
Best Arm Identification with Safety Constraints

Zhenlin Wang

National University of Singapore
wang_zhenlin@u.nus.edu

Andrew Wagenmaker

University of Washington
ajwagen@cs.washington.edu

Kevin Jamieson

University of Washington
jamieson@cs.washington.edu

Abstract

The best arm identification problem in the multi-armed bandit setting is an excellent model of many real-world decision-making problems, yet it fails to capture the fact that in the real-world, safety constraints often must be met while learning. In this work we study the question of best-arm identification in safety-critical settings, where the goal of the agent is to find the best *safe* option out of many, while exploring in a way that guarantees certain, initially unknown safety constraints are met. We first analyze this problem in the setting where the reward and safety constraint takes a linear structure, and show nearly matching upper and lower bounds. We then analyze a much more general version of the problem where we only assume the reward and safety constraint can be modeled by monotonic functions, and propose an algorithm in this setting which is guaranteed to learn safely. We conclude with experimental results demonstrating the effectiveness of our approaches in scenarios such as safely identifying the best drug out of many in order to treat an illness.

1 INTRODUCTION

Consider a dosing trial where a scientist is trying to determine, out of d different drugs, which is most effective at treating a particular illness. For each drug, the scientist can run an experiment where they administer a particular dose of a drug to a patient and observe the effectiveness. After repeating this process on multiple patients for each drug, the scientist must

then give a recommendation as to which drug is most effective. Critically, in the experimental process, the safety of the participants must be ensured. While we may know a low, baseline safe dosage for each drug, the drug may only become effective at higher doses, some of which may be above a certain dose threshold in which the treatment begins to become unsafe, and could cause negative side effects. The goal of the experimenter is not only to determine which drug is most effective, but to also guarantee that only safe dosage levels are applied to the patients in the experiments.

We can cast this problem as a *best-arm identification problem with safety constraints*. In the classical multi-armed bandit (MAB) setting, at each round, the agent may choose from one of d options and obtains a noisy realization of the option’s *reward*. The goal of best-arm identification (BAI) is to determine which of the d options has the largest mean reward using the minimum number of samples. This problem has been extensively studied and is an effective model of many real-world problems. However, the standard MAB setting is unable to incorporate either different “doses” or safety constraints.

In this work we set out to address this question and answer how an agent ought to sample in a multi-armed bandit model where the goal is to *safely* identify the best option. At each timestep, the learner must choose both the “dosing level” for the arm pulled, and ensure that this dosing level is “safe”. Critically, we assume that the range of safe dosing levels is unknown, and this must be learned as well.

We first study this problem in the setting where the relationship between the reward obtained—the effectiveness—and the dosing level, as well as the safety of a dosing level, are *linear*. We propose a near-optimal action elimination-style algorithm, SAFEBAI-LINEAR, which successively refines its estimates of the safe dosing levels to ensure safety throughout execution, and gradually increases the dose applied to each option until it is able to determine the best option. We next consider the much more general setting where we only assume that the effectiveness and the safety are

monotonic functions of the dosing level—a more realistic assumption for modeling a drug response than the linear assumption. We propose an algorithm in this setting, SAFEBAI-MONOTONIC, which we show explores efficiently—guaranteeing only safe dose levels are applied while determining the optimal drug.

2 RELATED WORK

Best-Arm Identification in Multi-Armed Bandits. The BAI problem in multi-armed bandits has been extensively studied for decades, beginning with the works of Bechhofer (1958) and Paulson (1964). A variety of different approaches have been proposed, including action-elimination style algorithms (Even-Dar et al., 2002; Karnin et al., 2013) and upper confidence bound-style algorithms (Bubeck et al., 2009; Kalyanakrishnan et al., 2012; Jamieson et al., 2014). Furthermore, lower bounds have been proposed (Mannor & Tsitsiklis, 2004; Kaufmann et al., 2016) which have been shown to be near-tight. The linear response setting we consider is somewhat related to the BAI problem in the linear bandit setting (Soare et al., 2014) and the monotonic function setting we study is related to the continuous-armed bandit setting, also known as zeroth-order or derivative-free optimization (Flaxman et al., 2005; Jamieson et al., 2012; Valko et al., 2013; Rios & Sahinidis, 2013).

