability at least $1 - \delta$. For each arm $i \in [n]$ we consider two alternative instances

$$\nu'_i = \{\mu_1, \dots, \mu'_i, \dots, \mu_n\}$$

and

$$\nu''_i = \{\mu_1, \dots, \mu''_i, \dots, \mu_n\}$$

such that only the mean of arm i differs compared to ν but $M_\epsilon(\nu) \neq M_\epsilon(\nu'_i)$ and $M_\epsilon(\nu) \neq M_\epsilon(\nu''_i)$. Therefore, E occurs with probability at most δ on these alternate instances.

For ν'_i , if $i \leq k$, let $\mu'_i = (1 - \epsilon - \eta)\mu_1$. Then $i \in M_\epsilon(\nu)$ but $i \notin M_\epsilon(\nu'_i)$. If $k < n$ and $i \geq k + 1$, let $\mu'_i = (1 - \epsilon + \eta)\mu_1$. Then $i \notin M_\epsilon(\nu)$ but $i \in M_\epsilon(\nu'_i)$.

More subtly, for ν''_i , for any $i \in [n] \setminus \{k\}$, let $\mu''_i = \frac{\mu_k}{1 - \epsilon - \eta}$. In particular, arm i is now the best arm. Under this definition, $\mu''_i - \epsilon > \mu_k$. Therefore, $k \notin M_\epsilon(\nu''_i)$ but $k \in M_\epsilon(\nu)$.

The above holds for all $\eta > 0$. Let N_i denote the random variable of the number of samples of arm i and \mathbb{E}_ν denote expectation with respect to instance ν . Using the fact that we have assumed the distributions are Gaussian, considering ν'_i , by Lemma 1 of [6], taking $\eta \rightarrow 0$, we have that for any δ -PAC algorithm,

$$\mathbb{E}_\nu[N_i] \geq \frac{2 \log(1/2.4\delta)}{(\mu_i - (1 - \epsilon)\mu_1)^2} = \frac{2 \log(1/2.4\delta)}{(\epsilon\mu_1 - \Delta_i)^2}.$$

Additionally, by the same Lemma, considering ν''_i and again taking $\eta \rightarrow 0$ we have that for $i \neq k$

$$\mathbb{E}_\nu[N_i] \geq \frac{2 \log(1/2.4\delta)}{\left(\mu_i - \frac{\mu_k}{1 - \epsilon}\right)^2} = \frac{2 \log(1/2.4\delta)}{\left(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i\right)^2},$$

where the later equality holds since $\frac{\mu_k}{1 - \epsilon} = \mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon}$ by definition of $\tilde{\alpha}_\epsilon$. Next recall that $\tilde{\alpha}_\epsilon := \min_{i \in M_\epsilon} \mu_i - (1 - \epsilon)\mu_1 = \mu_k - (1 - \epsilon)\mu_1$, we have that $\mu_k = \tilde{\alpha}_\epsilon + (1 - \epsilon)\mu_1$. Hence, $\frac{\mu_k}{1 - \epsilon} = \mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon}$. Then, for $i = k$

$$\begin{aligned} \frac{1}{\left(\frac{\mu_k}{1 - \epsilon} - \mu_k\right)^2} &\leq \frac{1}{(\mu_k - (1 - \epsilon)\mu_1)^2} = \frac{1}{\tilde{\alpha}_\epsilon^2} \\ \iff \tilde{\alpha}_\epsilon &\leq \frac{\mu_k}{1 - \epsilon} - \mu_k = \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} + \mu_1 - \mu_k = \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} + \Delta_k \\ \stackrel{(\Delta_k \geq 0)}{\iff} \tilde{\alpha}_\epsilon &\leq \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} \end{aligned}$$

which is always true since $\epsilon > 0$. Therefore,

$$\frac{1}{(\mu_k - (1 - \epsilon)\mu_1)^2} = \max \left(\frac{1}{(\mu_k - (1 - \epsilon)\mu_1)^2}, \frac{1}{\left(\frac{\mu_k}{1 - \epsilon} - \mu_k\right)^2} \right).$$

Hence, for all arms i ,

$$\mathbb{E}_\nu[N_i] \geq 2 \max \left(\frac{1}{(\mu_i - (1 - \epsilon)\mu_1)^2}, \frac{1}{\left(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i\right)^2} \right) \log(1/2.4\delta).$$

Summing over all i gives a lower bound for this problem in the multiplicative case. \square

D Theorem 4.1: Lower bounds in the moderate confidence regime

In this section, we prove a tighter lower bound that includes *moderate confidence* terms independent of the value of δ similar to those that appear in the upper bound on the sample complexity of FAREAST, Theorem 4.2.

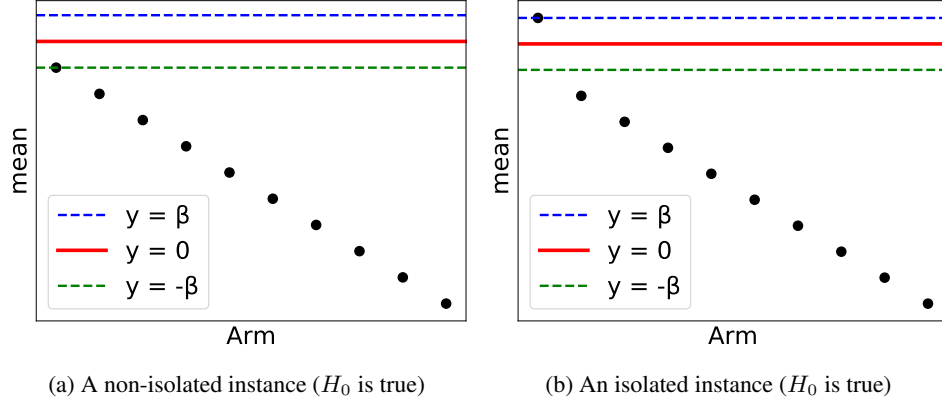


Figure 14: Example of an isolated and non-isolated instance

Outline. To give a tight lower bound in the isolated setting, we break our argument into pieces performing a series of reductions that link the all- ϵ problem to a hypothesis test, and then the hypothesis test to the problem of identifying the best-arm.

Step 1. Finding an isolated arm. We first consider the following problem. Imagine that you are given an *isolated* instance, depicted in Figure 14b where there are n distributions, with one of them at mean β and the rest with mean $-\beta$. Theorem D.3, captures the sample complexity of any algorithm that can return i^* with probability greater than $1 - \delta$.

Step 2. Deciding if an instance is isolated. We then consider a composite hypothesis test on n distributions where the null hypothesis, H_0 , is that the mean of each distribution is less than $-\beta$ and the alternate hypothesis, H_1 , is that there exists *single* distribution i^* with mean β and the remainder have mean less than $-\beta$ (i.e. the instance is isolated). In Figure 14, we show a picture of an instance where the null is true and where the alternate is true. In Theorem D.6 we lower bound the complexity of performing this test. To link this to Step 1, we show that if you can solve this composite hypothesis test then you can find i^* , hence the lower bound of step 1 is a lower bound for the hypothesis test.

Step 3: Reducing ALL- ϵ to Step 2 Finally in step 3 we link this to the all- ϵ problem. Using the above, we lower bound the complexity of ALL- ϵ in Theorem 4.1 when $|G_{2\beta\epsilon}| = 1$. The key insight of our proof is that any algorithm that can solve the ALL- ϵ problem can be used to solve the hypothesis test in Step 2.

D.1 Step 1: Finding an Isolated Arm

Fix $n \in \mathbb{N}$, $0 < \beta$, and $\delta > 0$. We refer to a β -*isolated instance* $\nu = \{\rho_1, \dots, \rho_n\}$, as a collection of n , Gaussian distributions with variance one satisfying two properties. Firstly, there exists a single arm $i^* \in [n]$ with $\rho_{i^*} = \mathcal{N}(\beta, 1)$. We refer to this as the *isolated arm*. Secondly, for $i \neq i^*$, $\rho_i = \mathcal{N}(\mu_i, 1) \forall i \in [n] \setminus \{i^*\}$ have means $\mu_i \leq -\beta$. We introduce the additional notation $\Delta_{i,j} = \mu_i - \mu_j$.

Lemma D.1. Fix n , $0 < \beta$ and consider a set ν of n Gaussian random variables such that for a uniformly random chosen $i^* \in [n]$, $\rho_{i^*} = \mathcal{N}(\beta, 1)$ and $\rho_i = \mathcal{N}(\mu_i, 1)$ for $\mu_i \leq -\beta$ for all $i \neq i^*$. Any algorithm that correctly returns i^* with probability at least $1 - \delta$, pulls arm i^* at least

$$\frac{1}{2\beta^2} \log(1/2.4\delta)$$

times in expectation.

Proof. Consider the oracle setting where the value of i^* is known and the algorithm only seeks to confirm that $\mu_{i^*} > -\beta$. Lemma 1 of [6] implies that any δ -PAC algorithm requires at least $\frac{1}{2\beta^2} \log(1/2.4\delta)$ samples in expectation. \square

The above bound controls the number of samples that any algorithm must gather from i^* , and is independent of n . The proof considered an oracle setting where the value of i^* is known, and

one only wishes to confirm that $\mu_{i^*} > -\beta$ with probability at least $1 - \delta$. To lower bound the number of samples drawn from $[n] \setminus \{i^*\}$, we need significantly more powerful tools. In particular, to rule out trivial algorithms that always output a fixed index, we consider a permutation model, as in [9, 11, 15, 25]. Informally, we consider an additional expectation in the lower bound over a random permutation π of the arms where π is sampled uniformly from the set of all permutations. In particular, we will use a *Simulator* argument, as in [9, 15]. In what follows, we will let $\pi : [n] \rightarrow [n]$ denote a permutation selected uniformly at random from the set of $n!$ permutations. For instance ν , let $\pi(\nu)$ denote the permuted instance such that the i^{th} distribution is mapped to $\pi(i)$, by a slight overloading of the definition of $\pi(\cdot)$. In what follows, we proceed similarly to the proof of Theorem 1 in [15].

Theorem D.2. Fix n , $0 < \beta$, and $\delta < 1/16$ and consider a set ν of n Gaussian random variables with variance 1 such that for $i^* \in [n]$, $\rho_{i^*} = \mathcal{N}(\beta, 1)$ and $\rho_i = \mathcal{N}(\mu_i, 1)$ for $\mu_i \leq -\beta$ for all $i \neq i^*$. Let π be a uniformly chosen permutation of $[n]$ and $\pi(\nu)$ be the permutation applied to instance ν . Let T be the random variable denoting the total number of samples at termination by an algorithm. Any δ -PAC algorithm to detect $\pi(i^*)$ on $\pi(\nu)$ requires

$$\mathbb{E}_\pi \mathbb{E}_{\pi(\nu)} [T] \geq \frac{1}{16} \sum_{k \neq i^*} \frac{1}{\Delta_{i^*, k}^2}$$

samples in expectation from arms in $[n] \setminus \{i^*\}$.

Proof. Fix a permutation π . Let $\pi(\nu)$ be the permutation applied to ν and $\pi(i)$ be the index of i under the permuted instance, $\pi(\nu)$. Let \mathcal{A} be any algorithm that detects and returns $\pi(i^*)$ on $\pi(\nu)$ with probability at least $1 - \delta$. We will take $\mathbb{P}_\mathcal{A}$ and $\mathbb{E}_\mathcal{A}$ to denote probability and expectation with respect any internal randomness in \mathcal{A} . Throughout, we will take $\rho_i = \mathcal{N}(\mu_i, 1)$ to denote the i^{th} distribution of ν . $\mu_{i^*} > 0$ and $\mu_i < 0$ for all $i \neq i^*$. Additionally, let $\Delta_{ij} = \mu_i - \mu_j$

Fix $k \neq i^*$. To bound the necessary number of samples for arm k , we turn to the Simulator [9]. We begin by defining an alternate instance $\nu'_k = \{\rho'_1, \dots, \rho'_n\}$ as

$$\rho'_j = \begin{cases} \rho_j, & j \neq i^* \\ \rho_k, & j = i^* \\ \rho_{i^*}, & j = k \end{cases}$$

Note that ν'_k is identical to ν except that the distributions of i^* and k are swapped.

Let E be the event that \mathcal{A} returns $\pi(i^*)$. We may bound the total variation distance between the joint distribution on $\mathcal{A} \times \pi(\nu)$ and $\mathcal{A} \times \pi(\nu'_k)$ as

$$\begin{aligned} TV(\mathbb{P}_{\mathcal{A} \times \pi(\nu)}, \mathbb{P}_{\mathcal{A} \times \pi(\nu'_k)}) &= \sup_A \left| \mathbb{P}_{\mathcal{A} \times \pi(\nu)}(A) - \mathbb{P}_{\mathcal{A} \times \pi(\nu'_k)}(A) \right| \\ &\geq \left| \mathbb{P}_{\mathcal{A} \times \pi(\nu)}(E) - \mathbb{P}_{\mathcal{A} \times \pi(\nu'_k)}(E) \right| \\ &\geq 1 - 2\delta. \end{aligned}$$

Let Ω_t denote the multiset of the transcript of samples up to time t .

$$\Omega_t = \{i_s \in [n] \text{ for } 1 \leq s \leq t\}$$

and define the events

$$W_j(\Omega_t) := \left\{ \sum_{i_t \in \Omega_t} \mathbb{1}(i_t = j) \leq \tau \right\}$$

for a τ to be defined later. With the definitions of $W_j(\Omega_t)$, we define a simulator $\text{Sim}(\nu, \Omega_t)$ with respect to ν . Let $\text{Sim}(\nu, \Omega_t)_i$ denote the distribution of arm i on $\text{Sim}(\nu, \Omega_t)$.

$$\text{Sim}(\nu, \Omega_t)_j = \begin{cases} \rho_j, & \text{if } j \notin \{i^*, k\} \\ \rho_j, & \text{if } j \in \{i^*, k\} \text{ and } W_{i^*}(\Omega_t) \cap W_k(\Omega_t) \\ \rho_{i^*}, & \text{if } j \in \{i^*, k\} \text{ and } (W_{i^*}(\Omega_t) \cap W_k(\Omega_t))^c \end{cases}$$

Furthermore, we define $\text{Sim}(\nu'_k, \Omega_t)$ with respect to ν'_k as

$$\text{Sim}(\nu'_k, \Omega_t)_j = \begin{cases} \rho'_j, & \text{if } j \notin \{i^*, k\} \\ \rho'_j, & \text{if } j \in \{i^*, k\} \text{ and } W_{i^*}(\Omega_t) \cap W_k(\Omega_t) \\ \rho_{i^*}, & \text{if } j \in \{i^*, k\} \text{ and } (W_{i^*}(\Omega_t) \cap W_k(\Omega_t))^c \end{cases}$$

For ease of notation, let $\text{Sim}(\pi(\nu), \Omega_t)$ be the same simulator defined on $\pi(\nu)$ and with respect to events $W_{\pi(i^*)}(\Omega_t)$ and $W_{\pi(k)}(\Omega_t)$. Note that in the simulator of ν'_k , if $(W_{i^*}(\Omega_t) \cap W_k(\Omega_t))^c$ is true, then i^* and k draw samples according to instance ν not ν'_k .

Definition 4. (Truthfulness of an event W , [15]) For an algorithm \mathcal{A} , we say that an event W is truthful on a simulator $\text{Sim}(\eta)$ with respect to an instance η if for all events E in the filtration \mathcal{F}_T generated by playing algorithm \mathcal{A} on instance η

$$\mathbb{P}_\eta(E \cap W) = \mathbb{P}_{\text{Sim}(\eta)}(E \cap W)$$

By our definition of both simulators, if $(W_{i^*}(\Omega_t) \cap W_k(\Omega_t))^c$ is true, then $\text{Sim}(\nu, \Omega_t)_i = \text{Sim}(\nu'_k, \Omega_t)_i \forall i \in [n]$. Contrarily, if $W_{i^*}(\Omega_t) \cap W_k(\Omega_t)$ is true, then $\text{Sim}(\nu, \Omega_t) = \nu$ and $\text{Sim}(\nu'_k, \Omega_t) = \nu'_k$. Similarly, on $W_{\pi(i^*)}(\Omega_t) \cap W_{\pi(k)}(\Omega_t)$, $\text{Sim}(\pi(\nu), \Omega_t) = \pi(\nu)$ and $\text{Sim}(\pi(\nu'_k), \Omega_t) = \pi(\nu'_k)$. Therefore, by the proof of Theorem 1 of [15], $W_{\pi(k)}(\Omega_t)$ is truthful on $\text{Sim}(\pi(\nu), \Omega_t)$ and $W_{\pi(i^*)}(\Omega_t)$ is truthful on $\text{Sim}(\pi(\nu'_k), \Omega_t)$.

Let i_t be the arm queried at time $t \in \mathbb{N}$ by \mathcal{A} . Following the proof of Theorem 1 of [15], we may bound the KL-Divergence between $\text{Sim}(\pi(\nu), \Omega_t)$ and $\text{Sim}(\pi(\nu'_k), \Omega_t)$ as

$$\begin{aligned} & \max_{i_1, \dots, i_T} \sum_{t=1}^T KL(\text{Sim}(\pi(\nu), \{i_s\}_{s=1}^t), \text{Sim}(\pi(\nu'_k), \{i_s\}_{s=1}^t)) \\ & \leq \tau KL(\pi(\nu)_{\pi(i^*)}, \pi(\nu'_k)_{\pi(i^*)}) + \tau KL(\pi(\nu)_{\pi(k)}, \pi(\nu'_k)_{\pi(k)}) \\ & = \tau \frac{\Delta_{i^*, k}^2}{2} + \tau \frac{\Delta_{i^*, k}^2}{2} \\ & = \tau \Delta_{i^*, k}^2. \end{aligned}$$

For any instance η , an algorithm \mathcal{A} is defined to be symmetric if

$$\mathbb{P}_{\mathcal{A}, \eta}((i_1, \dots, i_T) = (I_1, \dots, I_T)) = \mathbb{P}_{\mathcal{A}, \pi(\eta)}((\pi(i_1), \dots, \pi(i_T)) = (\pi(I_1), \dots, \pi(I_T))).$$

Semantically, this implies that the proportion of times \mathcal{A} pulls any arm i on the non-permuted instance η is the same as the proportion of times it pulls $\pi(i)$ on the permuted instance, $\pi(\eta)$.

In particular, the expected complexity of a symmetric algorithm is independent of the permutation π . By Lemma 1 of [9], if any algorithm \mathcal{B} (not necessarily symmetric) achieves an expected stopping time τ where the expectation is taken over all the randomness in the permutation and in the instance, then there is a symmetric algorithm that achieves the same expected stopping time. Hence, we may assume that \mathcal{A} is symmetric and capture the same set of possible stopping times. If \mathcal{A} is not symmetric, we may form an algorithm \mathcal{A}' by permuting the input, passing it to \mathcal{A} , getting the output of \mathcal{A} on the permuted input, and then undoing the permutation before return an answer.

Since $W_{\pi(k)}(\Omega_t)$ and $W_{\pi(i^*)}(\Omega_t)$ are truthful on $\text{Sim}(\pi(\nu), \Omega_t)$ and $\text{Sim}(\pi(\nu'_k), \Omega_t)$ respectively, by Lemma 2 of [9], we have that

$$\begin{aligned} & \mathbb{P}_{\mathcal{A}, \pi(\nu)}(W_{\pi(k)}(\Omega_t)) + \mathbb{P}_{\mathcal{A}, \pi(\nu'_k)}(W_{\pi(i^*)}(\Omega_t)) \\ & \geq TV(\mathbb{P}_{\mathcal{A}, \pi(\nu)}, \mathbb{P}_{\mathcal{A}, \pi(\nu'_k)}) - Q(KL(\text{Sim}(\pi(\nu), \Omega_t), \text{Sim}(\pi(\nu'_k), \Omega_t))) \end{aligned}$$

for $Q(x) = \min\{1 - 1/2e^{-x}, \sqrt{x/2}\}$. Since \mathcal{A} is symmetric, for any permutation π , we have that

$$\mathbb{P}_{\mathcal{A}, \pi(\nu)}(W_{\pi(k)}(\Omega_t)) + \mathbb{P}_{\mathcal{A}, \pi(\nu'_k)}(W_{\pi(i^*)}(\Omega_t)) = \mathbb{P}_{\mathcal{A}, \nu}(W_k(\Omega_t)) + \mathbb{P}_{\mathcal{A}, \nu'_k}(W_{i^*}(\Omega_t)) = 2\mathbb{P}_{\mathcal{A}, \nu}(W_k(\Omega_t)).$$

The first equality holds since event W_i depend only on the number of times that arm i is pulled. Since \mathcal{A} is symmetric, the probability that \mathcal{A} pulls arm i at most τ times on instance ν is equal to the probability that \mathcal{A} pulls $\pi(i)$ at most τ times on instance $\pi(\nu)$. The second equality is true using symmetry as well since instances ν and ν'_k are themselves equal up to a permutation.

Combining the above with the previous bounds on the total variation and KL divergence, we have that

$$\mathbb{P}_{\mathcal{A} \times \nu}(N_k > \tau) = \mathbb{P}_{\mathcal{A} \times \nu}(W_k(\Omega_t)) \geq \frac{1}{2} \left(1 - 2\delta - \sqrt{\frac{\tau \Delta_{i^*,k}^2}{2}} \right)$$

Plugging in $\tau = 1/(2\Delta_{i^*,k}^2)$, we see that $\mathbb{P}_{\mathcal{A} \times \nu}(N_k > 1/(2\Delta_{i^*,k}^2)) \geq 1/2(1/2 - 2\delta)$. Since k was arbitrary, we may repeat this argument for each k in $[n] \setminus \{i^*\}$. Combining this with Markov's inequality, we see that

$$\begin{aligned} \mathbb{E}_{\mathcal{A} \times \nu} \left[\sum_{k \neq i^*} N_k \right] &\geq \frac{1}{4}(1/2 - 2\delta) \sum_{k \neq i^*} \frac{1}{\Delta_{i^*,k}^2} \\ &> \frac{1}{16} \sum_{k \neq i^*} \frac{1}{\Delta_{i^*,k}^2} \end{aligned}$$

where the final inequality follows from $\delta < 1/16$. The above holds for any δ -PAC algorithm \mathcal{A} . \square

We now state our strong lower bound on the expected number of samples for any algorithm that can find an isolated arm.

Theorem D.3. Fix n , $0 < \beta$, and $\delta < 1/16$ and consider a set ν of n Gaussian random variables with variance 1 such that for a uniformly random chosen $i^* \in [n]$, $\rho_{i^*} = \mathcal{N}(\beta, 1)$ and $\rho_i = \mathcal{N}(\mu_i, 1)$ for $\mu_i \leq -\beta$ for all $i \neq i^*$. Let π be a uniformly chosen permutation of $[n]$ and $\pi(\nu)$ be the permutation applied to instance ν . Any δ -PAC algorithm to detect $\pi(i^*)$ on $\pi(\nu)$ requires

$$\frac{1}{16} \sum_{k \neq i^*} \frac{1}{\Delta_{i^*,k}^2} + \frac{1}{2\beta^2} \log(1/2.4\delta)$$

samples in expectation, where the expectation is taken both over the randomness in the permutation, the randomness in $\pi(\nu)$, and any internal randomness to the algorithm.

Proof. By Lemma D.1, arm i^* must be sampled $\frac{1}{2\beta^2} \log(1/2.4\delta)$ times. By Theorem D.2, arms in $[n] \setminus \{i^*\}$ must collectively be sampled $\frac{1}{16} \sum_{k \neq i^*} \frac{1}{\Delta_{i^*,k}^2}$ times. Joining these two results gives the stated result. \square

D.2 Step 2. Deciding if an instance is isolated

Next, we consider a composite hypothesis test that is related to the question of finding an isolated arm. As we will show, this test has the interesting property that the alternate hypothesis may be declared in significantly fewer samples than the null.

Definition 5 (β -Isolated Hypothesis Test). Fix $0 < \epsilon$ and $0 < \beta$. Consider an instance $\nu = \{\rho_1, \dots, \rho_n\}$ where $\rho_i = \mathcal{N}(\mu_i, 1)$. By sampling individual distributions ρ_i , one wishes to perform the following composite hypothesis test:

Null Hypothesis H_0 : $\mu_i < -\beta$ for all $i \in [n]$.

Alternate Composite Hypothesis H_1 : $\exists i^* : \mu_{i^*} = \beta > 0$ and $\mu_i \leq -\beta$ for all $i \neq i^*$.

For any instance ν , we say “ H_1 is true on ν ” if $\exists i^* : \mu_{i^*} = \beta > 0$ and otherwise we say “ H_0 is true on ν .” Next, we bound the sample complexity of any algorithm to perform the β -isolated hypothesis test with probability at least $1 - \delta$ in the case that H_0 is true.

Figure 2 shows an two example instance. One where H_0 is true and one where H_1 is true.

Lemma D.4. Fix n , β , and δ and consider a set ν of n standard normal random variables where H_0 is true. Any algorithm to correctly declare H_0 in the β -isolated hypothesis test problem with probability at least $1 - \delta$ requires

$$\sum_{i=1}^n \frac{2}{(\beta - \mu_i)^2} \log \left(\frac{1}{2.4\delta} \right)$$

samples in expectation.

Proof. Notice that for each $i \in [n]$, we may construct an alternate instance ν_i by changing the distribution of ρ_i to be $\mathcal{N}(\beta, 1)$ and leaving others unchanged. On ν_i , H_1 is instead true. To distinguish between ν and ν_i , necessary to declare H_0 versus H_1 , by Lemma 1 of [6], any δ -PAC algorithm requires $\mathbb{E}_\nu[N_i] \geq \frac{2}{(\beta - \mu_i)^2} \log(1/2.4\delta)$ where \mathbb{E}_ν denotes expectation with respect to the instance ν and N_i denotes the number of samples of arm i . Repeating this argument for each $i \in [n]$ gives the desired result. \square

To lower bound the expected sample complexity of any algorithm to perform the β -isolated hypothesis test in the setting where H_1 is true, we consider a reduction to the problem studied in Step 1, Section D.1. For the reduction to an algorithm that can find an isolated arm, we show that if there is an algorithm to declare H_1 in fewer than $O\left(\sum_{i=1}^n \frac{1}{\Delta_{i^*,k}^2}\right)$ samples, then one can design an algorithm akin to binary search that returns i^* in fewer than $O\left(\sum_{i=1}^n \frac{1}{\Delta_{i^*,k}^2}\right)$ samples, contradicting Lemma D.2.

Lemma D.5. Fix n , β , and $\delta < 1/16$. Let π be a random permutation. Consider an instance ν where H_1 is true. In this setting, any algorithm to correctly declare H_1 in the β -Isolated Hypothesis Testing problem on $\pi(\nu)$ with probability at least $1 - \delta$ requires $\frac{1}{32} \sum_{j \neq i^*} \Delta_{i^*,k}^{-2}$ samples in expectation.

Proof. Fix $\delta > 0$ and let i^* denote the single distribution such that $\rho_{i^*} = \mathcal{N}(\beta, 1)$ where $\beta > 0$. In particular, only i^* has a positive mean. Assume for contradiction that there is an algorithm $\mathcal{A}(\pi(\nu), \delta, \beta)$ that correctly declares H_1 on $\pi(\nu)$ in at most $\frac{1}{32} \sum_{k \neq i^*} \Delta_{i^*,k}^{-2}$ samples in expectation with probability at least $1 - \delta$ on any instance $\pi(\nu)$ of n distributions if H_1 is true. Otherwise, if H_0 is true, assume that \mathcal{A} correctly declares H_0 in an arbitrary number of samples in expectation, $N_{H_0}(\nu)$ lower bounded by Lemma D.4. As in the proof of Theorem D.2, if any algorithm \mathcal{B} (not necessarily symmetric) achieves an expected stopping time τ where the expectation is taken over all the randomness in the permutation and in the instance, by Lemma 1 of [9], there is a symmetric algorithm that achieves the same expected stopping time. Hence, we may assume that \mathcal{A} is symmetric and capture the same set of possible stopping times. For the remainder of this proof, we assume \mathcal{A} is symmetric. Therefore, its expected complexity is independent of the permutation π . Without loss of generality, assume that $n = 2^k$ for some $k \in \mathbb{N}$. Otherwise, we may hallucinate $(2^{\lceil \log_2(n) \rceil} - n)$ normal distributions, $\mathcal{N}(-\beta, 1)$, and form an instance ν' comprised of these additional distribution and those in ν . If so, anytime \mathcal{A} requests a sample from a distribution in $\nu' \setminus \nu$, draw a sample from $\mathcal{N}(-\beta, 1)$ and pass it to \mathcal{A} , only tracking the number of samples drawn from ν .