Perhaps the most comparable work to ours is Sui et al. (2015, 2018), which considers the problem of finding the value $x \in D$ such that $f(x)$ is maximized, for some f and D , where at each time step they observe $f(x_t) + \eta_t$, and must guarantee that $g(x_t) \geq h$ during exploration. In essence, this is a constrained BAI problem, but these works do not provide a tight upper bound (indeed, their algorithms are unverifiable in style—they provide no stopping condition guaranteeing the optimal value has been found), or an information-theoretic lower bound. Furthermore, their algorithms use a potentially wasteful exploration strategy which could overexplore suboptimal arms. Our work improves on all these facets—our algorithm is verifiable, providing a stopping criteria with an optimality guarantee, utilizes a much more efficient exploration scheme, and we prove an upper bound with nearly matching lower bound. To our knowledge, this is the only existing work which studies a problem akin to BAI where safety constraints are present.

Regret Minimization in Bandits with Safety Constraints. A related line of work is that of regret minimization in bandits with safety constraints. The majority of work in this setting considers the more general linear bandits problem where now the agent at each timestep must choose a vector $x_t \in \mathcal{X} \subseteq \mathbb{R}^d$ and

observes $y_t = \theta^\top x_t + \eta_t$. A variety of different formulations of the safety constraint have been proposed, but in general they require that $\mu^\top x_t \leq \gamma$ either almost surely or in expectation. Various algorithmic approaches have been applied, such as Thompson Sampling (Moradipari et al., 2020, 2021), as well as optimistic UCB-style approaches (Kazerouni et al., 2016; Amani et al., 2019; Pacchiano et al., 2021). The key difference between our problem and the existing work in this setting is that we are interested in best-arm identification, while existing works tend to focus on regret minimization, where the goal is to achieve large online reward, not simply to identify the best arm. In addition, we seek to obtain instance-dependent (gap-dependent) results while existing works only target minimax (worst-case) bounds.

Dose-Finding and Thresholding Bandits.

Dose-finding is a long-standing problem in the biomedical sciences where the goal is to determine the optimal dose of a drug to give a patient (Thall & Russell, 1998; Thall & Cook, 2004; Rogatko et al., 2005; Musuamba et al., 2017; Riviere et al., 2018). Much work has been done on developing statistically justified procedures to address this. Recently, the bandits community has begun to approach this problem from the perspective of structured multi-armed bandits—a setting which has become known as the *thresholding bandit* problem. Chen et al. (2014) consider a general version of the thresholding bandit problem, while several follow-up works (Locatelli et al., 2016; Garivier et al., 2017; Cheshire et al., 2020; Aziz et al., 2021) consider the particular application of this setting to the problem of identifying safe dosing levels in Phase 1 clinical trials. In contrast to our setting, this setting does not consider safety constraints—at any time, any dose level may be tested with no penalty—and, in addition, does not consider the identification of the best drug out of many, rather it can be seen as simply determining the largest safe dose level for a particular drug. In a sense, our setup then extends this dose-finding problem to its multi-dimensional analogue.

3 PRELIMINARIES

Notation. Throughout, $[n] = \{1, 2, 3, \dots, n\}$. We will use $\tilde{O}(\cdot)$ to hide logarithmic terms and absolute constants. We let $\mathbb{R}_+^d = \{x \in \mathbb{R}^d : x_i \geq 0, \forall i \in [d]\}$.

Problem Setting: Linear. In the linear response case, we are interested in the setting where at every time step t , the learner chooses a coordinate $i_t \in [d]$ and a value $a_t \in [a_{0,i_t}, M_{i_t}]$, and observes

$$y_t = a_t \theta_{i_t} + \eta_t, \quad z_t = a_t \mu_{i_t} + \eta_t$$

for η_t and w_t both σ^2 -sub-Gaussian and mean 0. We assume that $\theta \in \mathbb{R}_+^d$ and $\mu \in \mathbb{R}_+^d$ are initially unknown and, furthermore, that the learner must always choose, with high probability, a_t and i_t satisfying $a_t \mu_{i_t} \leq \gamma$, for some known γ . To make this tractable, we assume that the learner is given $a_{0,i} > 0$ and $M_i > 0$, $a_{0,i} \leq M_i$, and that $a_{0,i} \mu_i \leq \gamma$ for all i . We will refer to $i \in [d]$ as a *coordinate*, and $a \in [a_{0,i}, M_i]$ as an *arm*.

The goal of the learner is to identify the optimal coordinate; that is, to find $i^* \in [d]$ defined as:

$$i^* := \arg \max_i \max_{a \in [a_{0,i}, M_i]} a \theta_i \quad \text{s.t.} \quad a \mu_i \leq \gamma.$$

In other words, we want to find the coordinate that has the largest safe value. Here we will only be concerned with identifying this coordinate, not determining what that largest safe value is (once the optimal coordinate is identified, it can be repeatedly played until the largest safe value is identified). For simplicity we assume that there is a unique optimal coordinate.

We will let \mathcal{S}_i denote the set of all safe values for coordinate i : $\mathcal{S}_i = \{a \in [a_{0,i}, M_i] : a \mu_i \leq \gamma\}$. We will define the gap for coordinate i as:

$$\Delta_i := \min \left\{ \frac{\gamma \theta_{i^*}}{\mu_{i^*}}, \theta_{i^*} M_{i^*} \right\} - \min \left\{ \frac{\gamma \theta_i}{\mu_i}, \theta_i M_i \right\}.$$

Note that $\gamma \theta_i / \mu_i = \max_{a \in \mathcal{S}_i} a \theta_i$ is the largest safe effect for coordinate i , in the case when $M_i \notin \mathcal{S}_i$, and $\theta_i M_i$ is the largest safe effect when $M_i \in \mathcal{S}_i$. Our definition of the gap then corresponds to the difference in maximum safe value for the optimal coordinate and the maximum safe value for coordinate i .

Problem Setting: Monotonic. While the linear setting is useful for precisely quantifying the complexity of learning, often in real-world settings the actual structure is nonlinear. To address this, we introduce the following more general version of the problem, where now we assume our observations take the form

$$y_t = f_{i_t}(a_t) + \eta_t, \quad z_t = g_{i_t}(a_t) + w_t$$

for some functions $\{(f_i, g_i)\}_{i \in [d]}$. We assume that f_i and g_i are defined over all \mathbb{R} and, to simplify the analysis, allow the learner to play any values $a \in \mathbb{R}$ (provided they are safe). To guarantee safety while learning, we must make additional assumptions on the structure of f_i and g_i . Henceforth, we will assume that f_i and g_i satisfy the following.

Assumption 1. For all $i \in [d]$, $g_i(\cdot)$ is 1-Lipschitz and strictly monotonically increasing. Furthermore, for all $i \in [d]$ and $a \in \mathbb{R}$, $f_i(a) \in [0, 1]$ and $f_i(\cdot)$ is nondecreasing.