Step a). In what follows, we use \mathcal{A} to develop a method for isolated-arm identification. To do so, we show that one may use \mathcal{A} to perform binary search for the distribution i^* such that $\rho_{i^*} = \mathcal{N}(\beta, 1)$ and this leads to a contradiction of Theorem D.2. For ease of exposition, for a set $\mathcal{S} \subset [n]$, let $\nu(\mathcal{S}) := \{i \in \mathcal{S} : \rho_i\}$, the subset of instance ν of distributions whose indices are in \mathcal{S} .

If H_1 is true on $\nu(\mathcal{S})$, by assumption, with probability at least $1 - \delta$, \mathcal{A} correctly declares H_1 on $\nu(\mathcal{S})$ in at most $\frac{1}{32} \sum_{i \in \mathcal{S} \setminus \{i^*\}} \Delta_{i^*,k}^{-2}$ samples in expectation. Similarly, if H_0 is true on $\nu(\mathcal{S})$, the sample complexity is $N_{H_0}(\nu(\mathcal{S}))$ in expectation.

Algorithm 3 Binary search for Isolated Arm Identification

Require: $\delta > 0, \beta > 0$, instance ν such that H_1 is true, algorithm \mathcal{A}

```

1: Let Low = 1 and High =  $n$ 
2: for  $i = 1, \dots, \log_2(n)$  do
3:   1) Choose sets  $\mathcal{S}_1, \mathcal{S}_2$  uniformly at random such that  $\mathcal{S}_1 \cup \mathcal{S}_2 = \mathcal{S}$ ,  $\mathcal{S}_1 \cap \mathcal{S}_2 = \emptyset$ , and
       $\mathbb{P}(i \in \mathcal{S}_1) = \mathbb{P}(i \in \mathcal{S}_2)$  for all  $i \in \mathcal{S}$ 
4:   2) In parallel, run  $\mathcal{A}_1 = \mathcal{A}(\nu(\mathcal{S}_1), \beta, \delta/2 \log_2(n))$  and  $\mathcal{A}_2 = \mathcal{A}(\nu(\mathcal{S}_2), \beta, \delta/2 \log_2(n))$ 
5:   3) If either terminates, terminate the other
6:   if  $\mathcal{A}_1$  declares  $H_1$  or  $\mathcal{A}_2$  declares  $H_0$  then
7:      $\mathcal{S} = \mathcal{S}_1$ 
8:   else
9:      $\mathcal{S} = \mathcal{S}_2$ 
return  $i^* \in \mathcal{S}$  (note:  $|\mathcal{S}| = 1$  at this point)
  
```

In step 1, we choose 2 random subsets of \mathcal{S} , \mathcal{S}_1 and \mathcal{S}_2 that partition \mathcal{S} such that each arm is assigned with equal probability to either \mathcal{S}_1 or \mathcal{S}_2 independently.

In step 2) if the loop, we separately run \mathcal{A} in parallel on $\nu(\mathcal{S}_1)$ and $\nu(\mathcal{S}_2)$, each with failure probability $\delta/2 \log_2(n)$. We alternate between passing a sample to \mathcal{A}_1 and to \mathcal{A}_2 .

In Step 3), we terminate \mathcal{A}_1 if \mathcal{A}_2 terminates and vice versa. If, for instance, \mathcal{A}_1 terminates and declares H_0 , we may infer H_1 on \mathcal{S}_2 . Alternately, if \mathcal{A}_1 declares H_1 on \mathcal{S}_1 , we may infer H_0 on \mathcal{S}_2 as there is a single positive mean, μ_{i^*} . This process continues until $|\mathcal{S}_1| = |\mathcal{S}_2| = 1$, when there is a single distribution remaining in each. At this point, if \mathcal{A}_1 declares H_1 , then the single arm $i \in \mathcal{S}_1$ is the positive mean i^* . Otherwise, the single arm $j \in \mathcal{S}_2$ is.

First, we show that this algorithm is correct with probability at least $1 - \delta$. The algorithm errs if and only if in any round i , either \mathcal{A}_1 or \mathcal{A}_2 errs, each with occurs with probability at most $\delta/2 \log_2(n)$. Union bounding over the $\log_2(n)$ rounds, we see that the algorithm errs with probability at most δ . For the remainder of the proof, we will assume that in no round does either \mathcal{A}_1 or \mathcal{A}_2 incorrectly declare H_0 or H_1 if the reverse is true for the given instances $\nu(\mathcal{S}_1)$ and $\nu(\mathcal{S}_2)$.

Now we introduce some notation for the remainder of this proof. As the set \mathcal{S} , \mathcal{S}_1 , and \mathcal{S}_2 change in each round, let $\mathcal{S}(r)$, $\mathcal{S}_1(r)$, and $\mathcal{S}_2(r)$ denote their values in round r for $r = 1, \dots, \log_2(n)$. Define $\mathcal{A}_1(r)$ and $\mathcal{A}_2(r)$ similarly. We stop $\mathcal{A}_1(r)$ if $\mathcal{A}_2(r)$ terminates and vice versa.

Let T_r denote the random variable of the total number of samples of drawn in round r . Let $T_{r,1}$ be the number of samples drawn by $\mathcal{A}_1(r)$, and $T_{r,2}$ be the number of samples drawn by $\mathcal{A}_2(r)$.

Next, define $\mathcal{S}^*(r)$ be the set in $\{\mathcal{S}_1(r), \mathcal{S}_2(r)\}$ that contains i^* , i.e. let $\mathcal{S}^*(r)$ denote $\mathcal{S}_1(r)$ if $i^* \in \mathcal{S}_1(r)$ and $\mathcal{S}_2(r)$ otherwise for all r . Similarly, let $\mathcal{A}^*(r)$ denote $\mathcal{A}_1(r)$ if $i^* \in \mathcal{S}_1(r)$ and $\mathcal{A}_2(r)$ otherwise. Define T_{r,\mathcal{A}^*} to be the random number of samples given to $\mathcal{A}^*(r)$. Hence, $T_{r,\mathcal{A}^*} = T_{r,1}$ or $T_{r,\mathcal{A}^*} = T_{r,2}$.

By Step 2, $\mathcal{A}_1(r)$ and $\mathcal{A}_2(r)$ are run in parallel. Hence, $T_{r,1} = T_{r,2}$ deterministically. Furthermore, $T_r = T_{r,1} + T_{r,2}$ deterministically. Therefore,

$$T_{r,\mathcal{A}^*} = \frac{T_{r,1} + T_{r,2}}{2} = \frac{T_r}{2}.$$

Therefore, the expected number of samples in round r , taken over the randomness in the set $\mathcal{S}^*(r)$, the randomness in the instance $\nu(\mathcal{S}^*(r))$, and any randomness in $\mathcal{A}^*(r)$ is

$$\begin{aligned}
\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r), \nu(\mathcal{S}^*(r))} [T_r] &= 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r), \nu(\mathcal{S}^*(r))} [T_{r,\mathcal{A}^*}] \\
&= 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r)} [\mathbb{E}_{\nu(\mathcal{S}^*(r))} [T_{r,\mathcal{A}^*} | \mathcal{S}^*(r)]] \\
&= 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r)} \left[\min \left(\frac{1}{32} \sum_{j \in \mathcal{S}^*(r) \setminus \{i^*\}} \frac{1}{\Delta_{i^*,j}^2}, N_{H_0}(\nu(\mathcal{S}^*(r)^c)) \right) \right]
\end{aligned}$$

$$\begin{aligned}
&\leq 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r)} \left[\frac{1}{32} \sum_{j \in \mathcal{S}^*(r) \setminus \{i^*\}} \frac{1}{\Delta_{i^*,j}^2} \right] \\
&= 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r-1)} \left[\mathbb{E}_{\mathcal{S}^*(r)} \left[\frac{1}{32} \sum_{j \neq i^*} \mathbb{1}[j \in \mathcal{S}^*(r)] \frac{1}{\Delta_{i^*,j}^2} \middle| \mathcal{S}^*(r-1) \right] \right] \\
&= 2\mathbb{E}_{\mathcal{S}^*(1), \dots, \mathcal{S}^*(r-1)} \left[\frac{1}{32} \cdot \left(\frac{1}{2}\right) \sum_{j \neq i^*} \mathbb{1}[j \in \mathcal{S}^*(r-1)] \frac{1}{\Delta_{i^*,j}^2} \right] \\
&\vdots \\
&= 2\mathbb{E}_{\mathcal{S}^*(1)} \left[\frac{1}{32} \cdot \left(\frac{1}{2}\right)^{r-1} \sum_{j \neq i^*} \mathbb{1}[j \in \mathcal{S}^*(1)] \frac{1}{\Delta_{i^*,j}^2} \right] \\
&= \frac{1}{16} \cdot \left(\frac{1}{2}\right)^r \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}.
\end{aligned}$$

Therefore, we may bound the expected total number of samples for the above binary search algorithm to return i^* as

$$\mathbb{E} \left[\sum_{r=1}^{\log_2(n)} T_r \right] = \sum_{r=1}^{\log_2(n)} \mathbb{E}[T_r] \leq \frac{1}{16} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2} \sum_{r=1}^{\log_2(n)} \left(\frac{1}{2}\right)^r \leq \frac{1}{16} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}.$$

However, this contradicts Theorem D.2 for $\delta < 1/16$. Hence no such algorithm \mathcal{A} exists and any algorithm to declare H_1 on instance ν requires at least $\frac{1}{32} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}$ samples in expectation. \square

Theorem D.6. Fix n , β , and $\delta < 1/16$ and consider an instance ν . If H_0 is true on ν , any algorithm requires at least

$$\sum_{j=1}^n \frac{2}{(\beta - \mu_j)^2} \log \left(\frac{1}{2.4\delta} \right)$$

samples in expectation to perform the β -isolated Hypothesis Test. If H_1 is true on ν , any algorithm requires at least

$$\frac{1}{4\beta^2} \log \left(\frac{1}{2.4\delta} \right) + \frac{1}{64} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}$$

samples in expectation to perform the β -isolated Hypothesis Test.

Proof. If H_0 is true for ν , the result follows immediately from Lemma D.4. Otherwise, assume H_1 is true for ν and let i^* be the single distribution such that $\rho_{i^*} = \mathcal{N}(\beta, 1)$. Similar to the proof of Lemma D.1, one may consider an alternate instance ν' where $\rho_{i^*} = \mathcal{N}(-\beta, 1)$ and all other distributions are unchanged. Therefore, on ν' , H_0 is true and any algorithm that is correct with probability at least $1 - \delta$ must be able to distinguish between these two instances. By Lemma 1 of [6], any algorithm that is correct with probability at least $1 - \delta$ must therefore sample $i^* \frac{1}{2\beta^2} \log \left(\frac{1}{2.4\delta} \right)$ times in expectation. Combining this with the result of Lemma D.5, any algorithm that is correct with probability at least $1 - \delta$ must collect at least

$$\max \left\{ \frac{1}{32} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}, \frac{1}{2\beta^2} \log \left(\frac{1}{2.4\delta} \right) \right\} \geq \frac{1}{4\beta^2} \log \left(\frac{1}{2.4\delta} \right) + \frac{1}{64} \sum_{j \neq i^*} \frac{1}{\Delta_{i^*,j}^2}$$

samples in expectation. \square

D.3 Step 3: Reducing ALL- ϵ to isolated instance detection

In this section, we prove that for any instance ν for ALL- ϵ such that $|G_{\beta_\epsilon}(\nu)| = 1$ requires at least $O\left(\sum_{i=2}^n \frac{1}{\Delta_i^2}\right)$ samples in expectation. To do so, we prove a reduction from finding all ϵ -good arms to the β -Isolated Hypothesis Testing. In particular, we show that if one has a generic method to find all ϵ -good arms (with slack $\gamma = 0$), then one may use this to develop a method to perform the β -Isolated Hypothesis Test. Therefore, lower bounds on the this test apply to the problem of finding all ϵ -good arms as well.

Lemma D.7. Fix $\delta \leq 1/16$, $n \geq 2/\delta$, $\epsilon > 0$, $\beta \in (0, \epsilon/2)$. Let ν be an instance of n arms such that the i^{th} is distributed as $\mathcal{N}(\mu_i, 1)$, $|G_{2\beta_\epsilon}| = 1$, and there exists an arm in G_ϵ^c such that $\mu_1 - \epsilon - \mu_i = \beta$. Select a permutation $\pi : [n] \rightarrow [n]$ uniformly from the set of $n!$ permutations, and consider the permuted instance $\pi(\nu)$. Any algorithm that returns $G_\epsilon(\pi(\nu))$ on $\pi(\nu)$ with correctly probability at least $1 - \delta$ requires at least

$$\frac{1}{64} \sum_{i=2}^n \frac{1}{\Delta_i^2} + \frac{1}{4\beta_\epsilon^2} \log\left(\frac{1}{2.4\delta}\right)$$

samples in expectation, where the expectation is taken jointly over the randomness in ν and π .

Proof. Fix $0 < \delta < 1/16$, $n > 2/\delta$, $\epsilon > 0$, $0 < \beta < \epsilon/2$, and an arbitrary constant $c \in \mathbb{R}$. Consider a given instance $\nu = \{\rho_1, \dots, \rho_n\}$ such that $\mu_1 \in \{-\beta, \beta\}$, and $\mu_2, \dots, \mu_n < -\beta$. We wish to perform the β -isolated hypothesis test on $\pi(\nu)$. Assume for contradiction that there exists a generic algorithm $\mathcal{A}(\nu', \epsilon, \delta)$ such that if given a generic instance ν' where $|G_{2\beta_\epsilon}(\nu')| = 1$, it returns $G_\epsilon(\nu')$ with probability at least $1 - \delta$ in at most $\frac{1}{64} \sum_{i=2}^n \frac{1}{\Delta_i^2}$ samples where μ'_1 is the largest mean in ν' . Consider the following procedure that uses \mathcal{A} to perform the hypothesis test:

Algorithm 4 Using All- ϵ for β -isolated hypothesis test

Require: $\delta > 0$, $\epsilon > 0$, $0 < \beta$, instance $\pi(\nu)$, constant c , and algorithm \mathcal{A}

- 1: **Step 1:** Choose an index $\hat{i} \in [n]$ uniformly
- 2: **Step 2:** Let ν' be the instance

$$\nu' = \begin{cases} \rho_{\pi(i)} + c & \text{if } i \neq \hat{i} \\ \mathcal{N}(c - \epsilon, 1) & \text{if } i = \hat{i} \end{cases}$$

- 3: **Step 3:** $G = \mathcal{A}(\nu', \epsilon, \delta/2)$
 - 4: **if** $\hat{i} \in G$ **then:**
 - 5: Declare H_0 and terminate
 - 6: **else**
 - 7: Declare H_1 and terminate
-

Note that as $n \geq 2/\delta$, $\mathbb{P}(\hat{i} = \pi(1)) \leq \delta/2$. The method replaces $\rho_{\hat{i}}$ with $\mathcal{N}(c - \epsilon, 1)$. All other means μ_i are shifted up by c . The test then runs \mathcal{A} on this new instance ν' with failure probability $\delta/2$. If H_0 is true on $\pi(\nu)$, all distributions have means less than $-\beta$, and \hat{i} therefore is ϵ -good on instance ν' . If H_1 is true on $\pi(\nu)$, then $\rho_{\pi(1)} = \mathcal{N}(\beta, 1)$ and \hat{i} is not ϵ -good on instance ν' . This method correctly performs the test if $\hat{i} \neq \pi(1)$ and \mathcal{A} does not fail, the joint event of which occurs with probability at most 2δ . Therefore, this test is correct with probability at least $1 - \delta$.

Let $\mathcal{T}_{\mathcal{A}}(\nu')$ denote the random variable of the number of samples drawn by \mathcal{A} on instance ν' and let T denote the random variable of the total number of samples drawn by this procedure before it terminates and declares H_0 or H_1 on ν' . Therefore, $\mathbb{E}_{\pi, \nu}[T] = \mathbb{E}_{\pi, \nu}[\mathcal{T}_{\mathcal{A}}(\nu')]$.

By Lemma 1 of [9], averaging over all permutations is equivalent to first permuting the instance ν and then passing it to \mathcal{A} and undoing the permutation when returning the answer. We therefore assume that \mathcal{A} is *symmetric* in that its expected sample complexity of \mathcal{A} is invariant to the permutation π . Otherwise, we may use \mathcal{A} to form a symmetric algorithm. Therefore, $\mathbb{E}_{\pi, \nu}[T] = \mathbb{E}_{\pi, \nu}[\mathcal{T}_{\mathcal{A}}(\nu')] =$

$\mathbb{E}_\nu[\mathcal{T}_A(\nu')]$. By Theorem D.6, if H_1 is true,

$$\mathbb{E}_{\pi, \nu}[T] \geq \frac{1}{64} \sum_{i=2}^n \frac{1}{\Delta_i^2} + \frac{1}{4\beta^2} \log \left(\frac{1}{2.4\delta} \right).$$

Hence,

$$\mathbb{E}_\nu[\mathcal{T}_A(\nu')] \geq \frac{1}{64} \sum_{i=2}^n \frac{1}{\Delta_i^2} + \frac{1}{4\beta^2} \log \left(\frac{1}{2.4\delta} \right).$$

Lastly, as the constant c was chosen arbitrarily, and β is an number in $(0, \epsilon/2)$ this argument applies to any ALL- ϵ instance ν' such that $\beta_\epsilon \in (0, \epsilon/2)$ and $|G_{2\beta_\epsilon}| = 1$ for n appropriate choice of c . \square

With the above proof, we restate the following moderate confidence lower bound on the sample complexity of returning all ϵ -good stated in Section 4. In particular, this bound highlights *moderate confidence* terms that are independent of δ . Moderate confidence terms have been studied in works such as [9, 25]. Despite being independent of δ , these terms can have important effects in real world scenarios. The following bound demonstrates that there are instances for which moderate confidence terms are necessary for finding all ϵ -good arms. Moderate confidence terms likewise appear in the upper bound of the complexity of FAREAST, Theorem 4.2.

Theorem D.8. Fix $\delta \leq 1/16$, $n \geq 2/\delta$, and $\epsilon > 0$. Let ν be an instance of n arms such that the i^{th} is distributed as $\mathcal{N}(\mu_i, 1)$, $|G_{2\beta_\epsilon}| = 1$, and $\beta_\epsilon < \epsilon/2$. Select a permutation $\pi : [n] \rightarrow [n]$ uniformly from the set of $n!$ permutations, and consider the permuted instance $\pi(\nu)$. Any algorithm that returns $G_\epsilon(\pi(\nu))$ on $\pi(\nu)$ with correctly probability at least $1 - \delta$ requires at least

$$c_2 \sum_{i=1}^n \max \left(\frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \log \left(\frac{1}{2.4\delta} \right) + c_2 \sum_{i=1}^n \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2}$$

samples in expectation over the randomness in ν and π for a universal constant c_2 .

Proof. We may equivalently consider the same instance with all means shifted down by $\epsilon - 2\beta$ since a method for that instance could be used to return all ϵ good arms in the stated instance. By Lemma D.7, $c_2 \frac{n}{\beta^2}$ samples are necessary in expectation. By Theorem 2.1,

$$2 \sum_{i=1}^n \max \left(\frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \log \left(\frac{1}{2.4\delta} \right)$$

samples are necessary in expectation. By Lemma D.7,

$$\begin{aligned} \frac{1}{64} \sum_{i=2}^n \frac{1}{\Delta_i^2} + \frac{1}{4\beta_\epsilon^2} \log \left(\frac{1}{2.4\delta} \right) &\geq \frac{1}{64} \sum_{i=2}^n \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} + \frac{1}{4\beta_\epsilon^2} \log \left(\frac{1}{2.4\delta} \right) \\ &\geq \frac{1}{64} \sum_{i=1}^n \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \end{aligned}$$

samples are necessary in expectation taken over the randomness in the permutation and in the instance. In particular, the maximum and therefore the average is a valid bound. Therefore, any algorithm requires

$$\sum_{i=1}^n \max \left(\frac{1}{(\mu_1 - \epsilon - \mu_i)^2}, \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \log \left(\frac{1}{2.4\delta} \right) + \frac{1}{128} \sum_{i=1}^n \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2}$$

samples in expectation. \square

E An optimal method for finding all additive and multiplicative ϵ -good arms

E.1 The FAREAST Algorithm

Below, we present an algorithm called FAREAST (Fast Arm Removal Elimination Algorithm for a Sampled Threshold) that achieves the lower bound when $\gamma = 0$. Similar to (ST)², it relies on

anytime-correct confidence widths, $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$. The algorithm proceeds in rounds, and creates a filter for *good* arms and a filter for *bad* arms. The good filter detects arms in G_ϵ of M_ϵ and adds them to a set G_k . Similarly, the bad filter detects arms in G_ϵ^c or M_ϵ^c and adds them to a set B_k . At any given time, we may represent the set of arms that have *not* been declared as either in G_ϵ/M_ϵ or $G_\epsilon^c/M_\epsilon^c$ as $(G_k \cup B_k)^c$. In either the additive or multiplicative case, the algorithm terminates when it can certify that $G_\epsilon \subset G_k$ and $G_k \cap G_{\epsilon+\gamma}^c = \emptyset$ or $M_\epsilon \subset G_k$ and $G_k \cap M_{\epsilon+\gamma}^c = \emptyset$, respectively—i.e., when G_k contains all additive or multiplicative ϵ -good arms and none worse than $(\epsilon + \gamma)$ -good.

In each round, the bad filter uses `MedianElimination` [12] which given an instance ν , a value of ϵ , and a failure probability κ , returns an ϵ -good arm with probability at least $1 - \kappa$. In the k^{th} round, for an arm i in $(G_k \cup B_k)^c$, the bad filter uses `MedianElimination` to find a 2^{-k} good arm i_k with failure probability $\kappa = O(1)$ and then samples both arms i and i_k $\tilde{O}(2^{2k} \log(1/\delta))$ times. Let $\hat{\mu}_i$ and $\hat{\mu}_{i_k}$ denote the empirical means. For instance, in the additive case, if $\hat{\mu}_{i_k} - \hat{\mu}_i \geq \epsilon + 2^{-k+1}$, we may declare that $i \in G_\epsilon^c$, and the bad filter adds i to the set B_k . This allows the bad filter to commit to a single arm and sample it sufficiently to remove arms in G_ϵ^c .

The good filter is a simple elimination scheme. It maintains an upper bound U_t and lower bound L_t on $\mu_1 - \epsilon$. If an arm's upper bound drops below L_t (line 20), the good filter eliminates that arm, otherwise, if an arm's lower bound rises above U_t (19), the good filter adds the arm to G_k , but only eliminates this arm if its upper bound falls below the highest lower bound. This ensures that μ_1 is never eliminated and U_t and L_t are always valid bounds. This scheme works as an independent algorithm and achieves the sample complexity as (ST)², though worse empirical performance. We analyze this method in Appendix E.5. Indeed, this gives an additional high probability guarantee on the number of samples drawn by FAREAST in both the **additive** and **multiplicative** regimes. As the sampling is split across rounds, the good filter always samples the least sampled arm, breaking ties arbitrarily. The number of samples given to the good filter in each round is such that both filters receive identically many samples. Note that this is a random quantity since the number of arms in $(G_k \cup B_k)^c$ in round k is random. Despite this, we prove a lower bound on the number of samples drawn per round which ensures the Good Filter always receives a positive number of samples in each round. Note that by design elimination only occurs when all arms in the active set have received equal numbers of samples. This is crucial as it prevents the good filter from over-sampling bad arms and vice versa. In our proof, we show that in some round, unknown to the algorithm, $G_k = G_\epsilon$, ie all good arms have been found, and this takes no more than $O(\sum_{i=1}^n \max\{(\mu_1 - \epsilon - \mu_i)^{-2}, (\mu_1 + \alpha_\epsilon - \mu_i)^{-2}\} \log(n/\delta))$ samples, matching the lower bound.

The algorithm stops on either of three conditions. First, if $G_k \cup B_k = [n]$, every arm has been declared as either in G_ϵ or G_ϵ^c (or M_ϵ or M_ϵ^c). Second, if $\mathcal{A} \subset G_k$, the Good Filter has found every arm in G_ϵ and FAREAST can terminate. This is the same stopping condition as EAST itself. In either case, FAREAST returns the set $G_k = G_\epsilon$ exactly. The third condition allows for γ slack. The good filter maintains upper and lower bounds U_t and L_t on the threshold in both the additive and multiplicative cases. In the additive case, if $U_t - L_t < \gamma/2$, then all arms in $G_{\epsilon+\gamma}^c$ have been added to B_k , and FAREAST may return $G_k \cup \mathcal{A}$. The condition for the multiplicative case is similar, though slightly more complicated. Throughout, we will use **red text** to denote pieces specific to the additive case and **blue text** to denote pieces specific to the multiplicative case.

```

1 FAREAST
2 Input:  $\epsilon, \delta$ , Instance  $\nu$ , slack  $\gamma \geq 0$ . If multiplicative,  $\epsilon \in (0, 1/2]$ 
3 Let  $G_0 = \emptyset$  be the set of arms declared as good and  $B_0 = \emptyset$  the set of arms declared as bad.
4 Let  $\mathcal{A} = [n]$  be the active set,  $N_i = 0$  track the total number of samples of arm  $i$  by the Good Filter.
5 Let  $t = 0$  denote the total number of times that line 19 is true in the Good Filter.
6 Let  $C_{\delta/2n}(t)$  be an anytime  $\delta/2n$ -correct confidence width on  $t$  samples.
7 Let  $H_{\text{ME}}(n, \epsilon, \kappa) = \lceil c' \frac{n}{\epsilon^2} \log(1/\kappa) \rceil$  be the complexity of MedianElimination.
8 for  $k = 1, 2, \dots$ 
9   Let  $\delta_k = \delta/2k^2$ ,  $\tau_k = \lceil 2^{2k+3} \log\left(\frac{8n}{\delta_k}\right) \rceil$ , Initialize  $G_k = G_{k-1}$  and  $B_k = B_{k-1}$ 
10  // Bad Filter: find bad arms in  $G_\epsilon^c$  or  $M_\epsilon^c$ 
11  Let  $i_k = \text{MedianElimination}(\nu, 2^{-k}, 1/16)$ , sample  $i_k$   $\tau_k$  times, and compute  $\hat{\mu}_{i_k}$ 
12  for  $i \notin G_{k-1} \cup B_{k-1}$ :
13    Sample  $\mu_i$   $\tau_k$  times and compute  $\hat{\mu}_i$ 
14    If  $\hat{\mu}_{i_k} - \hat{\mu}_i \geq \epsilon + 2^{-k+1}$  or  $(1 - \epsilon)\hat{\mu}_{i_k} - \hat{\mu}_i > 2^{-(k+1)}(2 - \epsilon)$ :
15      Add  $i$  to  $B_k$ 
16  // Good Filter: find good arms in  $G_\epsilon$  or  $M_\epsilon$ 
17  for  $s = 1, \dots, H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k \cdot (|G_{k-1} \cup B_{k-1}|^c + 1)$ :
18    Pull arm  $I_s \in \arg \min_{j \in \mathcal{A}} \{N_j\}$  and set  $N_{I_s} \leftarrow N_{I_s} + 1$ .
19    if  $\min_{j \in \mathcal{A}} \{N_j\} = \max_{j \in \mathcal{A}} \{N_j\}$ :
20       $t = t + 1$ 
21      For  $i \in \mathcal{A}$  denote  $\hat{\mu}_i(t)$  the average of the first  $t$  samples of arm  $i$ .
22      Let  $U_t = \max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/2n}(t) - \epsilon$  or  $U_t = (1 - \epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/2n}(t))$ 
23      Let  $L_t = \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t) - \epsilon$  or  $L_t = (1 - \epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t))$ 
24      for  $i \in \mathcal{A}$ :
25        if  $\hat{\mu}_i(t) - C_{\delta/2n}(t) \geq U_t$ :
26          Add  $i$  to  $G_k$ 
27        if  $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq L_t$ : // Bad arms are removed from  $\mathcal{A}$ 
28          Remove  $i$  from  $\mathcal{A}$ 
29        if  $i \in G_k$  and  $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$ : // Good arms removed
30          Remove  $i$  from  $\mathcal{A}$ 
31      If  $\mathcal{A} \subset G_k$  or  $G_k \cup B_k = [n]$ :
32        Output: the set  $G_k$  // Stopping condition for returning  $G_\epsilon$  exactly.
33      If  $U_t - L_t < \frac{1}{2}\gamma$  or  $U_t - L_t < \frac{\gamma}{2-\epsilon}L_t$ :
34        Output: the set  $\mathcal{A} \cup G_k$  // Stopping condition for  $\gamma > 0$ .