Note that this assumption implies that g_i is invertible. We assume that f_i and g_i are initially unknown to the learner, but that the learner is told that g_i is 1-Lipschitz. In this setting, our safety constraint is:

$$g_{i_t}(a_t) \leq \gamma$$

and we assume that for each i , the learner is given an initial value $a_{0,i}$ such that $g_i(a_{0,i}) \leq \gamma$. To guarantee that we find the best coordinate, we need a slightly stronger query model which allows arms to be queried that are marginally above this threshold. To this end, we introduce a value ϵ_{safe} , and allow our learner to query points a_t which satisfy:

$$g_{i_t}(a_t) \leq \gamma + \epsilon_{\text{safe}}.$$

However, we are still interested in finding the best value satisfying $g_i(a_t) \leq \gamma$, and define the best coordinate as (note that $f_i(g_i^{-1}(\gamma))$ is the maximum achievable safe value for coordinate i):

$$i^* := \arg \max_i f_i(g_i^{-1}(\gamma))$$

Similar to the linear case, our goal is only to identify the best coordinate, not the value of $f_{i^*}(g_{i^*}^{-1}(\gamma))$. We can define a notion of the *gap* as the difference between the maximum safe reward achievable by the best coordinate and the maximum safe reward achievable by coordinate i :

$$\Delta_i := f_{i^*}(g_{i^*}^{-1}(\gamma)) - f_i(g_i^{-1}(\gamma)).$$

Algorithm Classes. Formally, we will define “safe” algorithms in the following way.

Definition 1 (δ -Safe Algorithm). *We say that an algorithm \mathcal{A} is δ -safe if, with probability $1 - \delta$, \mathcal{A} only pulls arms a_t satisfying $g_{i_t}(a_t) \leq \gamma$.*

In the linear case, the above condition $g_{i_t}(a_t) \leq \gamma$ is equivalent to $a_t \mu_{i_t} \leq \gamma$.

Definition 2 (δ -PAC Algorithm). *We say that an algorithm \mathcal{A} is δ -PAC if it outputs a coordinate \hat{i} such that $\mathbb{P}[\hat{i} = i^*] \geq 1 - \delta$.*

Definition 3 (δ -PAC Safe Algorithm). *We say that an algorithm \mathcal{A} is a δ -PAC Safe algorithm if \mathcal{A} is both δ -Safe and δ -PAC.*

Our goal will be to obtain an algorithm that is δ -PAC Safe with the minimum possible sample complexity.

4 LOWER BOUND

We first present a lower bound on the complexity of safe BAI. For simplicity, we assume here that $\sigma^2 =$

1 and that $\eta_t \sim \mathcal{N}(0, \sigma^2)$, $w_t \sim \mathcal{N}(0, \sigma^2)$. We will consider a slightly different class of algorithms to prove our lower bound. Rather than learning which arms are safe, we give the algorithm a set of arms it can pull throughout execution.

Definition 4 ($(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC Algorithm). *We say an algorithm is $(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC if it is δ -PAC and with probability 1 only chooses values $a_t \in \mathcal{X}_i$.*

A $(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC algorithm is then given a set of values, \mathcal{X}_i , by an oracle that it can pull for each coordinate, and is only allowed to pull these arms throughout execution (critically, it is not told if these values are safe—it is only told that it is allowed to pull them). In other words, rather than requiring it to *learn* which values it is able to safely pull, it is simply told which values it can pull. In particular, we will set $\mathcal{X}_i = \mathcal{S}_i$, allowing the $(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC learner to pull any safe arm without verifying it is safe. In this case a $(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC learner is a strictly more powerful learner than a δ -PAC Safe learner—it is able to query *any* safe arm at *any* time, while a δ -PAC Safe learner may only query arms it has verified are safe. Thus, the task of learning the optimal safe coordinate for a $(\delta, \{\mathcal{S}_i\}_{i \in [d]})$ -PAC learner is easier than for a δ -PAC Safe learner.