```

Remark 1. Note that the active set \mathcal{A} defined in line 4 of FAREAST is only used and updated internally by the Good Filter. In particular, it is not necessarily true that $(G_k \cup B_k)^c = \mathcal{A}$. Furthermore, a bad arm $i \in G_\epsilon^c$ maybe removed from \mathcal{A} even though it is not in B_k and vice versa as the Good Filter only seeks to detect good arms in G_ϵ and the Bad Filter only seeks to detect arms in G_ϵ^c . The same is true in the multiplicative case.

Remark 2. It is possible that when the loop in line 17 finishes in any given round, some arms in \mathcal{A} have received more samples than others. Because $I_s \in \arg \min_{j \in \mathcal{A}} \{N_j\}$ in line 18, this difference is no more than 1, and the arms with fewer samples are the first to be sampled in the next round. The condition on line 19 ensures that all arms have equal numbers of samples by the Good Filter (e.g., the N_i 's) when the Good Filter identifies good arms or eliminates arms from \mathcal{A} .

Now, we restate Theorem 4.2 for reference.

Theorem E.1. Fix $0 < \epsilon$, $0 < \delta < 1/8$, slack $\gamma \in [0, 8]$ and an instance ν of n arms such that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i . There exists an event E such that $\mathbb{P}(E) \geq 1 - \delta$, and on E , FAREAST terminates and returns G such that $G_\epsilon \subset G \subset G_{\epsilon+\gamma}$ in at most

$$c_4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\ \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right), \\ \left. \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \right\}, \\ \left. \frac{1}{\gamma^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta\gamma^2} \right) \right) \right\}$$

samples for a constant c_4 . Furthermore

$$\begin{aligned} \mathbb{E}[\mathbb{1}_E T] \leq c_3 \sum_{i \in G_\epsilon} \max & \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \\ & \left. \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right) \right\} \\ & + c_3 \sum_{i \in G_\epsilon^c} \frac{n}{(\mu_1 - \epsilon - \mu_i)^2} + \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right) \end{aligned}$$

for a sufficiently large constant c_3 where T denotes the number of samples.

Next, we present a theorem bounding the sample complexity of FAREAST for returning multiplicative ϵ -good arms. Recall that $\tilde{\alpha}_\epsilon := \min_{i \in M_\epsilon} \mu_i - (1 - \epsilon)\mu_1$ and $\tilde{\beta}_\epsilon := \min_{i \in M_\epsilon^c} (1 - \epsilon)\mu_1 - \mu_i$, the distance for the smallest good arm and best arm that is not good to the threshold $(1 - \epsilon)\mu_1$.

Theorem E.2. Fix $\epsilon \in (0, 1/2]$, $\gamma \in [0, \min(1, 6/\mu_1))$, $0 < \delta < 1/8$ and an instance ν of n arms such that $\max(\Delta_i, |\epsilon\mu_1 - \Delta_i|) \leq 6$. Assume that the highest mean is non-negative, i.e., $\mu_1 \geq 0$. There exists an event E such that $\mathbb{P}(E) \geq 1 - \delta$, and on E , FAREAST terminates and returns G such that $M_\epsilon \subset G \subset M_{\epsilon+\gamma}$ in at most

$$\begin{aligned} c_5 \sum_{i=1}^n \min & \left\{ \max \left\{ \frac{1}{((1 - \epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta((1 - \epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right. \\ & \frac{1}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i)^2} \right) \right), \\ & \left. \left. \frac{1}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1 - \epsilon} - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1 - \epsilon} - \mu_i)^2} \right) \right) \right\} \right\}, \\ & \frac{(1 - \epsilon + \gamma)^2}{\gamma^2 \mu_1^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{(1 - \epsilon + \gamma)^2 n}{\delta \gamma^2 \mu_1^2} \right) \right) \end{aligned}$$

samples for a sufficiently large constant c_5 . Furthermore

$$\begin{aligned} \mathbb{E}[\mathbb{1}_E T] \leq c_6 \sum_{i=1}^n \max & \left\{ \frac{1}{((1 - \epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta((1 - \epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \\ & \left. \frac{1}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1 - \epsilon} - \mu_i)^2} \right) \right) \right\} \\ & + c_6 \sum_{i \in M_\epsilon^c} \frac{n}{((1 - \epsilon)\mu_1 - \mu_i)^2} \end{aligned}$$

for a sufficiently large constant c_6 , where T denotes the number of samples.

E.2 Key ideas of the proof

The proof revolves around a central idea: there is an event in unknown round K_{Good} in which the final arm from G_ϵ or M_ϵ is added to G_k . We may split the total number of samples drawn as the number taken through round K_{Good} and the number taken from $K_{\text{Good}} + 1$ until termination if the algorithm does not terminate in round K_{Good} . Note that the Good filter and Bad filter are given the same number of samples in each round. The proof of FAREAST in the multiplicative regime is similar and deferred to Appendix E.4.

We begin by bounding the number of samples given to the Good filter when this event occurs that $G_k = G_\epsilon$. Next, since this happens at a random time within round K_{Good} , we bound the total number of additional samples in this round. Collectively, this gives us control over the number of samples drawn through round K_{Good} .

Next, we bound the number of samples from $K_{\text{Good}} + 1$ until termination. To do so, we analyze the expected number of samples drawn by the Bad filter before all arms in G_ϵ^c have been added to B_k . The total number of samples from $K_{\text{Good}} + 1$ until termination is no worse than twice this value. The proof is split into 12 steps and logically are organized as follows:

1. Step 0: We show that $G_k \subset G_\epsilon$ and $B_k \subset G_\epsilon^c$. In particular, this implies that $G_k \cup B_k = [n] \implies G_k = G_\epsilon$ so FAREAST terminates correctly.
2. Step 1: We split the total number of samples drawn by FAREAST into two sums that we will control individually.
3. Steps 2-4: We control the number of samples given to the Good filter before $G_k = G_\epsilon$.
4. Steps 5-6: Using the result of steps 2-4, we bound the total number of samples through round K_{Good} .
5. Steps 7-8: We use the result of step 6 to bound the total *expected* number of samples drawn by FAREAST, simplifying slightly in the process.
6. Step 9: We bound the number of samples that the Bad filter draws in adding a single bad arm to B_k .
7. Step 10: Repeating the argument in step 9, for every $i \in G_\epsilon^c$, we bound the total number of samples from round $K_{\text{Good}} + 1$ until termination. We finish by combining the bound on the number of samples drawn through K_{Good} with the bound from $K_{\text{Good}} + 1$ until termination. This controls the expected sample complexity of FAREAST.
8. Step 11: We provide a high probability bound on the sample complexity of FAREAST.

E.3 Proof of Theorem 4.2, FAREAST in the **additive** regime

Proof. Notation for the proof: Throughout, recall $\Delta_i = \mu_1 - \mu_i$. Recall that t counts the number of times the conditional in line 19 is true. By Line 19 of FAREAST, all arms in \mathcal{A} have received t samples when the loop in line 23 is executed for the t^{th} time. Within any round k , let $\mathcal{A}(t)$ and $G_k(t)$ denote the sets \mathcal{A} and G_k at this time since both sets can change in lines 27 and 29 and 25 respectively. Let t_k denote the maximum value of t in round k . By Lines 18 and 19 of FAREAST, the total number of samples given to the good filter when the conditional in line 19 is true for the t^{th} time is $\sum_{s=1}^t |\mathcal{A}(s)|$.

For $i \in G_\epsilon$, let T_i denote the random variable of the number of times arm i is sampled by the good filter before it is added to G_k in Line 25. For $i \in G_\epsilon^c$, let T_i denote the random variable of the number of times arm i is sampled by the good filter before it is removed from \mathcal{A} in Line 27. For any arm i , let T'_i denote the random variable of the number of times i is sampled by the good filter before $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$. Lastly, let T_γ denote the random variable of the number of times any arm is sampled by the good filter before $U_t - L_t < \gamma/2$.

Define the event

$$\mathcal{E}_1 = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/2n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\begin{aligned} \mathbb{P}(\mathcal{E}_1^c) &= \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/2n}(t) \right) \\ &\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/2n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{2n} = \frac{\delta}{2} \end{aligned}$$

Next, recall that $\hat{\mu}_i(t)$ denotes the empirical average of t samples of ρ_i . Consider the event,

$$\mathcal{E}_2 = \bigcap_{i \in G_\epsilon} \bigcap_{k \in \mathbb{N}} |(\hat{\mu}_{i_k}(\tau_k) - \hat{\mu}_i(\tau_k)) - (\mu_{i_k} - \mu_i)| \leq 2^{-k}$$

By Hoeffding's inequality,

$$\mathbb{P}(|(\hat{\mu}_j(\tau_k) - \hat{\mu}_i(\tau_k)) - (\mu_j - \mu_i)| > 2^{-k} | i_k = j) \leq \frac{\delta}{4nk^2}.$$

Then

$$\begin{aligned} & \mathbb{P}(|(\hat{\mu}_j(\tau_k) - \hat{\mu}_i(\tau_k)) - (\mu_j - \mu_i)| > 2^{-k}) \\ &= \sum_{j=1}^n \mathbb{P}(|(\hat{\mu}_j(\tau_k) - \hat{\mu}_i(\tau_k)) - (\mu_j - \mu_i)| > 2^{-k} | i_k = j) \mathbb{P}(i_k = j) \\ &\leq \frac{\delta}{4nk^2} \sum_{j=1}^n \mathbb{P}(i_k = j) \\ &= \frac{\delta}{4nk^2} \end{aligned}$$

Therefore, union bounding over the rounds $k \in \mathbb{N}$, $\mathbb{P}(\mathcal{E}_2^c) \leq \sum_{i \in G_\epsilon} \sum_{k=1}^{\infty} \frac{\delta}{4nk^2} \leq \frac{\delta}{2}$. Hence, $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1 - \delta$.

E.3.1 Step 0: Correctness.

On $\mathcal{E}_1 \cap \mathcal{E}_2$, first we prove that if there exists a random round k at which $G_k \cup B_k = [n]$ then $G_k = G_\epsilon$. Additionally, we prove that on $\mathcal{E}_1 \cap \mathcal{E}_2$, if $\mathcal{A} \subset G_k$, then $G_k = G_\epsilon$. Therefore, for either stopping condition for FAREAST in line 31, on the event $\mathcal{E}_1 \cap \mathcal{E}_2$, FAREAST correctly returns the set G_ϵ .

Claim 0: On $\mathcal{E}_1 \cap \mathcal{E}_2$, for all $k \in \mathbb{N}$, $G_k \subset G_\epsilon$.

Proof. Firstly we show $1 \in \mathcal{A}$ for all $t \in \mathbb{N}$, namely the best arm is never removed from \mathcal{A} . Note for any i

$$\hat{\mu}_1 + C_{\delta/2n}(t) \geq \mu_1 \geq \mu_i \geq \hat{\mu}_i(t) - C_{\delta/2n}(t) > \hat{\mu}_i(t) - C_{\delta/2n}(t) - \epsilon.$$

In particular this shows, $\hat{\mu}_1 + C_{\delta/2n}(t) > \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/2n}(t) - \epsilon = L_t$ and $\hat{\mu}_1 + C_{\delta/2n}(t) \geq \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/2n}(t)$ showing that 1 will never exit \mathcal{A} in line 28.

Secondly, we show that at all times t , $\mu_1 - \epsilon \in [L_t, U_t]$. By the above, since μ_1 never leaves \mathcal{A} ,

$$U_t = \max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/2n}(t) - \epsilon \geq \hat{\mu}_1(t) + C_{\delta/2n}(t) - \epsilon \geq \mu_1 - \epsilon$$

and for any i ,

$$\mu_1 - \epsilon \geq \mu_i - \epsilon \geq \hat{\mu}_i(t) - C_{\delta/2n}(t) - \epsilon$$

Hence $\mu_1 - \epsilon \geq \max_i \hat{\mu}_i(t) - C_{\delta/2n}(t) - \epsilon = L_t$.

Next, we show that $G_k \subset G_\epsilon$ for all $k \geq 1, t \geq 1$. Suppose not. Then $\exists k, t \in \mathbb{N}$ and $\exists i \in G_\epsilon^c \cap G_k(t)$ such that,

$$\mu_i \geq \hat{\mu}_i(t) - C_{\delta/2n}(t) \geq U_t \geq \mu_1 - \epsilon > \mu_i,$$

with the last inequality following from the previous assertion, giving a contradiction. \square

Claim 1: On $\mathcal{E}_1 \cap \mathcal{E}_2$, for all $k \in \mathbb{N}$, $B_k \subset G_\epsilon^c$.

Proof. Next, we show $B_k \subset G_\epsilon^c$. Suppose not. Either a good arm was added to the bad set by the bad filter or by the good filter. First, consider the case, that the bad filter added an arm in G_ϵ to B_k for some k . By definition, $B_0 = \emptyset$ and $B_{k-1} \subset B_k$ for all k . Then there must exist $k \in \mathbb{N}$ and an $i \in G_\epsilon$ such that $i \in B_k$ and $i \notin B_{k-1}$. Following line 14 of the algorithm, this occurs if and only if

$$\hat{\mu}_{i_k} - \hat{\mu}_i \geq \epsilon + 2^{-k+1}.$$

On the event \mathcal{E}_2 , the above implies

$$\mu_{i_k} - \mu_i + 2^{-k} \geq \epsilon + 2^{-k+1},$$

and simplifying, we see that $\epsilon + 2^{-k} \leq \mu_{i_k} - \mu_i \leq \mu_1 - \mu_i$ which contradicts the assertion that $i \in G_\epsilon$.

Next, consider the case that the good filter incorrectly adds a good arm $i \in G_\epsilon$ to B_k in some round k . Then there must be a $t \in \mathbb{N}$ such that.

$$\mu_i \stackrel{\mathcal{E}_1}{\leq} \hat{\mu}_i + C_{\delta/2n}(t) < L_t \stackrel{\mathcal{E}_1}{\leq} \mu_1 - \epsilon$$

which contradicts $i \in G_\epsilon$. Hence, in both cases $B_k \subset G_\epsilon^c$ for all k . \square Combining the above claims, we see that $\mathcal{E}_1 \cap \mathcal{E}_2$ implies $(G_k \cup B_k = [n])$ and $G_k \cap B_k = \emptyset \implies G_k = G_\epsilon$. Since $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1 - \delta$, if FAREAST terminates, with probability at least $1 - \delta$, it correctly returns the set G_ϵ .

Claim 2: Next, we show that on \mathcal{E}_1 , $G_\epsilon \subset \mathcal{A}(t) \cup G(t)$ for all $t \in \mathbb{N}$.

In particular this implies that if $\mathcal{A} \subset G$, then $G_\epsilon \subset G$. Combining this with the previous claim gives $G \subset G_\epsilon \subset G$, hence $G = G_\epsilon$. On this condition, FAREAST terminates by line 33 and returns the set $\mathcal{A} \cup G = G$. Note that by definition, $G_\epsilon \subset G_{(\epsilon+\gamma)}$ for all $\gamma \geq 0$. Therefore FAREAST terminates correctly on this condition.

Proof. Suppose for contradiction that there exists $i \in G_\epsilon$ such that $i \notin \mathcal{A}(t) \cup G(t)$. This occurs only if i is eliminated in line 28. Hence, there exists a $t' \leq t$ such that $\hat{\mu}_i(t') + C_{\delta/n}(t') < L_{t'}$. Therefore, on the event \mathcal{E}_1 ,

$$\mu_1 - \epsilon \stackrel{\mathcal{E}_1}{\geq} L_{t'} = \max_{j \in \mathcal{A}} \hat{\mu}_j(t') - C_{\delta/n}(t') - \epsilon > \hat{\mu}_i(t') + C_{\delta/n}(t') \stackrel{\mathcal{E}_1}{\geq} \mu_i$$

which contradicts $i \in G_\epsilon$. \square

Claim 3: Finally, we show that on \mathcal{E}_1 , if $U_t - L_t \leq \gamma/2$, then $\mathcal{A} \cup G \subset G_{(\epsilon+\gamma)}$.

Combining with Claim 3 that $G_\epsilon \subset \mathcal{A} \cup G$, if FAREAST terminates on this condition by line 33, it does so correctly and returns all arms in G_ϵ .

Proof. Assume $U_t - L_t \leq \gamma/2$. Since all arms in $\mathcal{A}(t)$ have received exactly t samples, this implies that

$$\left(\max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t) + C_{\delta/n}(t) - \epsilon \right) - \left(\max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon \right) = 2C_{\delta/n}(t) \leq \gamma/2.$$

Suppose for contradiction that there exists $i \in G_{(\epsilon+\gamma)}^c$ such that $i \in \mathcal{A} \cup G$. Since $G_\epsilon \cap G_{(\epsilon+\gamma)}^c = \emptyset$ and we have previously shown that $G(t) \subset G_\epsilon$ for all t , we have that $i \in \mathcal{A} \setminus G$. Therefore, by the condition in line 27, $\hat{\mu}_i(t) + C_{\delta/n}(t) \geq L_t$. Hence, $\mu_i + 2C_{\delta/n}(t) \stackrel{\mathcal{E}_1}{\geq} \hat{\mu}_i(t) + C_{\delta/n}(t) \geq L_t$. By assumption, we have that $U_t - \gamma/2 \leq L_t$, and the event \mathcal{E}_1 implies that $U_t \geq \mu_1 - \epsilon$. Therefore, $\mu_i + 2C_{\delta/n}(t) \geq U_t - \gamma/2 \geq \mu_1 - \epsilon - \gamma/2$. Combining this with the inequality $2C_{\delta/n} \leq \gamma/2$, we have that

$$\gamma \geq 2C_{\delta/n}(t) + \gamma/2 \geq \mu_1 - \epsilon - \mu_i \stackrel{i \in G_{(\epsilon+\gamma)}^c}{>} \gamma$$

which is a contradiction. \square

E.3.2 Step 1: An expression for the total number of samples drawn and introducing several helper random variables

Next, we write an expression for the total number of samples drawn by FAREAST. In particular, we introduce two sums that we will spend the remainder of the proof controlling. Additionally, we show that the conditional in line 19 in the good filter is true at least once in each round. Based on this, we more precisely define the random variables T_i and T'_i introduced in the notation section in subsection E.3. Additionally, we introduce the time T_γ at which $U_t - L_t < \frac{1}{2}\gamma$.

Recall that the largest value of t in round k is denoted t_k . Let E_k^γ be the event that $U_t - L_t \geq \gamma/2$ for all t in round k :

$$E_k^\gamma := \{U_t - L_t \geq \gamma/2 : t \in (t_{k-1}, t_k]\}.$$

Note that if E_{k-1}^γ is false, then FAREAST terminates in round $k - 1$ by line 33. We may write the total number of samples drawn by the algorithm as

$$T = \sum_{k=1}^{\infty} 2\mathbb{1} [\mathcal{A} \not\subset G_{k-1} \text{ and } G_{k-1} \cup B_{k-1} \neq [n] \text{ and } E_{k-1}^\gamma]$$

$$(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)$$

Deterministically, $\mathbb{1} [\mathcal{A} \not\subset G_{k-1} \text{ and } G_{k-1} \cup B_{k-1} \neq [n] \text{ and } E_{k-1}^\gamma] \leq \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]]$

Applying this,

$$\begin{aligned} T &\leq \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ &= \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} \neq G_\epsilon] \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \end{aligned} \quad (10)$$

$$+ \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} = G_\epsilon] \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \quad (11)$$

In round k , line 18 of the Good Filter, whereby an arm is sampled, is evaluated

$$(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \geq (H_{\text{ME}}(n, 2^{-k}, 1/16) + 2\tau_k) \geq n$$

times since $H_{\text{ME}}(n, 2^{-k}, 1/16) \geq n$ for all k and $|(G_{k-1} \cup B_{k-1})^c| \geq 1$ unless $G_{k-1} \cup B_{k-1} = [n]$ which implies termination in round $k-1$. Each time line 18 is called, $N_{I_s} \leftarrow N_{I_s} + 1$. Since $|\arg \min_{j \in \mathcal{A}} \{N_j\}| \leq |\mathcal{A}| \leq n$, line 18 is called at most n times before $\min_{j \in \mathcal{A}} \{N_j\} = \max_{j \in \mathcal{A}} \{N_j\}$. When this occurs, the conditional in line 19 is true and $t \leftarrow t + 1$.

If $\min_{i \in \mathcal{A}(t)} \{N_i\} = \max_{i \in \mathcal{A}(t)} \{N_i\}$, then $N_i = t$ for any $i \in \mathcal{A}(t)$. By Step 0, only arms in G_ϵ are added to G_k . Therefore, T_i is defined as

$$T_i = \min \left\{ t : \begin{array}{ll} i \in G_k(t+1) & \text{if } i \in G_\epsilon \\ i \notin \mathcal{A}(t+1) & \text{if } i \in G_\epsilon^c \end{array} \right\} \stackrel{\mathcal{E}_1}{=} \min \left\{ t : \begin{array}{ll} \hat{\mu}_i - C_{\delta/2n}(t) \geq U_t & \text{if } i \in G_\epsilon \\ \hat{\mu}_i + C_{\delta/2n}(t) \leq L_t & \text{if } i \in G_\epsilon^c \end{array} \right\} \quad (12)$$

Define $T_i = \infty$ if this never occurs. Note that this may happen if FAREAST terminates due to the condition in line 32 that $U_t - L_t < \gamma/2$. Similarly, recall T'_i denotes the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$. Hence,

$$T'_i = \min \left\{ t : \hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/2n}(t) \right\} \quad (13)$$

Define $T'_i = \infty$ if this never occurs. Note that this may happen if FAREAST terminates due to the condition in line 32 that $U_t - L_t < \gamma/2$. Finally, we define the time T_γ such that $U_t - L_t < \frac{1}{2}\gamma$.

$$T_\gamma = \min \left\{ t : U_t - L_t < \frac{1}{2}\gamma \right\} \quad (14)$$

By design, no arm is sampled more than T_γ times by the good filter, controlling the cases that T_i or T'_i are infinite.

E.3.3 Step 2: Bounding T_i and T'_i for $i \in G_\epsilon$

Step 2a: For $i \in G_\epsilon$, we have that $T_i \leq h(0.25(\epsilon - \Delta_i), \delta/2n)$.

Proof. Note that, $4C_{\delta/2n}(t) \leq \mu_i - (\mu_1 - \epsilon)$, true when $t > h(0.25(\epsilon - \Delta_i), \frac{\delta}{2n})$, implies that for all j ,

$$\begin{aligned} \hat{\mu}_i(t) - C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\geq} \mu_i - 2C_{\delta/2n}(t) \\ &\geq \mu_1 + 2C_{\delta/2n}(t) - \epsilon \\ &\geq \mu_j + 2C_{\delta/2n}(t) - \epsilon \\ &\stackrel{\mathcal{E}_1}{\geq} \hat{\mu}_j(t) + C_{\delta/2n}(t) - \epsilon \end{aligned}$$

so in particular, $\hat{\mu}_i(t) - C_{\delta/2n}(t) \geq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/2n}(t) - \epsilon = U_t$. \square

Additionally, we define a time T_{\max} when all good arms have entered G_k .

Step 2b: Defining $T_{\max} := \min\{t : G_k(t) = G_\epsilon\} = \max_{i \in G_\epsilon} T_i$, we also have that $T_{\max} \leq h(0.25\alpha_\epsilon, \delta/2n)$ (in other words, if $t > h(0.25\alpha_\epsilon, \delta/2n)$ (i.e. line 23 has been run t times), then we have that $G_k(t) = G_\epsilon$).

Proof. Recall that $\alpha_\epsilon = \min_{i \in G_\epsilon} \mu_i - \mu_1 + \epsilon = \min_{i \in G_\epsilon} \epsilon - \Delta_i$. By Step 1a, $T_i \leq h(0.25(\epsilon - \Delta_i), \frac{\delta}{2n})$. Furthermore, $h(\cdot, \cdot)$ is monotonic in its first argument, such that if $0 < x' < x$, then $h(x', \delta) > h(x, \delta)$ for any $\delta > 0$. Therefore $T_{\max} = \max_{i \in G_\epsilon} T_i \leq \max_{i \in G_\epsilon} h(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}) = h(0.25\alpha_\epsilon, \frac{\delta}{2n})$. \square

Step 2c: For $i \in G_\epsilon$, we have that $T'_i \leq h(0.25\Delta_i, \delta/2n)$.

Proof. Note that $4C_{\delta/2n}(t) \leq \mu_1 - \mu_i$, true when $t > h(0.25\Delta_i, \frac{\delta}{2n})$, implies that

$$\begin{aligned} \hat{\mu}_i(t) + C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\leq} \mu_i + 2C_{\delta/2n}(t) \\ &\leq \mu_1 - 2C_{\delta/2n}(t) \\ &\stackrel{\mathcal{E}_1}{\leq} \hat{\mu}_1(t) - C_{\delta/2n}(t). \end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Hence, $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/2n}(t)$. \square

E.3.4 Step 3: Bounding T_i for $i \in G_\epsilon^c$

Next, we bound T_i for $i \in G_\epsilon^c$. $i \in G_\epsilon^c$ is eliminated from \mathcal{A} if it has received at least T_i samples.

Claim: $T_i \leq h(0.25(\epsilon - \Delta_i), \frac{\delta}{2n})$ for $i \in G_\epsilon^c$

Proof. Note that, $4C_{\delta/2n}(t) \leq \mu_1 - \epsilon - \mu_i$, true when $t > h(0.25(\epsilon - \Delta_i), \frac{\delta}{2n})$, implies that

$$\begin{aligned} \hat{\mu}_i(t) + C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\leq} \mu_i + 2C_{\delta/2n}(t) \\ &\leq \mu_1 - 2C_{\delta/2n}(t) - \epsilon \\ &\stackrel{\mathcal{E}_1}{\leq} \hat{\mu}_1(t) - C_{\delta/2n}(t) - \epsilon \end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Therefore $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t) - \epsilon = L_t$. \square

E.3.5 Step 4: bounding the total number of samples given to the good filter at time $t = T_{\max}$

Note that for a time $t = T$, the total number of samples given to the good filter is $\sum_{s=1}^T |\mathcal{A}(s)|$. Therefore, the total number of samples up to time T_{\max} is $\sum_{t=1}^{T_{\max}} |\mathcal{A}(t)|$.