Theorem 1. *Fix an instance $\theta \in \mathbb{R}_+^d$ and $\mu \in \mathbb{R}_+^d$, $\gamma > 0$, and $M_i = \infty$ for all i , and let $\mathcal{S}_i(\mu, \gamma)$ denote the safe values for coordinate i on this instance. Let τ denote the stopping time for any $(\delta, \{\mathcal{S}_i(\mu, \gamma)\}_{i \in [d]})$ -PAC algorithm. Then on this instance we will have that*

$$\mathbb{E}_{\theta, \mu}[\tau] \geq \frac{2}{3} \log \frac{1}{2.4\delta} \cdot \sum_{i \neq i^*} \frac{1 + \theta_{i^*}^2 / \mu_{i^*}^2 + \theta_i^2 / \mu_i^2}{\Delta_i^2}.$$

Theorem 1 states that the lower bound in the safe BAI setting scales as the familiar “sum over inverse gaps squared” lower bound of the standard multi-armed bandit setting. The key difference here is the $\theta_{i^*}^2 / \mu_{i^*}^2 + \theta_i^2 / \mu_i^2$ term. For each i , we can break up the cost associated with coordinate i into terms $\frac{1}{\Delta_i^2}$ and $\frac{\theta_{i^*}^2 / \mu_{i^*}^2 + \theta_i^2 / \mu_i^2}{\Delta_i^2}$. The first term is due to showing that coordinate i is suboptimal and is present in the standard multi-armed bandit lower bound. The second term, in contrast, is not present in the standard multi-armed bandit lower bound and arises in this setting as the cost of learning where the safety threshold is for coordinate i . A proof of Theorem 1 is given in Appendix B.

5 LINEAR RESPONSE

Given this lower bound, we next propose an algorithm, SAFEBAI-LINEAR, in the linear response case.

Algorithm 1 Safe Best-Arm Identification for Linear Functions (SAFEBAI-LINEAR)

- 1: **input:** Confidence δ , noise variance σ^2 , safety tolerance γ , value bounds $\{(a_{0,i}, M_i)\}_{i \in [d]}$
 - 2: **initialize** $\hat{a}_s^0(i) \leftarrow a_{0,i}$, $\hat{a}_u^0(i) \leftarrow M_i$, $i \in [d]$, $\mathcal{X}_0 \leftarrow [d]$, $\ell \leftarrow 1$
 - 3: **while** $|\mathcal{X}_{\ell-1}| > 1$ **do**
 - 4: $\epsilon_\ell \leftarrow 2^{-\ell}$
 - 5: **for** $i \in \mathcal{X}_{\ell-1}$ **do**
 - 6: $N_\ell \leftarrow \lceil 2\sigma^2 \log \frac{8d\ell^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$
 - 7: Pull $\hat{a}_s^{\ell-1}(i)$ N_ℓ times, observe:

$$y_t = \hat{a}_s^{\ell-1}(i)\theta_i + \eta_t, z_t = \hat{a}_s^{\ell-1}(i)\mu_i + w_t$$
 - 8: $\hat{\theta}_{i,\ell} \leftarrow \frac{1}{\hat{a}_s^{\ell-1}(i)N_\ell} \sum_{t=1}^{N_\ell} y_t$
 - 9: $\hat{\mu}_{i,\ell} \leftarrow \frac{1}{\hat{a}_s^{\ell-1}(i)N_\ell} \sum_{t=1}^{N_\ell} z_t$
 - 10: **if** $\hat{a}_s^{\ell-1}(i) < M_i$ **then**
 - 11: $\hat{a}_s^\ell(i) \leftarrow \min\{\max\{\frac{\gamma}{\hat{\mu}_{i,\ell} + \epsilon_\ell / \hat{a}_s^{\ell-1}(i)}, a_{0,i}\}, M_i\}$
 - 12: $\hat{a}_u^\ell(i) \leftarrow \min\{\frac{\gamma}{\hat{\mu}_{i,\ell} - \epsilon_\ell / \hat{a}_s^{\ell-1}(i)}, M_i\}$
 - 13: **else**
 - 14: $\hat{a}_s^\ell(i) \leftarrow M_i$, $\hat{a}_u^\ell(i) \leftarrow M_i$
 - 15: $\mathcal{X}_\ell \leftarrow \{i \in \mathcal{X}_{\ell-1} : \hat{a}_u^\ell(i)(\hat{\theta}_{i,\ell} + \frac{\epsilon_\ell}{\hat{a}_s^{\ell-1}(i)}) \geq \max_j \hat{a}_s^\ell(j)(\hat{\theta}_{j,\ell} - \frac{\epsilon_\ell}{\hat{a}_s^{\ell-1}(j)})\}$
 - 16: $\ell \leftarrow \ell + 1$
 - 17: **return** $\mathcal{X}_{\ell-1}$
-

SAFEBAI-LINEAR operates by maintaining estimates of the maximum safe value for each active coordinate i at epoch ℓ , $\hat{a}_s^\ell(i)$. At every epoch it pulls $\hat{a}_s^\ell(i)$ repeatedly for each active arm, using the observations to refine its estimate of the maximum safe value as well as the minimum unsafe value, $\hat{a}_u^\ell(i)$, and to form an estimate of θ_i , $\hat{\theta}_{i,\ell}$. It then uses these estimates to construct a lower bound on the maximum safe value of coordinate i , $\hat{a}_s^\ell(i)(\hat{\theta}_{i,\ell} - \frac{\epsilon_\ell}{\hat{a}_s^{\ell-1}(i)})$, and an upper bound on the maximum safe value, $\hat{a}_u^\ell(i)(\hat{\theta}_{i,\ell} + \frac{\epsilon_\ell}{\hat{a}_s^{\ell-1}(i)})$, and eliminates coordinates that are provably suboptimal. By only pulling the provably safe values, $\hat{a}_s^\ell(i)$, and halving the tolerance at every epoch, SAFEBAI-LINEAR is able to safely refine its estimates of the problem parameters. Furthermore, by playing the largest verifiably safe value, it is able to effectively reduce the signal-to-noise ratio, guaranteeing efficient, safe convergence to the optimal coordinate.

5.1 Sample Complexity

Towards presenting the sample complexity of Algorithm 1, we define the function:

$$\xi_a(x) := 2^a \sqrt{\log_2 \max\{x, 2\}}.$$

For $a > 0$, $\xi_a(x)$ grows sub-polynomially yet super-logarithmically in x . The following result gives a quantification of its growth.

Proposition 1. *For $a > 0$, we can bound $\xi_a(x) \leq \min_z \max\{x, 2\}^z + 2^{a^2/z}$ and $\xi_a(x) \geq 1$.*

We are now ready to give our sample complexity result.

Theorem 2. *Recall that \mathcal{S}_i denotes the safe values for coordinate i , and M_i denotes the maximum playable value for coordinate i . Let $C_\gamma := \xi_{\sqrt{32}}(2\gamma)$ and define the following:*

1. **Case 1** ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$):

$$N_{1,i} := \frac{1 + \theta_{i^*}^2/\mu_{i^*}^2 + \theta_i^2/\mu_i^2}{\Delta_i^2} + C_\gamma \max \left\{ \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \right\}.$$

2. **Case 2** ($M_i \in \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$):

$$N_{2,i} := \frac{1 + M_i^2\theta_i^2/\gamma^2}{\Delta_i^2} + \frac{M_i^2\mu_i^2/\gamma^2}{(\gamma - M_i\mu_i)^2} + C_\gamma \max \left\{ \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right), \xi_4\left(\frac{M_i\mu_i}{\gamma - M_i\mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \right\}.$$

3. **Case 3** ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$):

$$N_{3,i} := \frac{1 + M_{i^*}^2\theta_{i^*}^2/\gamma^2}{\Delta_i^2} + \frac{M_{i^*}^2\mu_{i^*}^2/\gamma^2}{(\gamma - M_{i^*}\mu_{i^*})^2} + C_\gamma \max \left\{ \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right), \xi_4\left(\frac{M_{i^*}\mu_{i^*}}{\gamma - M_{i^*}\mu_{i^*}}\right), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \right\}.$$