Let $S_i = \min\{t : i \notin \mathcal{A}(t+1)\}$. Hence,

$$\sum_{t=1}^{T_{\max}} |\mathcal{A}(t)| = \sum_{t=1}^{T_{\max}} \sum_{i=1}^n \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \sum_{t=1}^{T_{\max}} \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \min\{T_{\max}, S_i\}$$

For arms $i \in G_\epsilon^c$, $S_i = T_i$ by definition. For $i \in G_\epsilon$, $S_i = \max(T_i, T'_i)$ by line 28 of the algorithm. Then

$$\begin{aligned} \sum_{i=1}^n \min\{T_{\max}, S_i\} &= \sum_{i \in G_\epsilon} \min\{T_{\max}, \max(T_i, T'_i)\} + \sum_{i \in G_\epsilon^c} \min\{T_{\max}, T_i\} \\ &\leq \sum_{i \in G_\epsilon} \min\{T_{\max}, \max(T_i, T'_i)\} + |G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| T_{\max} + \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} T_i \end{aligned}$$

$$\begin{aligned}
&= \sum_{i \in G_\epsilon} \max \{T_i, \min(T'_i, T_{\max})\} + |G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| T_{\max} + \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} T_i \\
&\stackrel{(a)}{\leq} \sum_{i \in G_\epsilon} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + |G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right).
\end{aligned}$$

Equality (a) follows from $T_{\max} \leq h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right)$ by Step 1b, $T_i \leq h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right)$ in Steps 2a and 3, and $T'_i \leq h \left(0.25\Delta_i, \frac{\delta}{2n} \right)$ in Step 2c.

E.3.6 Step 5: Bounding the number of samples in round k versus $k - 1$

Now we show that the total number of samples taken in round k is no more than 9 times the number taken in the previous round.

Claim: For $k > 1$

$$\begin{aligned}
&(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\leq 9 (H_{\text{ME}}(n, 2^{-k+1}, 1/16) + \tau_{k-1} + \tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|)
\end{aligned}$$

Proof. In round k , $(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)$ samples are drawn. Since $G_{k-1} \subset G_k$ and $B_{k-1} \subset B_k \forall k$ deterministically, we see that $|(G_{k-1} \cup B_{k-1})^c| \geq |(G_k \cup B_k)^c| \forall k$. By definition, $H_{\text{ME}}(n, 2^{-k-1}, 1/16) = 4H_{\text{ME}}(n, 2^{-k}, 1/16)$.

Next, recall $\tau_k = \left\lceil 2^{2k+3} \log \left(\frac{8}{\delta_k} \right) \right\rceil$. We bound τ_k / τ_{k-1} as

$$\begin{aligned}
\frac{\tau_k}{\tau_{k-1}} &= \frac{\left\lceil 2^{2k+3} \log \left(\frac{8}{\delta_k} \right) \right\rceil}{\left\lceil 2^{2k+1} \log \left(\frac{8}{\delta_{k-1}} \right) \right\rceil} = \frac{\left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil}{\left\lceil 2^{2k+1} \log \left(\frac{16n(k-1)^2}{\delta} \right) \right\rceil} \\
&\leq \frac{2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) + 1}{2^{2k+1} \log \left(\frac{16n(k-1)^2}{\delta} \right)} \leq \frac{4 \log \left(\frac{16nk^2}{\delta} \right)}{\log \left(\frac{16n(k-1)^2}{\delta} \right)} + 1 \\
&\leq 4 \frac{\log \left(\frac{16n}{\delta} \right) + 2 \log(k)}{\log \left(\frac{16n}{\delta} \right) + 2 \log(k-1)} + 1 = (*)
\end{aligned}$$

If $k = 2$, $(*) \leq 1 + 4 * \log(32) / \log(8) \leq 9$. Otherwise,

$$\begin{aligned}
(*) &= \frac{4(\log \left(\frac{16n}{\delta} \right) + 2 \log(k))}{\log \left(\frac{16n}{\delta} \right) + 2 \log(k-1)} + 1 \\
&\leq \frac{4 \log(k)}{\log(k-1)} + 1 \\
&\leq 4 \cdot 2 + 1 = 9
\end{aligned}$$

Putting these pieces together,

$$\begin{aligned}
&(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\leq (4H_{\text{ME}}(n, 2^{-k+1}, 1/16) + 9\tau_{k-1} + 9\tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|) \\
&\leq 9 (H_{\text{ME}}(n, 2^{-k+1}, 1/16) + \tau_{k-1} + \tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|)
\end{aligned}$$

□

E.3.7 Step 6: Bounding Equation (10)

Here, we introduce the round K_{Good} , when $G_{K_{\text{Good}}} = G_\epsilon$ at some point within the round. Using the result of the previous step, we may bound the total number of samples taken though this round, controlling Equation (10).

With the result of Step 5, we prove the following inequality.

Claim:

$$\begin{aligned} & \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} \neq G_\epsilon] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ & \leq c \sum_{i \in G_\epsilon} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + c |G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right) + c \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) \end{aligned} \quad (15)$$

for a constant c .

Proof. Recall $t_k = \max\{t : t \in k\}$ denotes the maximum value of t in round k and $T_{\text{max}} = \max_{i \in G_\epsilon} T_i$ denotes the minimum t such that $G_k(t) = G_\epsilon$. Define the random round

$$K_{\text{Good}} := \min\{k : G_k = G_\epsilon\} = \min\{k : t_k \geq T_{\text{max}}\}$$

By definition of K_{Good} ,

$$\begin{aligned} & \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} \neq G_\epsilon] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ & = \sum_{k=1}^{K_{\text{Good}}} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|). \end{aligned}$$

Next, applying Step 5, if $K_{\text{Good}} > 1$,

$$\begin{aligned} & \sum_{k=1}^{K_{\text{Good}}} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ & \leq 18 \sum_{k=1}^{K_{\text{Good}}-1} \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|). \end{aligned}$$

Observe that by lines 17 and 20 of FAREAST, for any round r and for any $t > t_{r-1}$,

$$\sum_{k=1}^{r-1} \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \leq \sum_{s=1}^t |\mathcal{A}(s)|.$$

By definition, for the round $K_{\text{Good}} - 1$, we see that $t_{(K_{\text{Good}}-1)} < T_{\text{max}}$. Applying the above inequality with the inequality proven in Step 4,

$$\begin{aligned} 18 \sum_{k=1}^{K_{\text{Good}}-1} \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) & \leq 18 \sum_{s=1}^{T_{\text{max}}} |\mathcal{A}(s)| \\ & \leq 18 \sum_{i \in G_\epsilon} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + 18 \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + 18 |G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| h \left(0.25\alpha_\epsilon, \frac{\delta}{2n} \right). \end{aligned}$$

Otherwise, if $K_{\text{Good}} = 1$, exactly $4c'n \log(16) + 32n \log(16n/\delta)$ samples are given to the good filter in round 1. One may use Lemma F.2 to invert $h(\cdot, \cdot)$ and show that the summation on the right had side of the above inequality is within a constant of this and the claim holds in this case as well for a different constant, potentially larger than 18. \square

E.3.8 Step 7: Bounding Equation (11)

Next, we bound $\sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)$.

$$\begin{aligned}
& \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&= \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{\epsilon} \cup B_{k-1})^c|) \\
&= \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |G_{\epsilon}^c \setminus B_{k-1}|) \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |G_{\epsilon}^c \setminus B_{k-1}|) \\
&\stackrel{\mathcal{E}_1, \mathcal{E}_2}{=} \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq G_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |G_{\epsilon}^c \setminus B_{k-1}|) \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq G_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq G_{\epsilon}^c] (\tau_k |G_{\epsilon}^c \setminus B_{k-1}|) \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq G_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} 2\tau_k |G_{\epsilon}^c \setminus B_{k-1}| \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq G_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in G_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&\leq \sum_{k=K_{\text{Good}}+1}^{\infty} 2|G_{\epsilon}^c \setminus B_{k-1}| (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in G_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in G_{\epsilon}^c} 2\mathbb{1}[i \notin B_{k-1}] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in G_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in G_{\epsilon}^c} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&= \sum_{i \in G_{\epsilon}^c} \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&\leq \sum_{i \in G_{\epsilon}^c} \sum_{k=1}^{\infty} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \tag{16}
\end{aligned}$$

E.3.9 Step 8: Bounding the expected total number of samples drawn by FAREAST

Now we take expectations over the number of samples drawn. These expectations are conditional on the high probability event $\mathcal{E}_1 \cap \mathcal{E}_2$. The bound in step 5 holds deterministically conditioned on this event.

Note τ_k and $H_{\text{ME}}(n, 2^{-k}, 1/16)$ are deterministic constants for any k . Let all expectations be jointly over the random instance ν and the randomness in FAREAST.

$$\begin{aligned}
& \mathbb{E}[T | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1] \leq \\
& \sum_{k=1}^{\infty} 2\mathbb{E}[\mathbb{1}[G_k \cup B_k \neq [n]] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1] (\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k |(G_{k-1} \cup B_{k-1})^c|)
\end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} \neq G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\quad + \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\stackrel{\text{Step 6}}{\leq} c \sum_{i \in G_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + c \sum_{i \in G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + c |G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}| h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \\
&\quad + \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} = G_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\stackrel{\text{Step 7}}{\leq} c \sum_{i \in G_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + c \sum_{i \in G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + c |G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}| h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \\
&\quad + \sum_{i \in G_{\epsilon}^c} \sum_{k=1}^{\infty} 2\mathbb{E}_{\nu} \left[\mathbb{1}[i \notin B_{k-1}] \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&\stackrel{(a)}{=} c \sum_{i \in G_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + c \sum_{i \in G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + c |G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}| h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \\
&\quad + \sum_{i \in G_{\epsilon}^c} \sum_{k=1}^{\infty} 2\mathbb{E}_{\nu} \left[\mathbb{1}[i \notin B_{k-1}] \mathbb{1}[\mathcal{E}_1] \right] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16))
\end{aligned}$$

where (a) follows from $\mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1 \cap \mathcal{E}_2] = \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1]$ for $i \in G_{\epsilon}^c$, since the event $\{i \in B_{k-1}\}$ is independent of \mathcal{E}_2 for all $i \in G_{\epsilon}^c$. This can be observed since \mathcal{E}_2 deals only with independent samples taken of arms in G_{ϵ} .

E.3.10 Step 9: Bounding $\sum_{k=1}^{\infty} \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16))$ for $i \in G_{\epsilon}^c$

Next, we bound the expectation remaining from step 8. In particular, this is the number of samples drawn by the bad filter to add arm $i \in G_{\epsilon}^c$ to B_k .

First, we bound the probability that for a given $i \in G_{\epsilon}^c$ and a given k $i \notin B_k$. Note that by Borel-Cantelli, this implies that the probability that i is never added to any B_k is 0.

Claim 1: For $i \in G_{\epsilon}^c, k \geq \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil \implies \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_k] | \mathcal{E}_1] \leq \left(\frac{1}{8} \right)^{k - \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil}$

Proof. $i \in B_k$ if either the good filter or the bad filter added it. Note that the behavior of the bad filter is independent of the event \mathcal{E}_1 . Hence,

$$\begin{aligned}
\mathbb{E}_{\nu} [\mathbb{1}[i \notin B_k] | \mathcal{E}_1] &= \mathbb{E}_{\nu} [\mathbb{1}[\hat{\mu}_i + C_{\delta/2n}(t_k) \geq L_{t_k}] \mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] | \mathcal{E}_1] \\
&\leq \mathbb{E}_{\nu} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] | \mathcal{E}_1] \\
&= \mathbb{E}_{\nu} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}]] .
\end{aligned}$$

Intuitively, the time at which an arm in G_ϵ^c enters B_k , which occurs if either the good filter adds it or the bad filter does, in expectation is at most the time at which the bad filter does on its own in expectation.

If $i \in B_{k-1}$ then $i \in B_k$ by definition. Otherwise, if $i \notin B_{k-1}$, by Hoeffding's Inequality conditional on the value of i_k and a sum over conditional probabilities as in step 0, with probability at least $1 - \frac{\delta}{4nk^2}$

$$|(\hat{\mu}_{i_k} - \hat{\mu}_i) - (\mu_{i_k} - \mu_i)| \leq 2^{-k}$$

If **MedianElimination** also succeeds, the joint event of which occurs with probability $\frac{15}{16} (1 - \frac{\delta}{4nk^2})$ by independence⁶,

$$\hat{\mu}_{i_k} - \hat{\mu}_i \geq \mu_{i_k} - \mu_i - 2^{-k} \geq \mu_1 - \mu_i - 2^{-k+1} = \Delta_i - 2^{-k+1}.$$

Then for $k \geq \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil$,

$$\hat{\mu}_{i_k} - \hat{\mu}_i \geq \Delta_i - 2^{-k+1} \geq \frac{1}{2}(\Delta_i + \epsilon) \geq \epsilon + 2^{-k+1},$$

which implies that $i \in B_k$ by line 15 of **FAREAST**. In particular, $\mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i \geq \epsilon + 2^{-k+1}] | i \notin B_{k-1}] \geq \frac{15}{16} (1 - \frac{\delta}{4nk^2})$. Furthermore, $i \notin B_0$ by definition. Additionally, recall that $\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}]$ is independent of \mathcal{E}_1 . Then for $k \geq \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil$,

$$\begin{aligned} \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}]] &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] (\mathbb{1}[i \notin B_{k-1}] + \mathbb{1}[i \in B_{k-1}]) | \mathcal{E}_1] \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &\quad + \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \in B_{k-1}] | \mathcal{E}_1] \end{aligned}$$

Deterministically, $\mathbb{1}[i \notin B_k] \mathbb{1}[i \in B_{k-1}] = 0$. Therefore,

$$\begin{aligned} &\mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &\quad + \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \in B_{k-1}] | \mathcal{E}_1] \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | i \notin B_{k-1}, \mathcal{E}_1] \mathbb{P}(i \notin B_{k-1} | \mathcal{E}_1) \\ &\quad + \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | i \in B_{k-1}, \mathcal{E}_1] \mathbb{P}(i \in B_{k-1} | \mathcal{E}_1) \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] \mathbb{1}[i \notin B_{k-1}] | i \notin B_{k-1}, \mathcal{E}_1] \mathbb{P}(i \notin B_{k-1} | \mathcal{E}_1) \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] | i \notin B_{k-1}, \mathcal{E}_1] \mathbb{E} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &= \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+1}] | i \notin B_{k-1}] \mathbb{E} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &\leq \left(\frac{1}{16} + \frac{\delta}{4nk^2} \right) \mathbb{E} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] \\ &\leq \left(\frac{1}{16} + \frac{\delta}{4nk^2} \right) \mathbb{E} [\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i < \epsilon + 2^{-k+2}]] \end{aligned}$$

where the final inequality follows by the same argument upper bounding $\mathbb{E} [\mathbb{1}[i \notin B_k] | \mathcal{E}_1]$. For $k < \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil$, trivially, $\mathbb{E} [\mathbb{1}[i \notin B_k]] \leq 1$. Recall $\delta \leq 1/8$. For $k \geq \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil$,

$$\mathbb{E} [\mathbb{1}[i \notin B_k] | \mathcal{E}_1] \leq \prod_{s=\left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil}^k \left(\frac{1}{16} + \frac{\delta}{2ns^2} \right) \leq \left(\frac{1}{8} \right)^{k - \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil}.$$

□

⁶Note that the success of **MedianElimination** and the concentration of $(\hat{\mu}_{i_k} - \hat{\mu}_i)$ around $(\mu_{i_k} - \mu_i)$ are independent of the events \mathcal{E}_1 and \mathcal{E}_2 conditioned on in Step 8.

Claim 2: For $j \in G_\epsilon^c$, $\sum_{k=1}^{\infty} 2\mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]|\mathcal{E}_1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \leq c'' \frac{n}{(\Delta_i - \epsilon)^2} + c''h(0.25(\Delta_i - \epsilon), \frac{\delta}{2n})$

Proof. This sum decomposes into two terms.

$$\begin{aligned} & \sum_{k=1}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]|\mathcal{E}_1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\ &= \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]|\mathcal{E}_1] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &+ \sum_{k=\lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]|\mathcal{E}_1] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \end{aligned}$$

We begin by bounding the first term.

$$\begin{aligned} & \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\leq \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\leq \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} \left(c'n 2^{2k} \log(16) + 2 + 2^{2k+4} \log \left(\frac{16nk^2}{\delta} \right) \right) \\ &\leq 2 \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) + \left(c'n \log(16) + 16 \log \left(\frac{16n}{\delta} \right) \right) \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} 2^{2k} \\ &\quad + 32 \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} 2^{2k} \log(k) \\ &\leq 2 \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) + \left(c'n \log(16) + 16 \log \left(\frac{16n}{\delta} \right) + 32 \log \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right) \sum_{k=1}^{\lfloor \log_2(\frac{4}{\Delta_i - \epsilon}) \rfloor} 2^{2k} \\ &\leq 2 \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) + \frac{16}{(\Delta_i - \epsilon)^2} \left(c'n \log(16) + 32 \log \left(\frac{16n}{\delta} \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right) \right) \end{aligned}$$

Next, we plug in the bound from claim 1 controlling the probability that $i \notin B_k$.

Using Claim 1, we bound the second sum as follows:

$$\begin{aligned} & \sum_{r=\lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}]|\mathcal{E}_1] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\leq \sum_{k=\lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil}^{\infty} \left(\frac{1}{8} \right)^{k - \lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil - 1} \left(c'n 2^{2k} \log(16) + 2 + 2^{2k+4} \log \left(\frac{16nk^2}{\delta} \right) \right) \\ &= c'n \log(16) \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} 2^{2(k + \lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil)} + 2 \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} \\ &\quad + 16 \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} 2^{2(k + \lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil)} \log \left(\frac{16n(k + \lceil \log_2(\frac{4}{\Delta_i - \epsilon}) \rceil)^2}{\delta} \right) \end{aligned}$$

$$\begin{aligned}
&\leq 3 + c'n \log(16) \sum_{k=1}^{\infty} 2^{-3k+3} 2^{2\left(k+\log_2\left(\frac{4}{\Delta_i-\epsilon}\right)+1\right)} \\
&\quad + 16 \sum_{k=1}^{\infty} 2^{-3k+3} 2^{2\left(k+\log_2\left(\frac{4}{\Delta_i-\epsilon}\right)+1\right)} \log \left(\frac{16n \left(k + \left\lceil \log_2 \left(\frac{4}{\Delta_i-\epsilon} \right) \right\rceil\right)^2}{\delta} \right) \\
&= 3 + \left(\frac{2^9 c'n \log(16)}{(\Delta_i - \epsilon)^2} + \frac{2^{13}}{(\Delta_i - \epsilon)^2} \log \left(\frac{16n}{\delta} \right) \right) \sum_{k=1}^{\infty} 2^{-k} \\
&\quad + \frac{2^{13}}{(\Delta_i - \epsilon)^2} \sum_{k=1}^{\infty} 2^{-k} \log \left(\left(k + \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil \right)^2 \right) \\
&\leq 3 + \frac{2^9 c'n \log(16)}{(\Delta_i - \epsilon)^2} + \frac{2^{13}}{(\Delta_i - \epsilon)^2} \log \left(\frac{16n}{\delta} \right) \\
&\quad + \frac{2^{14}}{(\Delta_i - \epsilon)^2} \sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil \right) \\
&= (**)
\end{aligned}$$

We may bound the final summand, $\sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil \right)$ as follows:

$$\sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{4}{\Delta_i - \epsilon} \right) \right\rceil \right) \leq \log \left(\frac{e}{2} \log_2 \left(\frac{256}{(\Delta_i - \epsilon)^2} \right) \right)$$

Plugging this back into (**), we have that

$$(**) \leq 3 + \frac{2^9 c'n \log(16)}{(\Delta_i - \epsilon)^2} + \frac{2^{13}}{(\Delta_i - \epsilon)^2} \log \left(\frac{16n}{\delta} \right) + \frac{2^{14}}{(\Delta_i - \epsilon)^2} \log \left(\frac{e}{2} \log_2 \left(\frac{256}{(\Delta_i - \epsilon)^2} \right) \right)$$

Combining the above with the bound on the first sum, we have that

$$\begin{aligned}
&\sum_{k=1}^{\infty} \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&\leq c'' \left(\frac{n}{(\Delta_i - \epsilon)^2} + \frac{c}{(\Delta_i - \epsilon)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{4}{(\Delta_i - \epsilon)^2} \right) \right) \right) \\
&= \frac{c''n}{(\Delta_i - \epsilon)^2} + c''h \left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n} \right)
\end{aligned}$$

for a sufficiently large, universal constant c'' and c from the definition of $h(\cdot, \cdot)$. \square

E.3.11 Step 10: Applying the result of Step 9 to the result of Step 8

We may repeat the result of step 9 for every $i \in G_{\epsilon}^c$ and plug this into the result of Step 8. From this point, we simplify to return the final result.

By Step 8, the total number of samples T drawn by FAREAST is bounded in expectation by

$$\begin{aligned}
\mathbb{E}[T | \mathcal{E}_1 \cap \mathcal{E}_2] &\leq c \sum_{i \in G_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + c \sum_{i \in G_{\epsilon+\alpha_{\epsilon}}^c} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + c |G_{\epsilon}^c \cap G_{\epsilon+\alpha_{\epsilon}}| h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \\
&\quad + 2 \sum_{i \in G_{\epsilon}^c} \sum_{k=1}^{\infty} \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] | \mathcal{E}_1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)).
\end{aligned}$$

Applying the bound from Step 9 to each $i \in G_{\epsilon}^c$, we have that

$$\mathbb{E}[T | \mathcal{E}_1 \cap \mathcal{E}_2] \leq c \sum_{i \in G_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(0.25\alpha_{\epsilon}, \frac{\delta}{2n} \right) \right] \right\}$$

$$\begin{aligned}
& + c \sum_{i \in G_{\epsilon+\alpha_\epsilon}^c} h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}\right) + c|G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}| h\left(0.25\alpha_\epsilon, \frac{\delta}{2n}\right) \\
& + 2c'' \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + h\left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n}\right).
\end{aligned}$$

For $i \in G_\epsilon^c \cap G_{\epsilon+\alpha_\epsilon}$, $\alpha_\epsilon = \min_{j \in G_\epsilon} \epsilon - \Delta_j \geq \Delta_i - \epsilon$. By monotonicity of $h(\cdot, \cdot)$, $h\left(0.25\alpha_\epsilon, \frac{\delta}{2n}\right) \leq \frac{c''n}{(\Delta_i - \epsilon)^2} + c''h\left(\Delta_i - \epsilon, \frac{\delta}{2n}\right)$. Therefore,

$$\begin{aligned}
\mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in G_\epsilon} \max \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}\right), \min \left[h\left(0.25\Delta_i, \frac{\delta}{2n}\right), h\left(0.25\alpha_\epsilon, \frac{\delta}{2n}\right) \right] \right\} \\
& + (2c'' + c) \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + h\left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n}\right).
\end{aligned}$$

Next, we use Lemma F.3 to bound the minimum of $h(\cdot, \dots)$ functions.

$$\begin{aligned}
& c \sum_{i \in G_\epsilon} \max \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}\right), \min \left[h\left(0.25\Delta_i, \frac{\delta}{2n}\right), h\left(0.25\alpha_\epsilon, \frac{\delta}{2n}\right) \right] \right\} \\
& + (2c'' + c) \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + h\left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n}\right) \\
& \leq c \sum_{i \in G_\epsilon} \max \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}\right), h\left(\frac{\Delta_i + \alpha_\epsilon}{8}, \frac{\delta}{2n}\right) \right\} \\
& + (2c'' + c) \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + h\left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n}\right)
\end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i , we have that $0.25\Delta_i \leq 2$, $0.25|\epsilon - \Delta_i| \leq 2$, and $0.25 \min(\alpha_\epsilon, \beta_\epsilon) \leq 2$. Therefore,

$$\begin{aligned}
\mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in G_\epsilon} \max \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n}\right), h\left(\frac{\Delta_i + \alpha_\epsilon}{8}, \frac{\delta}{2n}\right) \right\} \\
& + (2c'' + c) \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + h\left(0.25(\Delta_i - \epsilon), \frac{\delta}{2n}\right) \\
& \leq c \sum_{i \in G_\epsilon} \max \left\{ \frac{64}{(\epsilon - \Delta_i)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{384n}{\delta(\epsilon - \Delta_i)^2} \right) \right), \right. \\
& \quad \left. \frac{256}{(\Delta_i + \alpha_\epsilon)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{768n}{\delta(\Delta_i + \alpha_\epsilon)^2} \right) \right) \right\} \\
& + (2c'' + c) \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + \frac{64}{(\epsilon - \Delta_i)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{384n}{\delta(\epsilon - \Delta_i)^2} \right) \right) \\
& \leq c_3 \sum_{i \in G_\epsilon} \max \left\{ \frac{1}{(\epsilon - \Delta_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\epsilon - \Delta_i)^2} \right) \right), \right. \\
& \quad \left. \frac{1}{(\Delta_i + \alpha_\epsilon)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\Delta_i + \alpha_\epsilon)^2} \right) \right) \right\} \\
& + c_3 \sum_{i \in G_\epsilon^c} \frac{n}{(\Delta_i - \epsilon)^2} + \frac{1}{(\epsilon - \Delta_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\epsilon - \Delta_i)^2} \right) \right) \\
& = c_3 \sum_{i \in G_\epsilon} \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right.
\end{aligned}$$

$$\frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right) \Bigg\} \\ + c_3 \sum_{i \in G_\epsilon^c} \frac{n}{(\mu_1 - \epsilon - \mu_i)^2} + \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right)$$

for a sufficiently large constant c_4 .

E.3.12 Step 11: High probability sample complexity bound

Finally, the Good Filter is equivalent to EAST, Algorithm 5, except split across rounds. Note that the Good Filter is union bounded over $2n$ events whereas the bounds in EAST are union bounded over n events. The Good Filter and Bad Filter are given the same number of samples in each round, and the Good Filter can terminate within a round, conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$. Therefore, we can bound the complexity of FAREAST in terms of that of EAST run at failure probability $\delta/2$. If FAREAST terminates in the second round or later, the arguments in Steps 4 and 5 can be used to show that FAREAST draws no more than a factor of 18 more samples than EAST, though this estimate is highly pessimistic. If FAREAST terminates in round 1 (when gaps are large), we may still show that this is within a constant factor of the complexity of EAST, but the story is more complicated. In the first round, the bad filter draws at most $c'n \log(16) + 32n \log(8n/\delta)$ samples where c' is the constant from Median Elimination. Since we have assumed that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$, this sum is likewise within a constant factor of the complexity of EAST. Hence, by Theorem E.3,

$$T \leq c_4 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\ \frac{1}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right), \\ \left. \frac{1}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \right\} \\ \left. \frac{1}{\gamma^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta\gamma^2} \right) \right) \right\}$$

samples. □

E.4 Proof of Theorem E.2, FAREAST in the multiplicative regime

Proof. Notation for the proof: Throughout, recall $\Delta_i = \mu_1 - \mu_i$. Recall that t counts the number of times the conditional in line 19 is true. By Line 19 of FAREAST, all arms in \mathcal{A} have received t samples when the loop in line 23 is executed for the t^{th} time. Within any round k , let $\mathcal{A}(t)$ and $G_k(t)$ denote the sets \mathcal{A} and G_k at this time since both sets can change in lines 27 and 29 and 25 respectively. Let t_k denote the maximum value of t in round k . By Lines 18 and 19 of FAREAST, the total number of samples given to the good filter when the conditional in line 19 is true for the t^{th} time is $\sum_{s=1}^t |\mathcal{A}(s)|$.

For $i \in M_\epsilon$, let T_i denote the random variable of the number of times arm i is sampled before it is added to G_k in Line 25. For $i \in M_\epsilon^c$, let T_i denote the random variable of the number of times arm i is sampled before it is removed from \mathcal{A} in Line 27. For any arm i , let T'_i denote the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$.

Define the event

$$\mathcal{E}_1 = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/2n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\mathbb{P}(\mathcal{E}_1^c) = \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/2n}(t) \right)$$

$$\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/2n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{2n} = \frac{\delta}{2}$$

Next, recall that $\hat{\mu}_i(t)$ denotes the empirical average of t samples of ρ_i . Consider the event,

$$\mathcal{E}_2 = \bigcap_{i \in M_\epsilon} \bigcap_{k \in \mathbb{N}} |((1-\epsilon)\hat{\mu}_{i_k}(\tau_k) - \hat{\mu}_i(\tau_k)) - ((1-\epsilon)\mu_{i_k} - \mu_i)| \leq 2^{-(k+1)}(2-\epsilon)$$

By Hoeffding's inequality,

$$\mathbb{P} \left(|((1-\epsilon)\hat{\mu}_{i_k}(\tau_k) - \hat{\mu}_i(\tau_k)) - ((1-\epsilon)\mu_{i_k} - \mu_i)| \leq 2^{-(k+1)}(2-\epsilon) \mid i_k = j \right) \leq \frac{\delta}{4nk^2}.$$

Then

$$\begin{aligned} & \mathbb{P}(|((1-\epsilon)\hat{\mu}_{i_k}(\tau_k) - \hat{\mu}_i(\tau_k)) - ((1-\epsilon)\mu_{i_k} - \mu_i)| \leq 2^{-(k+1)}(2-\epsilon)) \\ &= \sum_{j=1}^n \mathbb{P}(|((1-\epsilon)\hat{\mu}_{i_k}(\tau_k) - \hat{\mu}_i(\tau_k)) - ((1-\epsilon)\mu_{i_k} - \mu_i)| \leq 2^{-(k+1)}(2-\epsilon) \mid i_k = j) \mathbb{P}(i_k = j) \\ &\leq \frac{\delta}{4nk^2} \sum_{j=1}^n \mathbb{P}(i_k = j) \\ &= \frac{\delta}{4nk^2} \end{aligned}$$

Therefore, union bounding over the rounds $k \in \mathbb{N}$, $\mathbb{P}(\mathcal{E}_2^c) \leq \sum_{i \in M_\epsilon} \sum_{k=1}^{\infty} \frac{\delta}{4nk^2} \leq \frac{\delta}{2}$. Hence, $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1 - \delta$.

E.4.1 Step 0: Correctness.

On $\mathcal{E}_1 \cap \mathcal{E}_2$, first we prove that if there exists a random round k at which $G_k \cup B_k = [n]$ then $G_k = M_\epsilon$. Additionally, we prove that on $\mathcal{E}_1 \cap \mathcal{E}_2$, if $\mathcal{A} \subset G_k$, then $G_k = M_\epsilon$. Therefore, for either stopping condition for FAREAST in line 31, on the event $\mathcal{E}_1 \cap \mathcal{E}_2$, FAREAST correctly returns the set M_ϵ .

Claim 0: On $\mathcal{E}_1 \cap \mathcal{E}_2$, for all $k \in \mathbb{N}$, $G_k \subset M_\epsilon$.

Proof. Firstly we show $1 \in \mathcal{A}$ for all $t \in \mathbb{N}$, namely the best arm is never removed from \mathcal{A} . Note for any i such that $\hat{\mu}_i(t) - C_{\delta/2n}(t) \geq 0$,

$$\hat{\mu}_1 + C_{\delta/2n}(t) \geq \mu_1 \geq \mu_i \geq \hat{\mu}_i(t) - C_{\delta/2n}(t) > (1-\epsilon)(\hat{\mu}_i(t) - C_{\delta/2n}(t)).$$

For i such that $\hat{\mu}_i(t) - C_{\delta/2n}(t) < 0$, if $\hat{\mu}_1 + C_{\delta/2n}(t) \geq 0$, then

$$\hat{\mu}_1 + C_{\delta/2n}(t) \geq 0 > (1-\epsilon)(\hat{\mu}_i(t) - C_{\delta/2n}(t)).$$

Note that $\hat{\mu}_1 + C_{\delta/2n}(t) < 0$ implies on the event \mathcal{E}_1 that $\mu_1 < 0$, which contradicts the assumption that $\mu_1 \geq 0$ made in the theorem. In particular this shows, $\hat{\mu}_1 + C_{\delta/2n}(t) > (1-\epsilon)(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/2n}(t)) = L_t$ and $\hat{\mu}_1 + C_{\delta/2n}(t) \geq \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/2n}(t)$ showing that 1 will never exit \mathcal{A} in line 28.

Secondly, we show that at all times t , $(1-\epsilon)\mu_1 \in [L_t, U_t]$. By the above, since μ_1 never leaves \mathcal{A} ,

$$U_t = (1-\epsilon)(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/2n}(t)) \geq (1-\epsilon)(\hat{\mu}_1(t) + C_{\delta/2n}(t)) \geq (1-\epsilon)\mu_1$$

and for any i ,

$$(1-\epsilon)\mu_1 \geq (1-\epsilon)\mu_i \geq (1-\epsilon)(\hat{\mu}_i(t) - C_{\delta/2n}(t))$$

Hence $(1-\epsilon)\mu_1 \geq (1-\epsilon)(\max_i \hat{\mu}_i(t) - C_{\delta/2n}(t)) = L_t$.

Next, we show that $G_k \subset M_\epsilon$ for all $k \geq 1, t \geq 1$. Suppose not. Then $\exists k, t \in \mathbb{N}$ and $\exists i \in M_\epsilon^c \cap G_k(t)$ such that,

$$\mu_i \geq \hat{\mu}_i(t) - C_{\delta/2n}(t) \geq U_t \geq (1-\epsilon)\mu_1 > \mu_i,$$

with the last inequality following from the previous assertion, giving a contradiction. \square

Claim 1: On $\mathcal{E}_1 \cap \mathcal{E}_2$, for all $k \in \mathbb{N}$, $B_k \subset M_\epsilon^c$.

Proof. Next, we show $B_k \subset M_\epsilon^c$. Suppose not. Then either the good filter or the bad filter added an arm in M_ϵ to B_k . Take $i \in M_\epsilon$. In the former, this implies that

$$\mu_i \stackrel{\mathcal{E}_1}{\leq} \hat{\mu}_i(t) + C_{\delta/2n}(t) < L_t \stackrel{\mathcal{E}_1}{\leq} (1 - \epsilon)\mu_1$$

which contradicts $i \in M_\epsilon$. Consider the alternate case that the bad filter adds i to B_k for some k . By definition, $B_0 = \emptyset$ and $B_{k-1} \subset B_k$ for all k . Then there must exist $k \in \mathbb{N}$ and an $i \in M_\epsilon$ such that $i \in B_k$ and $i \notin B_{k-1}$. Following line 14 of the algorithm, this occurs if and only if

$$(1 - \epsilon)\hat{\mu}_{i_k} - \hat{\mu}_i > 2^{-(k+1)}(2 - \epsilon).$$

On the event \mathcal{E}_2 , the above implies

$$(1 - \epsilon)\mu_{i_k} - \mu_i + 2^{-(k+1)}(2 - \epsilon) > 2^{-(k+1)}(2 - \epsilon),$$

and simplifying, we see that $0 < (1 - \epsilon)\mu_{i_k} - \mu_i \leq (1 - \epsilon)\mu_1 - \mu_i$ which contradicts the assertion that $i \in M_\epsilon$. Combining the above claims, we see that $\mathcal{E}_1 \cap \mathcal{E}_2$ implies $(G_k \cup B_k = [n])$ and $G_k \cap B_k = \emptyset \implies G_k = M_\epsilon$. Since $\mathbb{P}(\mathcal{E}_1 \cap \mathcal{E}_2) \geq 1 - \delta$, if FAREAST terminates, with probability at least $1 - \delta$, it correctly returns the set M_ϵ . \square

Claim 2: Next, we show that on \mathcal{E}_1 , $M_\epsilon \subset \mathcal{A}(t) \cup G(t)$ for all $t \in \mathbb{N}$.

In particular this implies that if $\mathcal{A} \subset G$, then $M_\epsilon \subset G$. Combining this with the previous claim gives $G \subset M_\epsilon \subset G$, hence $G = M_\epsilon$. On this condition, FAREAST terminates by line 33 and returns the set $\mathcal{A} \cup G = G$. Note that by definition, $M_\epsilon \subset M_{(\epsilon+\gamma)}$ for all $\gamma \geq 0$. Therefore FAREAST terminates correctly on this condition.

Proof. Suppose for contradiction that there exists $i \in M_\epsilon$ such that $i \notin \mathcal{A}(t) \cup G(t)$. This occurs only if i is eliminated in line 28. Hence, there exists a $t' \leq t$ such that $\hat{\mu}_i(t') + C_{\delta/n}(t') < L_{t'}$. Therefore, on the event \mathcal{E}_1 ,

$$(1 - \epsilon)\mu_1 \stackrel{\mathcal{E}_1}{\geq} L_{t'} = (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \hat{\mu}_j(t') - C_{\delta/n}(t') \right) > \hat{\mu}_i(t') + C_{\delta/n}(t') \stackrel{\mathcal{E}_1}{\geq} \mu_i$$

which contradicts $i \in M_\epsilon$. \square

Claim 3: Finally, we show that on \mathcal{E}_1 , if $U_t - L_t \leq \frac{\gamma}{2-\epsilon}L_t$, then $\mathcal{A} \cup G \subset M_{(\epsilon+\gamma)}$.

Combining with Claim 3 that $M_\epsilon \subset \mathcal{A} \cup G$, if FAREAST terminates on this condition by line 33, it does so correctly and returns all arms in M_ϵ and none in $M_{(\epsilon+\gamma)}^c$.

Proof. By Claim 0, $G \subset M_\epsilon \subset M_{(\epsilon+\gamma)}$. Hence, $G \cap M_{(\epsilon+\gamma)}^c = \emptyset$. Therefore, we wish to show that $\mathcal{A} \cap M_{(\epsilon+\gamma)}^c = \emptyset$ which implies that $G \cap \mathcal{A} \subset M_{(\epsilon+\gamma)}$. Assume $U_t - L_t < \frac{\gamma}{2-\epsilon}L_t$. Recall that

$$U_t = (1 - \epsilon) \left(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/2n}(t) \right)$$

and

$$L_t = (1 - \epsilon) \left(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/2n}(t) \right)$$

All arms in $\mathcal{A}(t)$ have received exactly t samples. Hence, $U_t - L_t = 2(1 - \epsilon)C_{\delta/2n}(t)$. On \mathcal{E}_1 , $L_t \leq (1 - \epsilon)\mu_1$. This implies that

$$2(1 - \epsilon)C_{\delta/2n}(t) < \frac{\gamma}{2 - \epsilon}L_t \leq \frac{1 - \epsilon}{2 - \epsilon}\gamma\mu_1,$$

and in particular,

$$2C_{\delta/2n}(t) < \frac{\gamma\mu_1}{2 - \epsilon}.$$

Therefore, we wish to show that when the above is true, then for any $i \in M_{(\epsilon+\gamma)}^c$, $L_t - (\hat{\mu}_i(t) + C_{\delta/n}(t)) > 0$, implying that $i \notin \mathcal{A}$.

$$L_t - (\hat{\mu}_i(t) + C_{\delta/n}(t)) = (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \hat{\mu}_j - C_{\delta/n}(t) \right) - (\hat{\mu}_i(t) + C_{\delta/n}(t))$$

$$\begin{aligned}
&\geq (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \mu_j - 2C_{\delta/n}(t) \right) - (\mu_i + 2C_{\delta/n}(t)) \\
&\stackrel{(a)}{\geq} (1 - \epsilon) (\mu_1 - 2C_{\delta/n}(t)) - ((1 - \epsilon - \gamma)\mu_1 + 2C_{\delta/n}(t)) \\
&= \gamma\mu_1 - 2(2 - \epsilon)C_{\delta/n}(t) \\
&> \gamma\mu_1 - (2 - \epsilon)\frac{\gamma\mu_1}{2 - \epsilon} \\
&= 0
\end{aligned}$$

which implies that $i \notin \mathcal{A}$. Inequality (a) follows jointly from the fact that $1 \in \mathcal{A}$ and the fact that all arms in \mathcal{A} have received t samples implies $\max_{j \in \mathcal{A}} \mu_j - 2C_{\delta/n}(t) = \mu_1 - 2C_{\delta/n}(t)$. Additionally, inequality (a) follows from $\mu_i \leq (1 - \epsilon - \gamma)\mu_1$ since $i \in M_{\epsilon+\gamma}^c$. \square

E.4.2 Step 1: An expression for the total number of samples drawn and introducing several helper random variables

Next, we write an expression for the total number of samples drawn by FAREAST. In particular, we introduce two sums that we will spend the remainder of the proof controlling. Additionally, we show that the conditional in line 19 in the good filter is true at least once in each round. Based on this, we more precisely define the random variables T_i and T'_i introduced in the notation section in section E.4. Additionally, we introduce the time T_γ at which $U_t - L_t < \frac{\gamma}{2-\epsilon}L_t$.

Recall that the largest value of t in round k is denoted t_k . Let E_k^γ be the event that $U_t - L_t \geq \frac{\gamma}{2-\epsilon}L_t$ for all t in round k :

$$E_k^\gamma := \left\{ U_t - L_t \geq \frac{\gamma}{2-\epsilon}L_t : t \in (t_{k-1}, t_k] \right\}.$$

Note that if E_{k-1}^γ is false, then FAREAST terminates in round $k - 1$ by line 33. We may write the total number of samples drawn by the algorithm as

$$\begin{aligned}
T &= \sum_{k=1}^{\infty} 2\mathbb{1} [\mathcal{A} \not\subset G_{k-1} \text{ and } G_{k-1} \cup B_{k-1} \neq [n] \text{ and } E_{k-1}^\gamma] \\
&\quad (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)
\end{aligned}$$

Deterministically, $\mathbb{1} [\mathcal{A} \not\subset G_{k-1} \text{ and } G_{k-1} \cup B_{k-1} \neq [n] \text{ and } E_{k-1}^\gamma] \leq \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]]$.

Applying this,

$$\begin{aligned}
T &\leq \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&= \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} \neq M_\epsilon] \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)
\end{aligned} \tag{17}$$

$$+ \sum_{k=1}^{\infty} 2\mathbb{1} [G_{k-1} = M_\epsilon] \mathbb{1} [G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \tag{18}$$

In round k , line 18 of the Good Filter, whereby an arm is sampled, is evaluated

$$(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \geq (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \geq n$$

times since $H_{\text{ME}}(n, 2^{-k}, 1/16) \geq n$ for all k and $|(G_{k-1} \cup B_{k-1})^c| \geq 1$ unless $G_{k-1} \cup B_{k-1} = [n]$ which implies termination in round $k - 1$. Each time line 18 is called, $N_{I_s} \leftarrow N_{I_s} + 1$. Since $|\arg \min_{j \in \mathcal{A}} \{N_j\}| \leq |\mathcal{A}| \leq n$, line 18 is called at most n times before $\min_{j \in \mathcal{A}} \{N_j\} = \max_{j \in \mathcal{A}} \{N_j\}$. When this occurs, the conditional in line 19 is true and $t \leftarrow t + 1$.

If $\min_{i \in \mathcal{A}(t)} \{N_i\} = \max_{i \in \mathcal{A}(t)} \{N_i\}$, then $N_i = t$ for any $i \in \mathcal{A}(t)$. By Step 0, only arms in M_ϵ are added to G_k . Therefore, T_i is defined as

$$T_i = \min \left\{ t : \begin{array}{ll} i \in G_k(t+1) & \text{if } i \in M_\epsilon \\ i \notin \mathcal{A}(t+1) & \text{if } i \in M_\epsilon^c \end{array} \right\} \stackrel{\mathcal{E}_1}{=} \min \left\{ t : \begin{array}{ll} \hat{\mu}_i - C_{\delta/2n}(t) \geq U_t & \text{if } i \in M_\epsilon \\ \hat{\mu}_i + C_{\delta/2n}(t) \leq L_t & \text{if } i \in M_\epsilon^c \end{array} \right\} \quad (19)$$

Define $T_i = \infty$ if this never occurs. Note that this may happen if FAREAST terminates due to the condition in line 32 that $U_t - L_t < \frac{\gamma}{2-\epsilon} L_t$. Similarly, recall T'_i denotes the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)$. Hence,

$$T'_i = \min \left\{ t : \hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/2n}(t) \right\} \quad (20)$$

Define $T'_i = \infty$ if this never occurs. Note that this may happen if FAREAST terminates due to the condition in line 32 that $U_t - L_t < \frac{\gamma}{2-\epsilon} L_t$. Finally, we define the time T_γ such that $U_t - L_t < \frac{\gamma}{2-\epsilon} L_t$.

$$T_\gamma = \min \left\{ t : U_t - L_t < \frac{\gamma}{2-\epsilon} L_t \right\} \quad (21)$$

By design, no arm is sampled more than T_γ times by the good filter, controlling the cases that T_i or T'_i are infinite.

E.4.3 Step 2: Bounding T_i and T'_i for $i \in M_\epsilon$

Step 2a: For $i \in M_\epsilon$, we have that $T_i \leq h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right)$.

Proof. Note that $\mu_i - 2C_{\delta/2n}(t) \geq (1-\epsilon)(\mu_1 + 2C_{\delta/2n}(t))$ may be rearranged as $(4-2\epsilon)C_{\delta/2n}(t) \leq \epsilon\mu_1 - \Delta_i$, and this is true when $t > h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right)$. This condition implies that for all j ,

$$\begin{aligned} \hat{\mu}_i(t) - C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\geq} \mu_i - 2C_{\delta/2n}(t) \\ &\geq (1-\epsilon)(\mu_1 + 2C_{\delta/2n}(t)) \\ &\geq (1-\epsilon)(\mu_j + 2C_{\delta/2n}(t)) \\ &\stackrel{\mathcal{E}_1}{\geq} (1-\epsilon)(\hat{\mu}_j(t) + C_{\delta/2n}(t)) \end{aligned}$$

so in particular, $\hat{\mu}_i(t) - C_{\delta/2n}(t) \geq (1-\epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/2n}(t)) = U_t$. \square

Additionally, we define a time T_{\max} when all good arms have entered G_k .

Step 2b: Defining $T_{\max} := \min\{t : G_k(t) = M_\epsilon\} = \max_{i \in M_\epsilon} T_i$, we also have that $T_{\max} \leq h(\tilde{\alpha}_\epsilon/(4-2\epsilon), \delta/2n)$ (in other words, if $t > h(\tilde{\alpha}_\epsilon/(4-2\epsilon), \delta/2n)$ (i.e. line 23 has been run t times, then we have that $G_k(t) = M_\epsilon$).

Proof. Recall that $\tilde{\alpha}_\epsilon = \min_{i \in M_\epsilon} \mu_i - \mu_1 + \epsilon = \min_{i \in M_\epsilon} \epsilon\mu_1 - \Delta_i$. By Step 1a, $T_i \leq h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right)$. Furthermore, $h(\cdot, \cdot)$ is monotonic in its first argument, such that if $0 < x' < x$, then $h(x', \delta) > h(x, \delta)$ for any $\delta > 0$. Therefore $T_{\max} = \max_{i \in M_\epsilon} T_i \leq \max_{i \in M_\epsilon} h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right) = h\left(\tilde{\alpha}_\epsilon/(4-2\epsilon), \frac{\delta}{2n}\right)$. \square

Step 2c: For $i \in M_\epsilon$, we have that $T'_i \leq h(0.25\Delta_i, \delta/2n)$.

Proof. Note that $4C_{\delta/2n}(t) \leq \mu_1 - \mu_i$, true when $t > h\left(0.25\Delta_i, \frac{\delta}{2n}\right)$, implies that

$$\begin{aligned} \hat{\mu}_i(t) + C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\leq} \mu_i + 2C_{\delta/2n}(t) \\ &\leq \mu_1 - 2C_{\delta/2n}(t) \\ &\stackrel{\mathcal{E}_1}{\leq} \hat{\mu}_1(t) - C_{\delta/2n}(t). \end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Hence, $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/2n}(t)$. \square

E.4.4 Step 3: Bounding T_i for $i \in M_\epsilon^c$

Next, we bound T_i for $i \in M_\epsilon^c$. $i \in M_\epsilon^c$ is eliminated from \mathcal{A} if it has received at least T_i samples.

Claim: $T_i \leq h\left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n}\right)$ for $i \in M_\epsilon^c$

Proof. Note that $\mu_i + 2C_{\delta/2n}(t) \leq (1-\epsilon)(\mu_1 - 2C_{\delta/2n}(t))$ may be rearranged as $(4-2\epsilon)C_{\delta/2n}(t) \leq \Delta_i - \epsilon\mu_1$, and this is true when $t > h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right)$. This condition implies that

$$\begin{aligned}\hat{\mu}_i(t) + C_{\delta/2n}(t) &\stackrel{\mathcal{E}_1}{\leq} \mu_i + 2C_{\delta/2n}(t) \\ &\leq (1-\epsilon)(\mu_1 - 2C_{\delta/2n}(t)) \\ &\stackrel{\mathcal{E}_1}{\leq} (1-\epsilon)(\hat{\mu}_1(t) - C_{\delta/2n}(t))\end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Therefore $\hat{\mu}_i(t) + C_{\delta/2n}(t) \leq (1-\epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/2n}(t)) = L_t$. \square

E.4.5 Step 4: bounding the total number of samples given to the good filter at time $t = T_{\max}$

Note that for a time $t = T$, the total number of samples given to the good filter is $\sum_{s=1}^T |\mathcal{A}(s)|$. Therefore, the total number of samples up to time T_{\max} is $\sum_{t=1}^{T_{\max}} |\mathcal{A}(t)|$.

Let $S_i = \min\{t : i \notin \mathcal{A}(t+1)\}$. Hence,

$$\sum_{t=1}^{T_{\max}} |\mathcal{A}(t)| = \sum_{t=1}^{T_{\max}} \sum_{i=1}^n \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \sum_{t=1}^{T_{\max}} \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \min\{T_{\max}, S_i\}$$

For arms $i \in M_\epsilon^c$, $S_i = T_i$ by definition. For $i \in M_\epsilon$, $S_i = \max(T_i, T'_i)$ by line 28 of the algorithm. Then

$$\begin{aligned}\sum_{i=1}^n \min\{T_{\max}, S_i\} &= \sum_{i \in M_\epsilon} \min\{T_{\max}, \max(T_i, T'_i)\} + \sum_{i \in M_\epsilon^c} \min\{T_{\max}, T_i\} \\ &\leq \sum_{i \in M_\epsilon} \min\{T_{\max}, \max(T_i, T'_i)\} + |M_\epsilon^c \cap M_{\epsilon+\tilde{\alpha}_\epsilon}| T_{\max} + \sum_{i \in M_{\epsilon+\tilde{\alpha}_\epsilon}^c} T_i \\ &= \sum_{i \in M_\epsilon} \max\{T_i, \min(T'_i, T_{\max})\} + |M_\epsilon^c \cap M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}| T_{\max} + \sum_{i \in M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}^c} T_i \\ &\stackrel{(a)}{\leq} \sum_{i \in M_\epsilon} \max\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right), \min\left[h\left(0.25\Delta_i, \frac{\delta}{2n}\right), h\left(\frac{\tilde{\alpha}_\epsilon}{4-2\epsilon}, \frac{\delta}{2n}\right)\right]\right\} \\ &\quad + \sum_{i \in M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}^c} h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right) + |M_\epsilon^c \cap M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}| h\left(\frac{\tilde{\alpha}_\epsilon}{4-2\epsilon}, \frac{\delta}{2n}\right).\end{aligned}$$

Equality (a) follows from $T_{\max} \leq h\left(\frac{\tilde{\alpha}_\epsilon}{4-2\epsilon}, \frac{\delta}{2n}\right)$ by Step 1b, $T_i \leq h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right)$ in Steps 2a and 3, and $T'_i \leq h\left(0.25\Delta_i, \frac{\delta}{2n}\right)$ in Step 2c.

E.4.6 Step 5: Bounding the number of samples in round k versus $k-1$

Now we show that the total number of samples taken in round k is no more than 9 times the number taken in the previous round.

Claim: For $k > 1$

$$(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \leq 9 (H_{\text{ME}}(n, 2^{-k+1}, 1/16) + \tau_{k-1} + \tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|)$$

Proof. In round k , $(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)$ samples are drawn. Since $G_{k-1} \subset G_k$ and $B_{k-1} \subset B_k \forall k$ deterministically, we see that $|(G_{k-1} \cup B_{k-1})^c| \geq |(G_k \cup B_k)^c| \forall k$.

By definition,

$$H_{\text{ME}}(n, 2^{-k-1}, 1/16) = 4H_{\text{ME}}(n, 2^{-k}, 1/16).$$

Next, recall $\tau_k = \left\lceil 2^{2k+3} \log\left(\frac{8}{\delta_k}\right) \right\rceil$. We bound τ_k/τ_{k-1} as

$$\begin{aligned} \frac{\tau_k}{\tau_{k-1}} &= \frac{\left\lceil 2^{2k+3} \log\left(\frac{8}{\delta_k}\right) \right\rceil}{\left\lceil 2^{2k+1} \log\left(\frac{8}{\delta_{k-1}}\right) \right\rceil} = \frac{\left\lceil 2^{2k+3} \log\left(\frac{16nk^2}{\delta}\right) \right\rceil}{\left\lceil 2^{2k+1} \log\left(\frac{16n(k-1)^2}{\delta}\right) \right\rceil} \\ &\leq \frac{2^{2k+3} \log\left(\frac{16nk^2}{\delta}\right) + 1}{2^{2k+1} \log\left(\frac{16n(k-1)^2}{\delta}\right)} \leq \frac{4 \log\left(\frac{16nk^2}{\delta}\right)}{\log\left(\frac{16n(k-1)^2}{\delta}\right)} + 1 \\ &\leq 4 \frac{\log\left(\frac{16n}{\delta}\right) + 2 \log(k)}{\log\left(\frac{16n}{\delta}\right) + 2 \log(k-1)} + 1 = (*) \end{aligned}$$

If $k = 2$, $(*) \leq 1 + 4 * \log(32)/\log(8) \leq 9$. Otherwise,

$$\begin{aligned} (*) &= \frac{4(\log\left(\frac{16n}{\delta}\right) + 2 \log(k))}{\log\left(\frac{16n}{\delta}\right) + 2 \log(k-1)} + 1 \\ &\leq \frac{4 \log(k)}{\log(k-1)} + 1 \\ &\leq 4 \cdot 2 + 1 = 9 \end{aligned}$$

Putting these pieces together,

$$\begin{aligned} &(H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ &\leq (4H_{\text{ME}}(n, 2^{-k+1}, 1/16) + 9\tau_{k-1} + 9\tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|) \\ &\leq 9 (H_{\text{ME}}(n, 2^{-k+1}, 1/16) + \tau_{k-1} + \tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|) \end{aligned}$$

□

E.4.7 Step 6: Bounding Equation (17)

Here, we introduce the round K_{Good} , when $G_{K_{\text{Good}}} = M_\epsilon$ at some point within the round. Using the result of the previous step, we may bound the total number of samples taken though this round, controlling Equation (17).

With the result of Step 5, we prove the following inequality.

Claim:

$$\begin{aligned} &\sum_{k=1}^{\infty} 2 \mathbb{1}[G_{k-1} \neq M_\epsilon] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ &\leq c \sum_{i \in M_\epsilon} \max \left\{ h\left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n}\right), \min \left[h\left(0.25\Delta_i, \frac{\delta}{2n}\right), h\left(0.25\frac{\tilde{\alpha}_\epsilon}{4 - 2\epsilon}, \frac{\delta}{2n}\right) \right] \right\} \\ &\quad + c |M_\epsilon^c \cap M_{\epsilon + \tilde{\alpha}_\epsilon/\mu_1}| h\left(\frac{\tilde{\alpha}_\epsilon}{4 - 2\epsilon}, \frac{\delta}{2n}\right) + c \sum_{i \in M_{\epsilon + \tilde{\alpha}_\epsilon/\mu_1}^c} h\left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n}\right) \end{aligned} \tag{22}$$

Proof. Recall $t_k = \max\{t : t \in k\}$ denotes the maximum value of t in round k and $T_{\text{max}} = \max_{i \in M_\epsilon} T_i$ denotes the minimum t such that $G_k(t) = M_\epsilon$. Define the random round

$$K_{\text{Good}} := \min\{k : G_k = M_\epsilon\} = \min\{k : t_k \geq T_{\text{max}}\}$$

By definition of K_{Good} ,

$$\begin{aligned} & \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} \neq M_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ &= \sum_{k=1}^{K_{\text{Good}}} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|). \end{aligned}$$

Next, applying Step 5, if $K_{\text{Good}} > 1$

$$\begin{aligned} & \sum_{k=1}^{K_{\text{Good}}} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ & \leq 18 \sum_{k=1}^{K_{\text{Good}}-1} \mathbb{1}[G_{k-2} \cup B_{k-2} \neq [n]] (H_{\text{ME}}(n, 2^{-k+1}, 1/16) + \tau_{k-1} + \tau_{k-1} |(G_{k-2} \cup B_{k-2})^c|). \end{aligned}$$

Observe that by lines 17 and 20 of FAREAST, for any round r and for any $t > t_{r-1}$,

$$\sum_{k=1}^{r-1} \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \leq \sum_{s=1}^t |\mathcal{A}(s)|.$$

By definition, for the round $K_{\text{Good}} - 1$, we see that $t_{(K_{\text{Good}}-1)} < T_{\text{max}}$. Applying the above inequality with the inequality proven in Step 4,

$$\begin{aligned} & 18 \sum_{k=1}^{K_{\text{Good}}-1} |(G_{k-1} \cup B_{k-1})^c| (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \leq 18 \sum_{s=1}^{T_{\text{max}}} |\mathcal{A}(s)| \\ & \leq 18 \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{4 - 2\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + 18 \sum_{i \in M_{\epsilon}^c \cap M_{\epsilon + \tilde{\alpha}_{\epsilon}/\mu_1}} h \left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n} \right) + 18 |M_{\epsilon}^c \cap M_{\epsilon + \tilde{\alpha}_{\epsilon}/\mu_1}| h \left(\frac{\tilde{\alpha}_{\epsilon}}{4 - 2\epsilon}, \frac{\delta}{2n} \right). \end{aligned}$$

Otherwise, if $K_{\text{Good}} = 1$, exactly $4c'n \log(16) + 32n \log(16n/\delta)$ samples are given to the good filter in round 1. One may use Lemma F.2 to invert $h(\cdot, \cdot)$ and show that the summation on the right had side of the above inequality is within a constant of this and the claim holds in this case as well for a different constant, potentially larger than 18. \square

E.4.8 Step 7: Bounding Equation (18)

Next, we bound $\sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|)$.

$$\begin{aligned} & \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\ &= \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(M_{\epsilon} \cup B_{k-1})^c|) \\ &= \sum_{k=1}^{\infty} 2\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |M_{\epsilon}^c \setminus B_{k-1}|) \\ &= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[G_{k-1} \cup B_{k-1} \neq [n]] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |M_{\epsilon}^c \setminus B_{k-1}|) \\ &\stackrel{\mathcal{E}_1, \mathcal{E}_2}{=} \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq M_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |M_{\epsilon}^c \setminus B_{k-1}|) \\ &= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq M_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq M_{\epsilon}^c] (\tau_k |M_{\epsilon}^c \setminus B_{k-1}|) \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq M_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} 2\tau_k |M_{\epsilon}^c \setminus B_{k-1}| \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[B_{k-1} \neq M_{\epsilon}^c] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in M_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&\leq \sum_{k=K_{\text{Good}}+1}^{\infty} 2|M_{\epsilon}^c \setminus B_{k-1}| (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in M_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in M_{\epsilon}^c} 2\mathbb{1}[i \notin B_{k-1}] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k) + \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in M_{\epsilon}^c} 2\tau_k \mathbb{1}[i \notin B_{k-1}] \\
&= \sum_{k=K_{\text{Good}}+1}^{\infty} \sum_{i \in M_{\epsilon}^c} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&= \sum_{i \in M_{\epsilon}^c} \sum_{k=K_{\text{Good}}+1}^{\infty} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
&\leq \sum_{i \in M_{\epsilon}^c} \sum_{k=1}^{\infty} 2\mathbb{1}[i \notin B_{k-1}] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \tag{23}
\end{aligned}$$

E.4.9 Step 8: Bounding the expected total number of samples drawn by FAREAST

Now we take expectations over the number of samples drawn. These expectations are conditional on the high probability event $\mathcal{E}_1 \cap \mathcal{E}_2$. The bound in step 5 holds deterministically conditioned on this event.

Note τ_k and $H_{\text{ME}}(n, 2^{-k}, 1/16)$ are deterministic constants for any k . Let all expectations be jointly over the random instance ν and the randomness in FAREAST.

$$\begin{aligned}
&\mathbb{E}[T | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1] = \\
&\sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_k \cup B_k \neq [n]] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&= \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} \neq M_{\epsilon}] \mathbb{1}[G_k \cup B_k \neq [n]] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\quad + \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_k \cup B_k \neq [n]] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\stackrel{\text{Step 6}}{\leq} c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{4 - 2\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\
&\quad + c \sum_{i \in M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}} h \left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n} \right) + c |M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}| h \left(\frac{\tilde{\alpha}_{\epsilon}}{4 - 2\epsilon}, \frac{\delta}{2n} \right) \\
&\quad + \sum_{k=1}^{\infty} 2\mathbb{E} \left[\mathbb{1}[G_{k-1} = M_{\epsilon}] \mathbb{1}[G_k \cup B_k \neq [n]] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1 \right] \\
&\quad (H_{\text{ME}}(n, 2^{-k}, 1/16) + \tau_k + \tau_k |(G_{k-1} \cup B_{k-1})^c|) \\
&\stackrel{\text{Step 7}}{\leq} c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{4 - 2\epsilon}, \frac{\delta}{2n} \right) \right] \right\}
\end{aligned}$$

$$\begin{aligned}
& + c \sum_{i \in M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}^c} h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right) + c|M_\epsilon^c \cap M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}| h\left(\frac{\tilde{\alpha}_\epsilon}{4-2\epsilon}, \frac{\delta}{2n}\right) \\
& + \sum_{i \in M_\epsilon^c} \sum_{k=1}^{\infty} 2\mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2] = 1] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\
& \stackrel{(a)}{=} c \sum_{i \in M_\epsilon} \max \left\{ h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right), \min \left[h\left(0.25\Delta_i, \frac{\delta}{2n}\right), h\left(\frac{\tilde{\alpha}_\epsilon}{4-2\epsilon}, \frac{\delta}{2n}\right) \right] \right\} \\
& + c \sum_{i \in M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}^c} h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n}\right) + c|M_\epsilon^c \cap M_{\epsilon+\tilde{\alpha}_\epsilon/\mu_1}| h\left(\frac{\tilde{\alpha}_\epsilon/\mu_1}{4-2\epsilon}, \frac{\delta}{2n}\right) \\
& + \sum_{i \in M_\epsilon^c} \sum_{k=1}^{\infty} 2\mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16))
\end{aligned}$$

where (a) follows from $\mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1 \cap \mathcal{E}_2]] = \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]]$ for $i \in M_\epsilon^c$, since the event $\{i \in B_{k-1}\}$ is independent of \mathcal{E}_2 for all $i \in M_\epsilon^c$. This can be observed since \mathcal{E}_2 deals only with independent samples taken of arms in M_ϵ .

E.4.10 Step 9: Bounding $\sum_{k=1}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16))$ for $i \in M_\epsilon^c$

Next, we bound the expectation remaining from step 8. In particular, this is the number of samples drawn by the bad filter to add arm $i \in M_\epsilon^c$ to B_k .

First, we bound the probability that for a given $i \in M_\epsilon^c$ and a given k $i \notin B_k$. Note that by Borel-Cantelli, this implies that the probability that i is never added to any B_k is 0.

Claim 1: For $i \in M_\epsilon^c$, $k \geq \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon\mu_1} \right) \right\rceil \implies \mathbb{E}_\nu [\mathbb{1}[i \notin B_k] | \mathbb{1}[\mathcal{E}_1]] \leq \left(\frac{1}{8}\right)^{k - \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon\mu_1} \right) \right\rceil}$

Proof. If $i \in B_k$, either the good or the bad filter may have added it. The behavior of the bad filter on arms in M_ϵ^c is independent of \mathcal{E}_1 . Hence.

$$\begin{aligned}
\mathbb{E}_\nu [\mathbb{1}[i \notin B_k] | \mathbb{1}[\mathcal{E}_1]] &= \mathbb{E}_\nu \left[\mathbb{1}[\hat{\mu}_i + C_{\delta/2n}(t) \geq L_{t_k}] \mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i \leq 2^{-(k+1)}(2-\epsilon)] | \mathbb{1}[\mathcal{E}_1] \right] \\
&\leq \mathbb{E}_\nu \left[\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i \leq 2^{-(k+1)}(2-\epsilon)] | \mathbb{1}[\mathcal{E}_1] \right] \\
&= \mathbb{E}_\nu \left[\mathbb{1}[\hat{\mu}_{i_k} - \hat{\mu}_i \leq 2^{-(k+1)}(2-\epsilon)] \right]
\end{aligned}$$

If $i \in B_{k-1}$ then $i \in B_k$ by definition. Otherwise, if $i \notin B_{k-1}$, by Hoeffding's Inequality conditional on the value of i_k and a sum over conditional probabilities as in step 0, with probability at least $1 - \frac{\delta}{4nk^2}$

$$|((1-\epsilon)\hat{\mu}_{i_k} - \hat{\mu}_i) - ((1-\epsilon)\mu_{i_k} - \mu_i)| \leq 2^{-(k+1)}$$

If **MedianElimination** also succeeds, the joint event of which occurs with probability $\frac{15}{16} \left(1 - \frac{\delta}{4nk^2}\right)$ by independence⁷,

$$\begin{aligned}
(1-\epsilon)\hat{\mu}_{i_k} - \hat{\mu}_i &\geq (1-\epsilon)\mu_{i_k} - \mu_i - 2^{-(k+1)} \\
&\geq (1-\epsilon)\mu_1 - \mu_i - 2^{-(k+1)}(2-\epsilon) \\
&= \Delta_i - \epsilon\mu_1 - 2^{-(k+1)}(2-\epsilon).
\end{aligned}$$

Then for $k \geq \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon\mu_1} \right) \right\rceil$,

$$(1-\epsilon)\hat{\mu}_{i_k} - \hat{\mu}_i \geq \Delta_i - \epsilon\mu_1 - 2^{-(k+1)}(2-\epsilon) \geq 2^{-(k+1)}(2-\epsilon),$$

⁷Note that the success of **MedianElimination** and the concentration of $(\hat{\mu}_{i_k} - \hat{\mu}_i)$ around $(\mu_{i_k} - \mu_i)$ are independent of the events \mathcal{E}_1 and \mathcal{E}_2 conditioned on in Step 8.

which implies that $i \in B_k$ by line 15 of FAREAST. In particular,

$$\mathbb{E} [\mathbb{1}[i \in B_k] | i \notin B_{k-1} \mathbb{1}[\mathcal{E}_1]] \geq \mathbb{E} [\hat{\mu}_{i_k} - \hat{\mu}_i > 2^{-(k+1)}(2 - \epsilon) | i \notin B_{k-1}, \mathbb{1}[\mathcal{E}_1]] \geq \frac{15}{16} \left(1 - \frac{\delta}{4nk^2}\right).$$

Furthermore, $i \notin B_0$ by definition. Then for $k \geq \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil$,

$$\begin{aligned} \mathbb{E} [\mathbb{1}[i \notin B_k] | \mathbb{1}[\mathcal{E}_1]] &= \mathbb{E} [\mathbb{1}[i \notin B_k] (\mathbb{1}[i \notin B_{k-1}] + \mathbb{1}[i \in B_{k-1}]) | \mathbb{1}[\mathcal{E}_1]] \\ &= \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] + \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \in B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \end{aligned}$$

Deterministically, $\mathbb{1}[i \notin B_k] \mathbb{1}[i \in B_{k-1}] = 0$. Therefore,

$$\begin{aligned} &\mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] + \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \in B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \\ &= \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \\ &= \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | i \notin B_{k-1}, \mathbb{1}[\mathcal{E}_1]] \mathbb{P}(i \notin B_{k-1} | \mathbb{1}[\mathcal{E}_1]) \\ &\quad + \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | i \in B_{k-1}, \mathbb{1}[\mathcal{E}_1]] \mathbb{P}(i \in B_{k-1} | \mathbb{1}[\mathcal{E}_1]) \\ &= \mathbb{E} [\mathbb{1}[i \notin B_k] \mathbb{1}[i \notin B_{k-1}] | i \notin B_{k-1}, \mathbb{1}[\mathcal{E}_1]] \mathbb{P}(i \notin B_{k-1} | \mathbb{1}[\mathcal{E}_1]) \\ &= \mathbb{E} [\mathbb{1}[i \notin B_k] | i \notin B_{k-1}, \mathbb{1}[\mathcal{E}_1]] \mathbb{E} [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \\ &\leq \left(\frac{1}{16} + \frac{\delta}{4nk^2} \right) \mathbb{E} [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]]. \end{aligned}$$

For $k < \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil$, trivially, $\mathbb{E} [\mathbb{1}[i \notin B_k] | \mathbb{1}[\mathcal{E}_1]] \leq 1$. Recall $\delta \leq 1/8$. For $k \geq \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil$,

$$\mathbb{E} [\mathbb{1}[i \notin B_k] | \mathbb{1}[\mathcal{E}_1]] \leq \prod_{s=\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil}^k \left(\frac{1}{16} + \frac{\delta}{2ns^2} \right) \leq \left(\frac{1}{8} \right)^{k - \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil}.$$

□

Claim 2: For $j \in M_\epsilon^c$, $\sum_{k=1}^{\infty} 2\mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \leq c'' \frac{4n(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} + c'' h \left(\frac{\Delta_i - \epsilon \mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right)$

Proof. This sum decomposes into two terms.

$$\begin{aligned} &\sum_{k=1}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\ &= \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\quad + \sum_{k=\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \end{aligned}$$

We begin by bounding the first term.

$$\begin{aligned} &\sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\leq \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\ &\leq \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} \left(c' n 2^{2k} \log(16) + 2 + 2^{2k+4} \log \left(\frac{16nk^2}{\delta} \right) \right) \end{aligned}$$

$$\begin{aligned}
&\leq 2 \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) + \left(c'n \log(16) + 16 \log \left(\frac{16n}{\delta} \right) \right) \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} 2^{2k} \\
&\quad + 32 \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} 2^{2k} \log(k) \\
&\leq 2 \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) + \left(c'n \log(16) + 16 \log \left(\frac{16n}{\delta} \right) + 32 \log \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right) \sum_{k=1}^{\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil} 2^{2k} \\
&\leq 2 \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) + \frac{(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \left(c'n \log(16) + 32 \log \left(\frac{16n}{\delta} \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right) \right)
\end{aligned}$$

Next, we plug in the bound from claim 1 controlling the probability that $i \notin B_k$.

Using Claim 1, we bound the second sum as follows:

$$\begin{aligned}
&\sum_{r=\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil}^{\infty} \mathbb{E}_\nu [\mathbb{1}[i \notin B_{k-1}] | \mathbb{1}[\mathcal{E}_1]] \left(H_{\text{ME}}(n, 2^{-k}, 1/16) + 2 \left\lceil 2^{2k+3} \log \left(\frac{16nk^2}{\delta} \right) \right\rceil \right) \\
&\leq \sum_{k=\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil}^{\infty} \left(\frac{1}{8} \right)^{k-\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil-1} \left(c'n 2^{2k} \log(16) + 2 + 2^{2k+4} \log \left(\frac{16nk^2}{\delta} \right) \right) \\
&= c'n \log(16) \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} 2^{2(k+\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil)} + 2 \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} \\
&\quad + 16 \sum_{k=1}^{\infty} \left(\frac{1}{8} \right)^{k-1} 2^{2(k+\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil)} \log \left(\frac{16n \left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right)^2}{\delta} \right) \\
&\quad + 16 \sum_{k=1}^{\infty} 2^{-3k+3} 2^{2(k+\left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil+1)} \log \left(\frac{16n \left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right)^2}{\delta} \right) \\
&= 3 + \left(c'n \log(16) \frac{2^5(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} + \frac{2^9(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \log \left(\frac{16n}{\delta} \right) \right) \sum_{k=1}^{\infty} 2^{-k} \\
&\quad + \frac{2^9(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \sum_{k=1}^{\infty} 2^{-k} \log \left(\left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right)^2 \right) \\
&\leq 3 + c'n \log(16) \frac{2^5(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} + \frac{2^9(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \log \left(\frac{16n}{\delta} \right) \\
&\quad + \frac{2^{10}(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right) \\
&= (**)
\end{aligned}$$

We may bound the final summand, $\sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right)$ as follows:

$$\sum_{k=1}^{\infty} 2^{-k} \log \left(k + \left\lceil \log_2 \left(\frac{2-\epsilon}{\Delta_i - \epsilon \mu_1} \right) \right\rceil \right) \leq \log \left(\frac{e}{2} \log_2 \left(\frac{16(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \right) \right)$$

Plugging this back into (**), we have that

$$(**) \leq 3 + c'n \log(16) \frac{2^5(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} + \frac{2^9(2-\epsilon)^2}{(\Delta_i - \epsilon \mu_1)^2} \log \left(\frac{16n}{\delta} \right)$$

$$+ \frac{2^{10}(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} \log \left(\frac{e}{2} \log_2 \left(\frac{16(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} \right) \right)$$

Combining the above with the bound on the first sum, we have that

$$\begin{aligned} & \sum_{k=1}^{\infty} \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)) \\ & \leq c'' \left(\frac{4n(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} + \frac{4c(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{4-2\epsilon}{(\Delta_i - \epsilon\mu_1)^2} \right) \right) \right) \\ & = \frac{4c''n(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} + c''h \left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right) \end{aligned}$$

for a sufficiently large, universal constant c'' and c from the definition of $h(\cdot, \cdot)$. \square

E.4.11 Step 10: Applying the result of Step 9 to the result of Step 8

We may repeat the result of step 9 for every $i \in M_{\epsilon}^c$ and plug this into the result of Step 8. From this point, we simplify to return the final result.

By Step 8, the total number of samples T drawn by FAREAST is bounded in expectation by

$$\begin{aligned} \mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}/\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + c \sum_{i \in M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}} h \left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{2n} \right) + c|M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}}| h \left(\frac{\tilde{\alpha}_{\epsilon}}{4-2\epsilon}, \frac{\delta}{2n} \right) \\ & \quad + 2 \sum_{i \in M_{\epsilon}^c} \sum_{k=1}^{\infty} \mathbb{E}_{\nu} [\mathbb{1}[i \notin B_{k-1}] \mathbb{1}[\mathcal{E}_1]] (2\tau_k + H_{\text{ME}}(n, 2^{-k}, 1/16)). \end{aligned}$$

Applying the bound from Step 9 to each $i \in M_{\epsilon}^c$, we have that

$$\begin{aligned} \mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{4-2\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + c \sum_{i \in M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}} h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{2n} \right) + c|M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}| h \left(\frac{\tilde{\alpha}_{\epsilon}}{4-2\epsilon}, \frac{\delta}{2n} \right) \\ & \quad + 2c'' \sum_{i \in M_{\epsilon}^c} \frac{4n(2-\epsilon)^2}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right). \end{aligned}$$

For $i \in M_{\epsilon}^c \cap M_{\epsilon+\tilde{\alpha}_{\epsilon}/\mu_1}$, $\tilde{\alpha}_{\epsilon} = \min_{j \in M_{\epsilon}} \epsilon\mu_1 - \Delta_j \geq \Delta_i - \epsilon\mu_1$. By monotonicity of $h(\cdot, \cdot)$, $h \left(\frac{\tilde{\alpha}_{\epsilon}}{4-2\epsilon}, \frac{\delta}{2n} \right) \leq \frac{c''n(4-2\epsilon)}{(\Delta_i - \epsilon\mu_1)^2} + c''h \left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right)$. Therefore,

$$\begin{aligned} \mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{4-2\epsilon}, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + (2c'' + c) \sum_{i \in M_{\epsilon}^c} \frac{n(4-2\epsilon)}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{2n} \right). \end{aligned}$$

Lastly, note that $\frac{1}{3(1-x)} \leq \frac{1}{2-x}$ for $x \leq 1/2$. By monotonicity of h , we may lower bound the denominators $\frac{1}{4-2\epsilon}$ and $\frac{1}{2(2-\epsilon+\gamma)}$ as $\frac{1}{6(1-\epsilon)}$ and $\frac{1}{6(1-\epsilon+\gamma)}$ respectively. Since $\epsilon \in (0, 1/2]$, $\frac{1}{4-2\epsilon} \leq 1/4$. Plugging this in, we see that

$$\begin{aligned} \mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in M_{\epsilon}} \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_{\epsilon}}{6(1-\epsilon)}, \frac{\delta}{2n} \right) \right] \right\} \\ & \quad + (2c'' + c) \sum_{i \in M_{\epsilon}^c} \frac{4n}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right). \end{aligned}$$

Next, we use Lemma F.3 to bound the minimum of $h(\cdot, \cdot)$ functions.

$$\begin{aligned}
& c \sum_{i \in M_\epsilon} \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{2n} \right), h \left(\frac{\tilde{\alpha}_\epsilon}{6(1-\epsilon)}, \frac{\delta}{2n} \right) \right] \right\} \\
& \quad + (2c'' + c) \sum_{i \in M_\epsilon^c} \frac{4n}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right) \\
& = c \sum_{i \in M_\epsilon} \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{2n} \right) \right\} \\
& \quad + (2c'' + c) \sum_{i \in M_\epsilon^c} \frac{4n}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right)
\end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $|\epsilon\mu_1 - \Delta_i| \leq 6$ for all i and $\epsilon \leq 1/2$ implies that $\frac{1}{6(1-\epsilon)}|\epsilon\mu_1 - \Delta_i| \leq 2$ and $\frac{1}{6(1-\epsilon)} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq 2$. $\Delta_i \leq 8$ for all i , gives $0.25\Delta_i \leq 2$. Lastly, $\gamma \leq 6/\mu_1$ implies that $\frac{\gamma\mu_1}{6(1-\epsilon+\gamma)} \leq 2$. Therefore,

$$\begin{aligned}
\mathbb{E}[T|\mathcal{E}_1 \cap \mathcal{E}_2] & \leq c \sum_{i \in M_\epsilon} \max \left\{ h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{2n} \right) \right\} \\
& \quad + (2c'' + c) \sum_{i \in M_\epsilon^c} \frac{4n}{(\Delta_i - \epsilon\mu_1)^2} + h \left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{2n} \right) \\
& \leq c \sum_{i \in M_\epsilon} \max \left\{ \frac{64}{(\epsilon\mu_1 - \Delta_i)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{384n}{\delta(\epsilon\mu_1 - \Delta_i)^2} \right) \right), \right. \\
& \quad \left. \frac{576}{\left(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} \right)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{1728n}{\delta \left(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} \right)^2} \right) \right) \right\} \\
& \quad + (2c'' + c) \sum_{i \in M_\epsilon^c} \frac{4n}{(\Delta_i - \epsilon\mu_1)^2} + \frac{64}{(\epsilon\mu_1 - \Delta_i)^2} \log \left(\frac{4n}{\delta} \log_2 \left(\frac{384n}{\delta(\epsilon\mu_1 - \Delta_i)^2} \right) \right) \\
& \leq c_6 \sum_{i=1}^n \max \left\{ \frac{1}{(\epsilon\mu_1 - \Delta_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\epsilon\mu_1 - \Delta_i)^2} \right) \right), \right. \\
& \quad \left. \frac{1}{\left(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} \right)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta \left(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} \right)^2} \right) \right) \right\} \\
& \quad + c_6 \sum_{i \in M_\epsilon^c} \frac{n}{(\Delta_i - \epsilon\mu_1)^2} \\
& = c_6 \sum_{i=1}^n \max \left\{ \frac{1}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \\
& \quad \left. \frac{1}{\left(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i \right)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta \left(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i \right)^2} \right) \right) \right\} \\
& \quad + c_6 \sum_{i \in M_\epsilon^c} \frac{n}{((1-\epsilon)\mu_1 - \mu_i)^2}
\end{aligned}$$

for a sufficiently large constant c_6 .

E.4.12 Step 11: High probability sample complexity bound

Finally, the Good Filter is equivalent to EAST, Algorithm 5, except split across rounds. EAST is an elimination algorithm. Note that the Good Filter is union bounded over $2n$ events whereas the bounds in EAST are union bounded over n events. The Good Filter and Bad Filter are given the same number of samples in each round, and the Good Filter can terminate within a round, conditioned on $\mathcal{E}_1 \cap \mathcal{E}_2$. Therefore, we can bound the complexity of FAREAST in terms of that of EAST run at failure probability $\delta/2$. If FAREAST terminates in the second round or later, the arguments in Steps 4 and 5 can be used to show that FAREAST draws no more than a factor of 18 more samples than EAST, though this estimate is highly pessimistic. If FAREAST terminates in round 1 (when gaps are large), we may still show that this is within a constant factor of the complexity of EAST, but the story is more complicated. In the first round, the bad filter draws at most $c'n \log(16) + 16(n+1) \log(8n/\delta)$ samples where c' is the constant from Median Elimination. Since we have assumed that $\max(\Delta_i, |\epsilon\mu_1 - \Delta_i|) \leq 6(1-\epsilon) \leq 6$, this sum is likewise within a constant factor of the complexity of EAST. Hence with probability at least $1 - \delta$, by Theorem E.4,

$$T \leq c_5 \sum_{i=1}^n \min \left\{ \max \left\{ \frac{1}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right. \\ \left. \frac{1}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right), \right. \\ \left. \frac{1}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \right) \right) \right\}, \\ \left. \frac{(1-\epsilon+\gamma)^2}{\gamma^2 \mu_1^2} \log \left(\frac{n}{\delta} \log_2 \left(\frac{(1-\epsilon+\gamma)^2 n}{\delta \gamma^2 \mu_1^2} \right) \right) \right\}$$

samples for a sufficiently large constant c_5 .

□

E.5 An elimination algorithm for all ϵ

First, we state an elimination algorithm EAST (Elimination Algorithm for a Sampled Threshold) and bound its sample complexity. EAST is equivalent to the good filter in FAREAST. At all times, EAST maintains an active set \mathcal{A} and samples all arms $i \in \mathcal{A}$, progressively eliminating arms from \mathcal{A} until termination occurs. Additionally, EAST maintains upper and lower bounds, denoted U_t and L_t , on the threshold, $\mu_1 - \epsilon$ in the additive case and $(1-\epsilon)\mu_1$ in the multiplicative case. If $\hat{\mu}_i(t) + C_{\delta/n}(t) < L_t$, EAST may infer that $i \notin G_\epsilon$ (resp. $i \notin M_\epsilon$) and accordingly removes i from \mathcal{A} . If $\hat{\mu}_i(t) - C_{\delta/n}(t) > U_t$, EAST may infer that $i \in G_\epsilon$ (resp. $i \in M_\epsilon$) and adds i to a set G of good arms it has found so far. However, a good arm $i \in G$ is only removed from \mathcal{A} , if EAST can also certify that it is not the best arm, namely if $\hat{\mu}_i(t) + C_{\delta/n}(t) < \max_j \hat{\mu}_j(t) - C_{\delta/n}(t)$. This ensures that $\mu_1 - \epsilon \in [L_t, U_t]$ at all times in the additive case, and similarly, $(1-\epsilon)\mu_1 \in [L_t, U_t]$ in the multiplicative case. If $\mathcal{A} \subset G$, EAST may declare that $G = G_\epsilon$ (resp. $G = M_\epsilon$) and terminates. Otherwise, the algorithm terminates when $U_t - L_t < \gamma/2$ and returns $\mathcal{A} \cup G$ in the additive case or when $U_t - L_t < \frac{\gamma}{2-\epsilon} L_t$ in the multiplicative case. This limits the number of samples of any arm and ensures that no arm worse than $(\epsilon + \gamma)$ -good is returned. We give pseudocode for EAST in Algorithm 5. Pieces specific to the additive case are shown in red, and pieces specific to the multiplicative case are shown in blue.

Recall that $\alpha_\epsilon = \min_{i \in G_\epsilon} \epsilon - \Delta_i$ and $\beta_\epsilon = \min_{i \in G_\epsilon^c} \Delta_i - \epsilon$.

Theorem E.3. Fix $\epsilon > 0$, $0 < \delta \leq 1/2$, $\gamma \in [0, 8]$ and an instance ν such that $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i . In the case that $G_\epsilon = [n]$, let $\alpha_\epsilon = \min(\alpha_\epsilon, \beta_\epsilon)$. With probability at least $1 - \delta$, EAST returns a set G such that $G_\epsilon \subset G \subset G_{(\epsilon+\gamma)}$ in at most

$$\sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\ \left. \frac{256}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right) \right\},$$

Algorithm 5 EAST : Elimination Algorithm for a Sampled Threshold

Require: $\epsilon, \delta > 0$, slack $\gamma \geq 0$, (if multiplicative, $0 < \epsilon \leq 1/2$)

```

1: Let  $\mathcal{A} \leftarrow [n]$  be the active set, and  $G \leftarrow \emptyset$  be the set of  $\epsilon$ -good arms found so far, Let  $t \leftarrow 0$ 
2: while  $\mathcal{A} \not\subset G$  and  $U_t - L_t \geq \gamma/2$  or  $U_t - L_t \geq \frac{\gamma}{2-\epsilon} L_t$  do
3:   Pull each arm  $i \in \mathcal{A}$  and update its empirical mean  $\hat{\mu}_i(t)$ , Update  $t \leftarrow t + 1$ 
4:   Update  $U_t \leftarrow \max_j \hat{\mu}_j(t) + C_{\delta/n}(t) - \epsilon$  or  $U_t \leftarrow (1 - \epsilon) (\max_j \hat{\mu}_j(t) + C_{\delta/n}(t))$ 
5:   Update  $L_t \leftarrow \max_j \hat{\mu}_j(t) - C_{\delta/n}(t) - \epsilon$  or  $L_t \leftarrow (1 - \epsilon) (\max_j \hat{\mu}_j(t) - C_{\delta/n}(t))$ 
6:   for  $i \in \mathcal{A}$  do
7:     if  $\hat{\mu}_i(t) - C_{\delta/n}(t) > U_t$  then
8:       add  $i$  to  $G$  ▷ Arm  $i$  is good
9:     if  $\hat{\mu}_i(t) + C_{\delta/n}(t) < L_t$  then
10:      Remove  $i$  from  $\mathcal{A}$  ▷ Arms in  $G_\epsilon^c$  or  $M_\epsilon^c$  are removed
11:     if  $i \in G$  and  $\hat{\mu}_i(t) + C_{\delta/n}(t) < \max_j \hat{\mu}_j(t) - C_{\delta/n}(t)$  then
12:      Remove  $i$  from  $\mathcal{A}$  ▷ Arms in  $G_\epsilon$  or  $M_\epsilon$  are removed
return  $G \cup \mathcal{A}$ 

```

$$\frac{256}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \Bigg\},$$

$$\frac{64}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta\gamma^2} \right) \right) \Bigg\}$$

samples.

Additionally, in the multiplicative case, recall that $\tilde{\alpha}_\epsilon = \min_{i \in G_\epsilon} \epsilon - \Delta_i$ and $\tilde{\beta}_\epsilon = \min_{i \in G_\epsilon^c} \Delta_i - \epsilon$. Next, we a theorem bounding the complexity of EAST in the **multiplicative** regime.

Theorem E.4. Fix $\epsilon, \delta \in (0, 1/2]$, $\gamma \in [0, \min(1, 6/\mu_1))$ and an instance ν such that $\max(\Delta_i, |\epsilon\mu_1 - \Delta_i|) \leq 6$ for all i . Assume that $\mu_1 \geq 0$. In the case that $M_\epsilon = [n]$, let $\tilde{\alpha}_\epsilon = \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)$. With probability at least $1 - \delta$, EAST returns a set G such that $M_\epsilon \subset G \subset M_{(\epsilon+\gamma)}$ in at most

$$\sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right.$$

$$\frac{576}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right),$$

$$\left. \frac{576}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \right) \right) \right\},$$

$$\frac{144(1-\epsilon+\gamma)}{\gamma^2\mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{432(1-\epsilon+\gamma)n}{\delta\gamma^2\mu_1^2} \right) \right) \Bigg\}$$

samples.

E.6 Proof of Theorem E.3 EAST in the **additive** regime

Proof. Notation for the proof: Throughout, recall $\Delta_i = \mu_1 - \mu_i$. Recall that t counts the number of times each arm in \mathcal{A} has been sampled and thus the number of times that the conditionals in Lines 10 and 11 have been evaluated. Let $\mathcal{A}(t)$ denote the state \mathcal{A} at this time before the arms have been eliminated from \mathcal{A} in lines 10 and 11. Let $G(t)$ be defined similarly. Therefore, the total number of samples drawn by EAST up to time t is $\sum_{s=1}^t |\mathcal{A}(s)|$.

For $i \in G_\epsilon$, let T_i denote the random variable of the number of times arm i is sampled before it is added to G in Line 8. For $i \in G_\epsilon^c$, let T_i denote the random variable of the number of times arm i is sampled before it is removed from \mathcal{A} in Line 10. For any arm i , let T'_i denote the random variable of the of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t)$.

Define the event

$$\mathcal{E} = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_{\delta}(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\begin{aligned} \mathbb{P}(\mathcal{E}^c) &= \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \\ &\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{n} = \delta \end{aligned}$$

Hence, $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$.

E.6.1 Step 0: Correctness

Claim 0: On \mathcal{E} , first we prove that $G(t) \subset G_{\epsilon}$ for all $t \in \mathbb{N}$.

In particular, this shows that EAST never incorrectly add arms in G_{ϵ}^c to the set G .

Proof. We begin by showing that on \mathcal{E} the best arm is never removed from \mathcal{A} for all t . Note for any i

$$\hat{\mu}_1 + C_{\delta/n}(t) \geq \mu_1 \geq \mu_i \geq \hat{\mu}_i(t) - C_{\delta/n}(t) > \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon.$$

In particular this shows, $\hat{\mu}_1 + C_{\delta/n}(t) > \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon = L_t^*$ and $\hat{\mu}_1 + C_{\delta/n}(t) \geq \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/n}(t)$ showing that 1 will never exit \mathcal{A} in line 11.

Secondly, we show that at all times t , $\mu_1 - \epsilon \in [L_t, U_t]$. By the above, since μ_1 never leaves \mathcal{A} ,

$$U_t = \max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/n}(t) - \epsilon \geq \hat{\mu}_1(t) + C_{\delta/n}(t) - \epsilon \geq \mu_1 - \epsilon$$

and for any i ,

$$\mu_1 - \epsilon \geq \mu_i - \epsilon \geq \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon$$

Hence $\mu_1 - \epsilon \geq \max_i \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon = L_t$.

Next, we show that $G(t) \subset G_{\epsilon}$ for all $t \geq 1$. Suppose not. Then $\exists t \in \mathbb{N}$ and $\exists i \in G_{\epsilon}^c \cap G(t)$ such that,

$$\mu_i \geq \hat{\mu}_i(t) - C_{\delta/n}(t) \geq U_t \geq \mu_1 - \epsilon > \mu_i,$$

with the last inequality following from the previous assertion, giving a contradiction. \square

Claim 1: Next, we show that on \mathcal{E} , $G_{\epsilon} \subset \mathcal{A}(t) \cup G(t)$ for all $t \in \mathbb{N}$.

In particular this implies that if $\mathcal{A} \subset G$, then $G_{\epsilon} \subset G$. Combining this with the previous claim gives $G \subset G_{\epsilon} \subset G$, hence $G = G_{\epsilon}$. On this condition, EAST terminates by line 2 and returns the set $\mathcal{A} \cup G = G$. Note that by definition, $G_{\epsilon} \subset G_{(\epsilon+\gamma)}$ for all $\gamma \geq 0$. Therefore EAST terminates correctly on this condition.

Proof. Suppose for contradiction that there exists $i \in G_{\epsilon}$ such that $i \notin \mathcal{A}(t) \cup G(t)$. This occurs only if i is eliminated in line 10. Hence, there exists a $t' \leq t$ such that $\hat{\mu}_i(t') + C_{\delta/n}(t') < L_{t'}$. Therefore, on the event \mathcal{E} ,

$$\mu_1 - \epsilon \stackrel{\mathcal{E}}{\geq} L_{t'} = \max_{j \in \mathcal{A}} \hat{\mu}_j(t') - C_{\delta/n}(t') - \epsilon > \hat{\mu}_i(t') + C_{\delta/n}(t') \stackrel{\mathcal{E}}{\geq} \mu_i$$

which contradicts $i \in G_{\epsilon}$. \square

Claim 2: Finally, we show that if $U_t - L_t \leq \gamma/2$, then $\mathcal{A} \cup G \subset G_{(\epsilon+\gamma)}$.

Combining with the previous that $G_{\epsilon} \subset \mathcal{A} \cup G$, if EAST terminates on this condition by line 2, it does so correctly.

Proof. Assume $U_t - L_t \leq \gamma/2$. This implies that

$$\left(\max_{i \in A(t)} \hat{\mu}_i(t) + C_{\delta/n}(t) - \epsilon\right) - \left(\max_{i \in A(t)} \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon\right) = 2C_{\delta/n}(t) \leq \gamma/2.$$

Suppose for contradiction that there exists $i \in G_{(\epsilon+\gamma)}^c$ such that $i \in \mathcal{A} \cup G$. Since $G_\epsilon \cap G_{(\epsilon+\gamma)}^c = \emptyset$ and we have previously shown that $G(t) \subset G_\epsilon$ for all t , we have that $i \in A \setminus G$. Therefore, by the condition in line 10, $\hat{\mu}_i(t) + C_{\delta/n}(t) \geq L_t$. Hence, $\mu_i + 2C_{\delta/n}(t) \stackrel{\mathcal{E}}{\geq} \hat{\mu}_i(t) + C_{\delta/n}(t) \geq L_t$. By assumption, we have that $U_t - \gamma/2 \leq L_t$, and the event \mathcal{E} implies that $U_t \geq \mu_1 - \epsilon$. Therefore, $\mu_i + 2C_{\delta/n}(t) \geq U_t - \gamma/2 \geq \mu_1 - \epsilon - \gamma/2$. Combining this with the inequality $2C_{\delta/n} \leq \gamma/2$, we have that

$$\gamma \geq 2C_{\delta/n}(t) + \gamma/2 \geq \mu_1 - \epsilon - \mu_i \stackrel{i \in G_{(\epsilon+\gamma)}^c}{>} \gamma$$

which is a contradiction. \square

Therefore, on the event \mathcal{E} , if EAST terminates due to either condition in line 2, it returns $\mathcal{A} \cup G$ such that $G_\epsilon \subset \mathcal{A} \cup G \subset G_{(\epsilon+\gamma)}$. Since $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$, EAST terminates correctly with probability at least $1 - \delta$.

E.6.2 Step 1: Controlling the total number of samples given by EAST to arms in G_ϵ

To keep track of the number of samples that arms are given by EAST, we introduce random variables T_i and T'_i for all $i \in [n]$. When arm i has been given $\max(T_i, T'_i)$ samples it is removed from \mathcal{A} in line 11.

By Step 0, only arms in G_ϵ are added to G . Therefore, T_i is defined as

$$T_i = \min \left\{ t : \begin{array}{ll} i \in G_k(t+1) & \text{if } i \in G_\epsilon \\ i \notin \mathcal{A}(t+1) & \text{if } i \in G_\epsilon^c \end{array} \right\} \stackrel{\mathcal{E}}{=} \min \left\{ t : \begin{array}{ll} \hat{\mu}_i - C_{\delta/n}(t) \geq U_t & \text{if } i \in G_\epsilon \\ \hat{\mu}_i + C_{\delta/n}(t) \leq L_t & \text{if } i \in G_\epsilon^c \end{array} \right\} \quad (24)$$

Similarly, recall T'_i denotes the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t)$. Hence,

$$T'_i = \min \left\{ t : \hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/n}(t) \right\} \quad (25)$$

Claim 0: For $i \in G_\epsilon$, we have that $T_i \leq h(0.25(\epsilon - \Delta_i), \delta/n)$.

Proof. Note that, $4C_{\delta/n}(t) \leq \mu_i - (\mu_1 - \epsilon)$, true when $t > h(0.25(\epsilon - \Delta_i), \frac{\delta}{n})$, implies that for all j ,

$$\begin{aligned} \hat{\mu}_i(t) - C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\geq} \mu_i - 2C_{\delta/n}(t) \\ &\geq \mu_1 + 2C_{\delta/n}(t) - \epsilon \\ &\geq \mu_j + 2C_{\delta/n}(t) - \epsilon \\ &\stackrel{\mathcal{E}}{\geq} \hat{\mu}_j(t) + C_{\delta/n}(t) - \epsilon \end{aligned}$$

so in particular, $\hat{\mu}_i(t) - C_{\delta/n}(t) \geq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/n}(t) - \epsilon = U_t$. \square

Claim 1: For $i \in G_\epsilon$, we have that $T'_i \leq h(0.25\Delta_i, \delta/n)$.

Proof. Note that $4C_{\delta/n}(t) \leq \mu_1 - \mu_i$, true when $t > h(0.25\Delta_i, \frac{\delta}{n})$, implies that

$$\begin{aligned} \hat{\mu}_i(t) + C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(t) \\ &\leq \mu_1 - 2C_{\delta/n}(t) \\ &\stackrel{\mathcal{E}}{\leq} \hat{\mu}_1(t) - C_{\delta/n}(t). \end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Hence, $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/n}(t)$. \square

E.6.3 Step 2: Controlling the total number of samples given by EAST to arms in G_ϵ^c

Claim: Next, we show that $T_i \leq h(0.25(\epsilon - \Delta_i), \frac{\delta}{n})$ for $i \in G_\epsilon^c$

Proof. Note that, $4C_{\delta/n}(t) \leq \mu_1 - \epsilon - \mu_i$, true when $t > h(0.25(\epsilon - \Delta_i), \frac{\delta}{n})$, implies that

$$\begin{aligned} \hat{\mu}_i(t) + C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(t) \\ &\leq \mu_1 - 2C_{\delta/n}(t) - \epsilon \\ &\stackrel{\mathcal{E}}{\leq} \hat{\mu}_1(t) - C_{\delta/n}(t) - \epsilon \end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Therefore $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t) - \epsilon = L_t$. \square

E.6.4 Step 3: Bounding the total number of samples drawn by EAST

With the results of Steps 1 and 2, we may bound the total sample complexity of EAST. Note that independently of the event \mathcal{E} , EAST terminates if $U_t - L_t \leq \gamma/2$. Let the random variable of the maximum number of samples given to any arm before this occurs be T_γ . Additionally, EAST may terminate if $\mathcal{A} \subset G$. Let the random variable of maximum number of samples given to any arm before this occurs be $T_{\alpha_\epsilon \beta_\epsilon}$. Note that due to the sampling procedure, the total number of samples drawn by EAST at termination may be written as $\sum_{t=1}^{\min(T_\gamma, T_{\alpha_\epsilon \beta_\epsilon})} |\mathcal{A}(t)|$.

Now we bound $\sum_{t=1}^{\min(T_\gamma, T_{\alpha_\epsilon \beta_\epsilon})} |\mathcal{A}(t)|$. Let $S_i = \min\{t : i \notin \mathcal{A}(t+1)\}$. Hence,

$$\sum_{t=1}^{\min(T_\gamma, T_{\alpha_\epsilon \beta_\epsilon})} |\mathcal{A}(t)| = \sum_{t=1}^{\min(T_\gamma, T_{\alpha_\epsilon \beta_\epsilon})} \sum_{i=1}^n \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \sum_{t=1}^{\min(T_\gamma, T_{\alpha_\epsilon \beta_\epsilon})} \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, S_i\}$$

For arms $i \in G_\epsilon^c$, $S_i = T_i$ by definition. For $i \in G_\epsilon$, $S_i = \max(T_i, T'_i)$ by line 11 of the algorithm. Then

$$\begin{aligned} \sum_{i=1}^n \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, S_i\} &= \sum_{i \in G_\epsilon} \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, \max(T_i, T'_i)\} + \sum_{i \in G_\epsilon^c} \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, T_i\} \\ &= \sum_{i \in G_\epsilon} \min\{T_\gamma, \min\{T_{\alpha_\epsilon \beta_\epsilon}, \max(T_i, T'_i)\}\} + \sum_{i \in G_\epsilon^c} \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, T_i\} \\ &= \sum_{i \in G_\epsilon} \min\{T_\gamma, \max\{T_i, \min(T'_i, T_{\alpha_\epsilon \beta_\epsilon})\}\} + \sum_{i \in G_\epsilon^c} \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, T_i\} \end{aligned}$$

We may define $T_\gamma := \min\{t : U_t - L_t \leq \gamma/2\}$. Note that $4C_{\delta/n}(t) \leq \gamma$, true when $t > h(0.25\gamma, \delta/n)$ implies that

$$U_t - L_t = (\max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t) + C_{\delta/n}(t) - \epsilon) - (\max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t) - C_{\delta/n}(t) - \epsilon) = 2C_{\delta/n}(t) \leq \gamma/2.$$

Therefore, we have that $T_\gamma \leq h(0.25\gamma, \delta/n)$.

Next, we may define $T_{\alpha_\epsilon \beta_\epsilon} = \min\{t : \mathcal{A}(t) \subset G_\epsilon\}$. By step 0, on the event \mathcal{E} , $\mathcal{A} \subset G$ implies that $G = G_\epsilon$. Therefore, $T_{\alpha_\epsilon \beta_\epsilon}$ may be equivalently defined as $T_{\alpha_\epsilon \beta_\epsilon} = \min\{t : G(t) = G_\epsilon \text{ and } G_\epsilon^c \cap \mathcal{A} = \emptyset\}$. Recalling the definition of T_i , we see that $T_{\alpha_\epsilon \beta_\epsilon} = \max_i(T_i)$.

Recall that by steps 1 and 2, $T_i \leq h(0.25(\epsilon - \Delta_i), \frac{\delta}{n})$ and $T'_i \leq h(0.25\Delta_i, \frac{\delta}{n})$. Furthermore, by monotonicity of $h(\cdot, \cdot)$, this implies that $T_{\alpha_\epsilon \beta_\epsilon} = h(0.25 \min(\alpha_\epsilon, \beta_\epsilon), \delta/n)$. Plugging this in, we see that

$$\begin{aligned} &\sum_{i \in G_\epsilon} \min\{T_\gamma, \max\{T_i, \min(T'_i, T_{\alpha_\epsilon \beta_\epsilon})\}\} + \sum_{i \in G_\epsilon^c} \min\{T_\gamma, T_{\alpha_\epsilon \beta_\epsilon}, T_i\} \\ &= \sum_{i \in G_\epsilon} \min\{T_\gamma, \max\{T_i, \min(T'_i, T_{\alpha_\epsilon \beta_\epsilon})\}\} + \sum_{i \in G_\epsilon^c} \min\{T_\gamma, T_i\} \\ &\leq \sum_{i \in G_\epsilon} \min \left\{ \max \left\{ h \left(0.25(\epsilon - \Delta_i), \frac{\delta}{n} \right), \min \left[h \left(0.25\Delta_i, \frac{\delta}{n} \right), h \left(0.25 \min(\alpha_\epsilon, \beta_\epsilon), \frac{\delta}{n} \right) \right] \right\} \right\}, \end{aligned}$$

$$\begin{aligned}
& h\left(0.25\gamma, \frac{\delta}{n}\right)\Big\} \\
& + \sum_{i \in G_\epsilon^c} \min \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{n}\right), h\left(0.25 \min(\alpha_\epsilon, \beta_\epsilon), \frac{\delta}{n}\right) \right\} \\
& = \sum_{i=1}^n \min \left\{ \max \left\{ h\left(0.25(\epsilon - \Delta_i), \frac{\delta}{n}\right), \min \left[h\left(0.25\Delta_i, \frac{\delta}{n}\right), h\left(0.25 \min(\alpha_\epsilon, \beta_\epsilon), \frac{\delta}{n}\right) \right] \right\}, \right. \\
& \quad \left. h\left(0.25\gamma, \frac{\delta}{n}\right) \right\}
\end{aligned}$$

where the final equality holds by definition for arms in G_ϵ . Next, by Lemma F.3, we may bound the minimum of $h(\cdot, \cdot)$ functions.

$$\begin{aligned}
& \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), \min \left[h\left(\frac{\Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\min(\alpha_\epsilon, \beta_\epsilon)}{4}, \frac{\delta}{n}\right) \right] \right\}, \right. \\
& \quad \left. h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right\} \\
& = \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), \right. \right. \\
& \quad \left. \min \left[h\left(\frac{\Delta_i}{4}, \frac{\delta}{n}\right), \max \left[h\left(\frac{\alpha_\epsilon}{4}, \frac{\delta}{n}\right), h\left(\frac{\beta_\epsilon}{4}, \frac{\delta}{n}\right) \right] \right] \right\}, \\
& \quad \left. h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right\} \\
& \leq \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), \right. \right. \\
& \quad \left. \max \left[h\left(\frac{\Delta_i + \alpha_\epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \beta_\epsilon}{8}, \frac{\delta}{n}\right) \right] \right\}, \\
& \quad \left. h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right\} \\
& = \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \alpha_\epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \beta_\epsilon}{8}, \frac{\delta}{n}\right) \right\}, \right. \\
& \quad \left. h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right\}
\end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $\max(\Delta_i, |\epsilon - \Delta_i|) \leq 8$ for all i , we have that $0.25\Delta_i \leq 2$, $0.25|\epsilon - \Delta_i| \leq 2$, and $0.25 \min(\alpha_\epsilon, \beta_\epsilon) \leq 2$. Therefore,

$$\begin{aligned}
& \sum_{i=1}^n \min \left\{ \max \left\{ h\left(\frac{\Delta_i - \epsilon}{4}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \alpha_\epsilon}{8}, \frac{\delta}{n}\right), h\left(\frac{\Delta_i + \beta_\epsilon}{8}, \frac{\delta}{n}\right) \right\}, \right. \\
& \quad \left. h\left(\frac{\gamma}{4}, \frac{\delta}{n}\right) \right\} \\
& \leq \sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{(\epsilon - \Delta_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta(\epsilon - \Delta_i)^2} \right) \right), \right. \right. \\
& \quad \frac{256}{(\Delta_i + \alpha_\epsilon)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\Delta_i + \alpha_\epsilon)^2} \right) \right), \\
& \quad \left. \frac{256}{(\Delta_i + \beta_\epsilon)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\Delta_i + \beta_\epsilon)^2} \right) \right) \right\},
\end{aligned}$$

$$\begin{aligned}
& \frac{64}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta \gamma^2} \right) \right) \Bigg\} \\
= & \sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{(\mu_1 - \epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 - \epsilon - \mu_i)^2} \right) \right), \right. \right. \\
& \frac{256}{(\mu_1 + \alpha_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 + \alpha_\epsilon - \mu_i)^2} \right) \right), \\
& \left. \frac{256}{(\mu_1 + \beta_\epsilon - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{768n}{\delta(\mu_1 + \beta_\epsilon - \mu_i)^2} \right) \right) \right\}, \\
& \left. \frac{64}{\gamma^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta \gamma^2} \right) \right) \right\}.
\end{aligned}$$

□

E.7 Proof of Theorem E.4, EAST in the multiplicative regime

Proof. Notation for the proof: Throughout, recall $\Delta_i = \mu_1 - \mu_i$. Recall that t counts the number of times each arm in \mathcal{A} has been sampled and thus the number of times that the conditionals in Lines 10 and 11 have been evaluated. Let $\mathcal{A}(t)$ denote the state \mathcal{A} at this time before the arms have been eliminated from \mathcal{A} in lines 10 and 11. Let $G(t)$ be defined similarly. Therefore, the total number of samples drawn by EAST up to time t is $\sum_{s=1}^t |\mathcal{A}(s)|$.

For $i \in M_\epsilon$, let T_i denote the random variable of the number of times arm i is sampled before it is added to G in Line 8. For $i \in M_\epsilon^c$, let T_i denote the random variable of the number of times arm i is sampled before it is removed from \mathcal{A} in Line 10. For any arm i , let T'_i denote the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t)$.

Define the event

$$\mathcal{E} = \left\{ \bigcap_{i \in [n]} \bigcap_{t \in \mathbb{N}} |\hat{\mu}_i(t) - \mu_i| \leq C_{\delta/n}(t) \right\}.$$

Using standard anytime confidence bound results, and recalling that $C_\delta(t) := \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}}$, we have

$$\begin{aligned}
\mathbb{P}(\mathcal{E}^c) &= \mathbb{P} \left(\bigcup_{i \in [n]} \bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \\
&\leq \sum_{i=1}^n \mathbb{P} \left(\bigcup_{t \in \mathbb{N}} |\hat{\mu}_i - \mu_i| > C_{\delta/n}(t) \right) \leq \sum_{i=1}^n \frac{\delta}{n} = \delta
\end{aligned}$$

Hence, $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$.

E.7.1 Step 0: Correctness

Claim 0: On \mathcal{E} , first we prove that $G(t) \subset M_\epsilon$ for all $t \in \mathbb{N}$.

In particular, this shows that EAST never incorrectly add arms in M_ϵ^c to the set G .

Proof. Firstly we show $1 \in \mathcal{A}$ for all $t \in \mathbb{N}$, namely the best arm is never removed from \mathcal{A} . Note for any i such that $\hat{\mu}_i(t) - C_{\delta/n}(t) \geq 0$,

$$\hat{\mu}_1 + C_{\delta/n}(t) \geq \mu_1 \geq \mu_i \geq \hat{\mu}_i(t) - C_{\delta/n}(t) > (1 - \epsilon)(\hat{\mu}_i(t) - C_{\delta/n}(t)).$$

For i such that $\hat{\mu}_i(t) - C_{\delta/n}(t) < 0$, if $\hat{\mu}_1 + C_{\delta/n}(t) \geq 0$, then

$$\hat{\mu}_1 + C_{\delta/n}(t) \geq 0 > (1 - \epsilon)(\hat{\mu}_i(t) - C_{\delta/n}(t)).$$

Note that $\hat{\mu}_1 + C_{\delta/n}(t) < 0$ implies on the event \mathcal{E} that $\mu_1 < 0$, which contradicts the assumption that $\mu_1 \geq 0$ made in the theorem. In particular this shows, $\hat{\mu}_1 + C_{\delta/n}(t) > (1 - \epsilon)(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/n}(t))$.

$C_{\delta/n}(t) = L_t$ and $\hat{\mu}_1 + C_{\delta/n}(t) \geq \max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/n}(t)$ showing that 1 will never exit \mathcal{A} in line 28.

Secondly, we show that at all times t , $(1 - \epsilon)\mu_1 \in [L_t, U_t]$. By the above, since μ_1 never leaves \mathcal{A} ,

$$U_t = (1 - \epsilon)(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/n}(t)) \geq (1 - \epsilon)(\hat{\mu}_1(t) + C_{\delta/n}(t)) \geq (1 - \epsilon)\mu_1$$

and for any i ,

$$(1 - \epsilon)\mu_1 \geq (1 - \epsilon)\mu_i \geq (1 - \epsilon)(\hat{\mu}_i(t) - C_{\delta/n}(t))$$

Hence $(1 - \epsilon)\mu_1 \geq (1 - \epsilon)(\max_i \hat{\mu}_i(t) - C_{\delta/n}(t)) = L_t$.

Next, we show that $G \subset M_\epsilon$ for all $k \geq 1, t \geq 1$. Suppose not. Then $\exists k, t \in \mathbb{N}$ and $\exists i \in M_\epsilon^c \cap G(t)$ such that,

$$\mu_i \geq \hat{\mu}_i(t) - C_{\delta/n}(t) \geq U_t \geq (1 - \epsilon)\mu_1 > \mu_i,$$

with the last inequality following from the previous assertion, giving a contradiction. \square

Claim 1: Next, we show that on \mathcal{E} , $M_\epsilon \subset \mathcal{A}(t) \cup G(t)$ for all $t \in \mathbb{N}$.

In particular this implies that if $\mathcal{A} \subset G$, then $M_\epsilon \subset G$. Combining this with the previous claim gives $G \subset M_\epsilon \subset G$, hence $G = M_\epsilon$. On this condition, EAST terminates and returns the set $\mathcal{A} \cup G = G$. Note that by definition, $M_\epsilon \subset M_{(\epsilon+\gamma)}$ for all $\gamma \geq 0$. Therefore EAST terminates correctly on this condition.

Proof. Suppose for contradiction that there exists $i \in M_\epsilon$ such that $i \notin \mathcal{A}(t) \cup G(t)$. This occurs only if i is eliminated in line 10. Hence, there exists a $t' \leq t$ such that $\hat{\mu}_i(t') + C_{\delta/n}(t') < L_{t'}$. Therefore, on the event \mathcal{E} ,

$$(1 - \epsilon)\mu_1 \stackrel{\mathcal{E}}{\geq} L_{t'} = (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \hat{\mu}_j(t') - C_{\delta/n}(t') \right) > \hat{\mu}_i(t') + C_{\delta/n}(t') \stackrel{\mathcal{E}}{\geq} \mu_i$$

which contradicts $i \in M_\epsilon$. \square

Claim 2: Finally, we show that on \mathcal{E} , if $U_t - L_t \leq \frac{\gamma}{2-\epsilon} L_t$, then $\mathcal{A} \cup G \subset M_{(\epsilon+\gamma)}$.

Combining with Claim 1 that $M_\epsilon \subset \mathcal{A} \cup G$, if EAST terminates on this condition, it does so correctly and returns all arms in M_ϵ and none in $M_{(\epsilon+\gamma)}^c$.

Proof. By Claim 0, $G \subset M_\epsilon \subset M_{\epsilon+\gamma}$. Hence, $G \cap M_{(\epsilon+\gamma)}^c = \emptyset$. Therefore, we wish to show that $\mathcal{A} \cap M_{(\epsilon+\gamma)}^c = \emptyset$ which implies that $G \cap \mathcal{A} \subset M_{\epsilon+\gamma}$. Assume $U_t - L_t < \frac{\gamma}{2-\epsilon} L_t$. Recall that

$$U_t = (1 - \epsilon) \left(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) + C_{\delta/n}(t) \right)$$

and

$$L_t = (1 - \epsilon) \left(\max_{i \in \mathcal{A}} \hat{\mu}_i(t) - C_{\delta/n}(t) \right)$$

All arms in $\mathcal{A}(t)$ have received exactly t samples. Hence, $U_t - L_t = 2(1 - \epsilon)C_{\delta/n}(t)$. On \mathcal{E} , $L_t \leq (1 - \epsilon)\mu_1$. This implies that

$$2(1 - \epsilon)C_{\delta/n}(t) < \frac{\gamma}{2 - \epsilon} L_t \leq \frac{1 - \epsilon}{2 - \epsilon} \gamma \mu_1,$$

and in particular,

$$2C_{\delta/n}(t) < \frac{\gamma \mu_1}{2 - \epsilon}.$$

Therefore, we wish to show that when the above is true, then for any $i \in M_{\epsilon+\gamma}^c$, $L_t - (\hat{\mu}_i(t) + C_{\delta/n}(t)) > 0$, implying that $i \notin \mathcal{A}$.

$$\begin{aligned} L_t - (\hat{\mu}_i(t) + C_{\delta/n}(t)) &= (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t) \right) - (\hat{\mu}_i(t) + C_{\delta/n}(t)) \\ &\geq (1 - \epsilon) \left(\max_{j \in \mathcal{A}} \mu_j - 2C_{\delta/n}(t) \right) - (\mu_i + 2C_{\delta/n}(t)) \end{aligned}$$

$$\begin{aligned}
&\stackrel{(a)}{\geq} (1 - \epsilon) (\mu_1 - 2C_{\delta/n}(t)) - ((1 - \epsilon - \gamma)\mu_1 + 2C_{\delta/n}(t)) \\
&= \gamma\mu_1 - 2(2 - \epsilon)C_{\delta/n}(t) \\
&> \gamma\mu_1 - (2 - \epsilon)\frac{\gamma\mu_1}{2 - \epsilon} \\
&= 0
\end{aligned}$$

which implies that $i \notin \mathcal{A}$. Inequality (a) follows jointly from the fact that $1 \in \mathcal{A}$ and the fact that all arms in \mathcal{A} have received t samples implies $\max_{j \in \mathcal{A}} \mu_j - 2C_{\delta/n}(t) = \mu_1 - 2C_{\delta/n}(t)$. Additionally, inequality (a) follows from $\mu_i \leq (1 - \epsilon - \gamma)\mu_1$ since $i \in M_{\epsilon+\gamma}^c$. \square

Therefore, on the event \mathcal{E} , if EAST terminates due to either condition in line 2, it returns $\mathcal{A} \cup G$ such that $M_\epsilon \subset \mathcal{A} \cup G \subset M_{(\epsilon+\gamma)}$. Since $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$, EAST terminates correctly with probability at least $1 - \delta$.

E.7.2 Step 1: Controlling the total number of samples given by EAST to arms in M_ϵ

To keep track of the number of samples that arms are given by EAST, we introduce random variables T_i and T'_i for all $i \in [n]$. When arm i has been given $\max(T_i, T'_i)$ samples it is removed from \mathcal{A} in line 11.

By Step 0, only arms in M_ϵ are added to G . Therefore, T_i is defined as

$$T_i = \min \left\{ t : \begin{array}{ll} i \in G(t+1) & \text{if } i \in M_\epsilon \\ i \notin \mathcal{A}(t+1) & \text{if } i \in M_\epsilon^c \end{array} \right\} \stackrel{\mathcal{E}}{=} \min \left\{ t : \begin{array}{ll} \hat{\mu}_i - C_{\delta/n}(t) \geq U_t & \text{if } i \in M_\epsilon \\ \hat{\mu}_i + C_{\delta/n}(t) \leq L_t & \text{if } i \in M_\epsilon^c \end{array} \right\} \quad (26)$$

Similarly, recall T'_i denotes the random variable of the number of times i is sampled before $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t)$. Hence,

$$T'_i = \min \left\{ t : \hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/n}(t) \right\} \quad (27)$$

Claim 0: For $i \in M_\epsilon$, we have that $T_i \leq h\left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{n}\right)$.

Proof. Note that $\mu_i - 2C_{\delta/n}(t) \geq (1 - \epsilon)(\mu_1 + 2C_{\delta/n}(t))$ may be rearranged as $(4 - 2\epsilon)C_{\delta/n}(t) \leq \epsilon\mu_1 - \Delta_i$, and this is true when $t > h\left(\frac{\epsilon\mu_1 - \Delta_i}{4 - 2\epsilon}, \frac{\delta}{n}\right)$. This condition implies that for all j ,

$$\begin{aligned}
\hat{\mu}_i(t) - C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\geq} \mu_i - 2C_{\delta/n}(t) \\
&\geq (1 - \epsilon)(\mu_1 + 2C_{\delta/n}(t)) \\
&\geq (1 - \epsilon)(\mu_j + 2C_{\delta/n}(t)) \\
&\stackrel{\mathcal{E}}{\geq} (1 - \epsilon)(\hat{\mu}_j(t) + C_{\delta/n}(t))
\end{aligned}$$

so in particular, $\hat{\mu}_i(t) - C_{\delta/n}(t) \geq (1 - \epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) + C_{\delta/n}(t)) = U_t$. \square

Claim 1: For $i \in M_\epsilon$, we have that $T'_i \leq h(0.25\Delta_i, \delta/n)$.

Proof. Note that $4C_{\delta/n}(t) \leq \mu_1 - \mu_i$, true when $t > h(0.25\Delta_i, \frac{\delta}{n})$, implies that

$$\begin{aligned}
\hat{\mu}_i(t) + C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(t) \\
&\leq \mu_1 - 2C_{\delta/n}(t) \\
&\stackrel{\mathcal{E}}{\leq} \hat{\mu}_1(t) - C_{\delta/n}(t).
\end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Hence, $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq \max_{j \in \mathcal{A}(t)} \hat{\mu}_j(t) - C_{\delta/n}(t)$. \square

E.7.3 Step 2: Controlling the total number of samples given by EAST to arms in M_ϵ^c

Next, we bound T_i for $i \in M_\epsilon^c$. $i \in M_\epsilon^c$ is eliminated from \mathcal{A} if it has received at least T_i samples.

Claim: $T_i \leq h\left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{n}\right)$ for $i \in M_\epsilon^c$

Proof. Note that $\mu_i + 2C_{\delta/n}(t) \leq (1-\epsilon)(\mu_1 - 2C_{\delta/n}(t))$ may be rearranged as $(4-2\epsilon)C_{\delta/n}(t) \leq \Delta_i - \epsilon\mu_1$, and this is true when $t > h\left(\frac{\Delta_i - \epsilon\mu_1}{4-2\epsilon}, \frac{\delta}{n}\right)$. This condition implies that

$$\begin{aligned}\hat{\mu}_i(t) + C_{\delta/n}(t) &\stackrel{\mathcal{E}}{\leq} \mu_i + 2C_{\delta/n}(t) \\ &\leq (1-\epsilon)(\mu_1 - 2C_{\delta/n}(t)) \\ &\stackrel{\mathcal{E}}{\leq} (1-\epsilon)(\hat{\mu}_1(t) - C_{\delta/n}(t))\end{aligned}$$

As shown in Step 0, $1 \in \mathcal{A}(t)$ for all $t \in \mathbb{N}$, and in particular $\hat{\mu}_1(t) \leq \max_{i \in \mathcal{A}(t)} \hat{\mu}_i(t)$. Therefore $\hat{\mu}_i(t) + C_{\delta/n}(t) \leq (1-\epsilon)(\max_{j \in \mathcal{A}} \hat{\mu}_j(t) - C_{\delta/n}(t)) = L_t$. \square

E.7.4 Step 3: Bounding the total number of samples drawn by EAST

With the results of Steps 1 and 2, we may bound the total sample complexity of EAST. Note that independently of the event \mathcal{E} , EAST terminates if $U_t - L_t \leq \frac{\gamma}{2-\epsilon}L_t$. Let the random variable of the maximum number of samples given to any arm before this occurs be $T_\gamma := \min\{t : U_t - L_t \leq \frac{\gamma}{2-\epsilon}L_t\}$. Additionally, EAST may terminate if $\mathcal{A} \subset G$. Let the random variable of maximum number of samples given to any arm before this occurs be $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}$. Note that due to the sampling procedure, the total number of samples drawn by EAST at termination may be written as $\sum_{t=1}^{\min(T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})} |\mathcal{A}(t)|$.

Now we bound $\sum_{t=1}^{\min(T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})} |\mathcal{A}(t)|$. Let $S_i = \min\{t : i \notin \mathcal{A}(t+1)\}$. Hence,

$$\sum_{t=1}^{\min(T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})} |\mathcal{A}(t)| = \sum_{t=1}^{\min(T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})} \sum_{i=1}^n \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \sum_{t=1}^{\min(T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})} \mathbb{1}[i \in \mathcal{A}(t)] = \sum_{i=1}^n \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, S_i\}$$

For arms $i \in M_\epsilon^c$, $S_i = T_i$ by definition. For $i \in M_\epsilon$, $S_i = \max(T_i, T'_i)$ by line 11 of the algorithm. Then

$$\begin{aligned}\sum_{i=1}^n \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, S_i\} &= \sum_{i \in M_\epsilon} \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, \max(T_i, T'_i)\} + \sum_{i \in M_\epsilon^c} \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, T_i\} \\ &= \sum_{i \in M_\epsilon} \min\{T_\gamma, \min\{T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, \max(T_i, T'_i)\}\} + \sum_{i \in M_\epsilon^c} \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, T_i\} \\ &= \sum_{i \in M_\epsilon} \min\{T_\gamma, \max\{T_i, \min(T'_i, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})\}\} + \sum_{i \in M_\epsilon^c} \min\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, T_i\}\end{aligned}$$

Next we bound T_γ .

Claim: On \mathcal{E} , $T_\gamma \leq h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)$.

Proof: $C_{\delta/n}(t) < \frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}$ is true when $t \geq h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)$. Note that

$$C_{\delta/n}(t) < \frac{\gamma\mu_1}{2(2-\epsilon+\gamma)} \iff 2C_{\delta/n}(t) < \frac{\gamma}{2-\epsilon}(\mu_1 - 2C_{\delta/n}(t)).$$

This implies that

$$\begin{aligned}U_t - L_t &= 2(1-\epsilon)C_{\delta/n}(t) \\ &< 2\frac{1-\epsilon}{2-\epsilon}\gamma(\mu_1 - 2C_{\delta/n}(t)) \\ &\leq \frac{1-\epsilon}{2-\epsilon}\gamma(\hat{\mu}_1(t) - C_{\delta/n}(t))\end{aligned}$$

$$\begin{aligned}
&\leq \frac{1-\epsilon}{2-\epsilon} \gamma \left(\max_{i \in \mathcal{A}} \hat{\mu}_i - C_{\delta/n}(t) \right) \\
&= \frac{\gamma}{2-\epsilon} L_t
\end{aligned}$$

□

Next, we may define $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon} = \min\{t : \mathcal{A}(t) \subset M_\epsilon\}$. By step 0, on the event \mathcal{E} , $\mathcal{A} \subset G$ implies that $G = M_\epsilon$. Therefore, $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}$ may be equivalently defined as $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon} = \min\{t : G(t) = M_\epsilon \text{ and } M_\epsilon^c \cap \mathcal{A} = \emptyset\}$. Recalling the definition of T_i , we see that $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon} = \max_i(T_i)$.

Recall that by steps 1 and 2, $T_i \leq h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{n}\right)$ and $T'_i \leq h\left(0.25\Delta_i, \frac{\delta}{n}\right)$. Furthermore, by monotonicity of $h(\cdot, \cdot)$, this implies that $T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon} = h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{4-2\epsilon}, \frac{\delta}{n}\right)$. Plugging this in, we see that

$$\begin{aligned}
&\sum_{i \in M_\epsilon} \min\left\{T_\gamma, \max\left\{T_i, \min(T'_i, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})\right\}\right\} + \sum_{i \in M_\epsilon^c} \min\left\{T_\gamma, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon}, T_i\right\} \\
&= \sum_{i \in M_\epsilon} \min\left\{T_\gamma, \max\left\{T_i, \min(T'_i, T_{\tilde{\alpha}_\epsilon \tilde{\beta}_\epsilon})\right\}\right\} + \sum_{i \in M_\epsilon^c} \min\{T_\gamma, T_i\} \\
&\leq \sum_{i \in M_\epsilon} \min\left\{\max\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{n}\right), \min\left[h\left(0.25\Delta_i, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{4-2\epsilon}, \frac{\delta}{n}\right)\right]\right\}, \right. \\
&\quad \left. h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\} \\
&\quad + \sum_{i \in M_\epsilon^c} \min\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{n}\right), h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\} \\
&= \sum_{i=1}^n \min\left\{\max\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{n}\right), \min\left[h\left(0.25\Delta_i, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{4-2\epsilon}, \frac{\delta}{n}\right)\right]\right\}, \right. \\
&\quad \left. h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\}
\end{aligned}$$

where the final equality holds by definition for arms in M_ϵ . Lastly, note that $\frac{1}{3(1-x)} \leq \frac{1}{2-x}$ for $x \leq 1/2$. By monotonicity of h , we may lower bound the denominators $\frac{1}{4-2\epsilon}$ and $\frac{1}{2(2-\epsilon+\gamma)}$ as $\frac{1}{6(1-\epsilon)}$ and $\frac{1}{6(1-\epsilon+\gamma)}$ respectively. Since $\epsilon \in (0, 1/2]$, we may likewise lower bound $\frac{1}{4-2\epsilon}$ as $1/4$. Plugging this in, we see that

$$\begin{aligned}
&\sum_{i=1}^n \min\left\{\max\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4-2\epsilon}, \frac{\delta}{n}\right), \min\left[h\left(0.25\Delta_i, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{4-2\epsilon}, \frac{\delta}{n}\right)\right]\right\}, \right. \\
&\quad \left. h\left(\frac{\gamma\mu_1}{2(2-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\} \\
&\leq \sum_{i=1}^n \min\left\{\max\left\{h\left(\frac{\epsilon\mu_1 - \Delta_i}{4}, \frac{\delta}{n}\right), \min\left[h\left(0.25\Delta_i, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{6(1-\epsilon)}, \frac{\delta}{n}\right)\right]\right\}, \right. \\
&\quad \left. h\left(\frac{\gamma\mu_1}{6(1-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\}
\end{aligned}$$

Next, by Lemma F.3, we may bound the minimum of $h(\cdot, \cdot)$ functions.

$$\begin{aligned}
&\sum_{i=1}^n \min\left\{\max\left\{h\left(\frac{\Delta_i - \epsilon\mu_1}{4}, \frac{\delta}{n}\right), \min\left[h\left(\frac{\Delta_i}{4}, \frac{\delta}{n}\right), h\left(\frac{\min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon)}{6(1-\epsilon)}, \frac{\delta}{n}\right)\right]\right\}, \right. \\
&\quad \left. h\left(\frac{\gamma\mu_1}{6(1-\epsilon+\gamma)}, \frac{\delta}{n}\right)\right\}
\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon \mu_i}{4}, \frac{\delta}{n} \right), \right. \right. \\
&\quad \min \left[h \left(\frac{\Delta_i}{4}, \frac{\delta}{n} \right), \max \left[h \left(\frac{\tilde{\alpha}_\epsilon}{6(1-\epsilon)}, \frac{\delta}{n} \right), h \left(\frac{\tilde{\beta}_\epsilon}{6(1-\epsilon)}, \frac{\delta}{n} \right) \right] \right] \left. \right\}, \\
&\quad \left. h \left(\frac{\gamma \mu_i}{6(1-\epsilon+\gamma)}, \frac{\delta}{n} \right) \right\} \\
&\leq \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon \mu_i}{4}, \frac{\delta}{n} \right), \right. \right. \\
&\quad \max \left[h \left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right) \right] \left. \right\}, \\
&\quad \left. h \left(\frac{\gamma \mu_i}{6(1-\epsilon+\gamma)}, \frac{\delta}{n} \right) \right\} \\
&= \sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon \mu_i}{4}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right) \right\}, \right. \\
&\quad \left. h \left(\frac{\gamma \mu_i}{6(1-\epsilon+\gamma)}, \frac{\delta}{n} \right) \right\}
\end{aligned}$$

Finally, we use Lemma F.2 to bound the function $h(\cdot, \cdot)$. Since $\delta \leq 1/2$, $\delta/n \leq 2e^{-e/2}$. Further, $|\epsilon \mu_1 - \Delta_i| \leq 6$ for all i and $\epsilon \leq 1/2$ implies that $\frac{1}{6(1-\epsilon)}|\epsilon \mu_1 - \Delta_i| \leq 2$ and $\frac{1}{6(1-\epsilon)} \min(\tilde{\alpha}_\epsilon, \tilde{\beta}_\epsilon) \leq 2$. $\Delta_i \leq 8$ for all i , gives $0.25\Delta_i \leq 2$. Lastly, $\gamma \leq 6/\mu_1$ implies that $\frac{\gamma \mu_1}{6(1-\epsilon+\gamma)} \leq 2$. Therefore,

$$\begin{aligned}
&\sum_{i=1}^n \min \left\{ \max \left\{ h \left(\frac{\Delta_i - \epsilon \mu_i}{4}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right), h \left(\frac{\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon}}{12}, \frac{\delta}{n} \right) \right\}, \right. \\
&\quad \left. h \left(\frac{\gamma \mu_i}{6(1-\epsilon+\gamma)}, \frac{\delta}{n} \right) \right\} \\
&\leq \sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{(\epsilon \mu_1 - \Delta_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta(\epsilon \mu_1 - \Delta_i)^2} \right) \right), \right. \right. \\
&\quad \frac{576}{(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\Delta_i + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon})^2} \right) \right), \\
&\quad \frac{576}{(\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon})^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\Delta_i + \frac{\tilde{\beta}_\epsilon}{1-\epsilon})^2} \right) \right) \left. \right\}, \\
&\quad \frac{144(1-\epsilon+\gamma)^2}{\gamma^2 \mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{432(1-\epsilon+\gamma)^2 n}{\delta \gamma^2 \mu_1^2} \right) \right) \left. \right\} \\
&= \sum_{i=1}^n \min \left\{ \max \left\{ \frac{64}{((1-\epsilon)\mu_1 - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{192n}{\delta((1-\epsilon)\mu_1 - \mu_i)^2} \right) \right), \right. \right. \\
&\quad \frac{576}{(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\mu_1 + \frac{\tilde{\alpha}_\epsilon}{1-\epsilon} - \mu_i)^2} \right) \right), \\
&\quad \frac{576}{(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{1728n}{\delta(\mu_1 + \frac{\tilde{\beta}_\epsilon}{1-\epsilon} - \mu_i)^2} \right) \right) \left. \right\}, \\
&\quad \frac{144(1-\epsilon+\gamma)^2}{\gamma^2 \mu_1^2} \log \left(\frac{2n}{\delta} \log_2 \left(\frac{432(1-\epsilon+\gamma)^2 n}{\delta \gamma^2 \mu_1^2} \right) \right) \left. \right\}.
\end{aligned}$$

□

F Technical Lemmas

Lemma F.1. *If $a > 1$, $b > e$, and $t > \max(a \log(2b \log(ab)), e)$, then $\frac{a \log(b \log(t))}{t} \leq 1$*

Proof. **Step 1:** Plug in $t = a \log(2b \log(ab))$ to the expression $\frac{a \log(b \log(t))}{t}$.

$$\frac{a \log(b \log(a \log(2b \log(ab))))}{a \log(2b \log(ab))} = \frac{\log(b \log(a \log(2b \log(ab))))}{\log(2b \log(ab))}$$

Since $\log(\cdot)$ increases monotonically, the above is less than 1 if $b \log(a \log(2b \log(ab))) \leq 2b \log(ab)$.

$$\begin{aligned} b \log(a \log(2b \log(ab))) &\leq 2b \log(ab) \\ \stackrel{(b>0)}{\iff} \log(a \log(2b \log(ab))) &\leq 2 \log(ab) \\ \iff a \log(2b \log(ab)) &\leq (ab)^2 \\ \iff \log(2b \log(ab)) &\leq ab^2 \\ \iff 2b \log(ab) &\leq e^{ab^2} \end{aligned}$$

which is true if $a, b > 1$.

Step 2: Next, for $t > a \log(2b \log(ab))$, we wish to show that the inequality $\frac{a \log(b \log(t))}{t} \leq 1$ still holds. To do so, it suffices to show that $f(t) = \frac{a \log(b \log(t))}{t}$ is decreasing for $t > a \log(2b \log(ab))$. To see this, take the derivative.

$$f'(t) = \frac{a}{t^2 \log(t)} - \frac{a \log(b \log(t))}{t^2} = \frac{a}{t^2} \left(\frac{1}{\log(t)} - \log(b \log(t)) \right)$$

This is negative when $\frac{1}{\log(t)} < \log(b \log(t))$. Let $u = b \log(t)$. The previous is equivalent to the condition $b < u \log(u)$. For $t > e$, $u > b$ and $b > e$. Hence $b < u \log(u)$ completing the proof. □

Lemma F.2. *For $\delta < 2e^{-e/2}$, $\Delta \leq 2$,*

$$t \geq \frac{4}{\Delta^2} \log \left(\frac{2}{\delta} \log_2 \left(\frac{12}{\delta \Delta^2} \right) \right) \implies C_\delta(t) = \sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}} \leq \Delta.$$

Proof.

$$\sqrt{\frac{4 \log(\log_2(2t)/\delta)}{t}} \leq \Delta \iff \frac{4 \frac{8}{\Delta^2} \log \left(\frac{1}{\delta \log(2)} \log(2t) \right)}{t} \leq 1.$$

If $\Delta \leq 2$, then $8/\Delta^2 \geq 2 > 1$. Similarly, if $\delta < 2e^{-e/2} < \frac{1}{e \log(2)}$, then $\frac{1}{\delta \log(2)} > e$. Hence, by Lemma F.1, setting $a = \frac{8}{\Delta^2}$ and $b = \frac{1}{\delta \log(2)}$, the above is true if

$$2t \geq \max \left(\frac{8}{\Delta^2} \log \left(\frac{2}{\delta \log(2)} \log \left(\frac{8}{\delta \Delta^2 \log(2)} \right) \right), e \right).$$

Trivially, $\delta \log(2) < 2$. Hence, $\delta < 2e^{-e/2}$ and $\Delta \leq 2$ implies

$$\frac{8}{\Delta^2} \log \left(\frac{2}{\delta \log(2)} \log \left(\frac{8}{\delta \Delta^2 \log(2)} \right) \right) \geq 2 \log \left(\frac{2}{\delta} \log_2 \left(\frac{2}{\delta \log(2)} \right) \right) \geq 2 \log(2/\delta) > e.$$

Therefore, we may simplify the maximum as

$$t \geq \frac{4}{\Delta^2} \log \left(\frac{2}{\delta} \log_2 \left(\frac{12}{\delta \Delta^2} \right) \right) \geq \frac{4}{\Delta^2} \log \left(\frac{2}{\delta} \log_2 \left(\frac{8}{\delta \Delta^2 \log(2)} \right) \right)$$

which implies the desired result. □

Lemma F.3. For any function $h(\cdot, \cdot) : \mathbb{R}^+ \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ that decreases monotonically in its first argument, we have that for any $a, b, c, \delta \in \mathbb{R}^+$

$$\min(h(a, \delta), h(b, \delta)) \leq h\left(\frac{a+b}{2}, \delta\right)$$

and

$$\min\{h(a, \delta), \max[h(b, \delta), h(c, \delta)]\} \leq \max\left\{h\left(\frac{a+b}{2}, \delta\right), h\left(\frac{a+c}{2}, \delta\right)\right\}.$$

Proof. First, we bound the expression $\min(h(a, \delta), h(b, \delta))$.

$$\min(h(a, \delta), h(b, \delta)) = h(\max(a, b), \delta) \leq h((a+b)/2, \delta)$$

Next, we bound, expressions of the form $\min\{h(a, \delta), \max[h(b, \delta), h(c, \delta)]\}$ using the above inequality.

$$\begin{aligned} \min\{h(a, \delta), \max[h(b, \delta), h(c, \delta)]\} &= \max\{\min[h(a, \delta), h(b, \delta)], \min[h(a, \delta), h(c, \delta)]\} \\ &\leq \max\{h((a+b)/2, \delta), h((a+c)/2, \delta)\}. \end{aligned}$$

□