4. **Case 4** ($M_i \in \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$):

$$N_{4,i} := \frac{1}{\Delta_i^2} + \frac{M_{i^*}^2\mu_{i^*}^2/\gamma^2}{(\gamma - M_{i^*}\mu_{i^*})^2} + \frac{M_i^2\mu_i^2/\gamma^2}{(\gamma - M_i\mu_i)^2} + C_\gamma \max_{j \in \{i, i^*\}} \xi_4\left(\frac{1}{a_{0,j}\mu_j}\right) \xi_4\left(\frac{M_j\mu_j}{\gamma - M_j\mu_j}\right).$$

Let $c(i) \in \{1, 2, 3, 4\}$ denote the case coordinate i falls in. Then, with probability at least $1 - \delta$, Algorithm 1 will output i^* , only pull safe arms, and terminate after collecting at most

$$\tilde{\mathcal{O}}\left(\log \frac{d}{\delta} \cdot \sum_{i \neq i^*} N_{c(i),i} + \log \frac{d}{\delta} \cdot \frac{d}{\gamma^8}\right)$$

samples.

Theorem 2 is a slightly simplified version of the more general Theorem 4 which we present in the appendix. While we only state the result for *best*-coordinate identification, SAFEBAI-LINEAR could also be applied to obtain an ϵ -good coordinate identification-style guarantee. When M_i is unsafe for all i (for example, if the maximum possible value M_i is unbounded and $\mu_i > 0$), we obtain the following.

Corollary 1. *Assume that M_i is unsafe for all i . Then, with probability at least $1 - \delta$, Algorithm 1 will output i^* , only pull safe arms, and terminate after collecting at most*

$$\tilde{\mathcal{O}}\left(\log \frac{d}{\delta} \cdot \sum_{i \neq i^*} \frac{1 + \theta_{i^*}^2/\mu_{i^*}^2 + \theta_i^2/\mu_i^2}{\Delta_i^2} + \log \frac{d}{\delta} \cdot \sum_{i \neq i^*} \max \left\{ \frac{1}{\gamma^8}, C_\gamma \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right), C_\gamma \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \right\}\right)$$

samples.

We note that the sample complexity stated in Corollary 1 exactly matches the lower bound given in Theorem 1, up to constants, log factors, and the lower order term scaling inversely in $a_{0,i}$ and γ , despite the fact that Theorem 1 was proved for a more powerful class of learners. It follows that SAFEBAI-LINEAR achieves near-optimal sample complexity for the problem while guaranteeing only safe arms are pulled.

5.2 Proof Sketch

At every epoch ℓ , for all coordinates $i \in \mathcal{X}_{\ell-1}$ we have not yet shown are suboptimal, we collect $N_\ell = \lceil 2\sigma^2 \log \frac{8d\ell^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$ samples at value $\hat{a}_s^{\ell-1}(i)$ on Line 7. Standard concentration then gives that, with high probability,

$$\hat{a}_s^{\ell-1}(i)|\mu_i - \hat{\mu}_{i,\ell}| \leq \epsilon_\ell, \quad \hat{a}_s^{\ell-1}(i)|\theta_i - \hat{\theta}_{i,\ell}| \leq \epsilon_\ell.$$

Note then that

$$\frac{\gamma}{\hat{\mu}_{i,\ell} + \epsilon_\ell/\hat{a}_s^{\ell-1}(i)} \leq \frac{\gamma}{\mu_i - \epsilon_\ell/\hat{a}_s^{\ell-1}(i) + \epsilon_\ell/\hat{a}_s^{\ell-1}(i)} = \frac{\gamma}{\mu_i}$$

so it follows that $\gamma/(\hat{\mu}_{i,\ell} + \epsilon_\ell/\hat{a}_s^{\ell-1}(i))$ is safe. As this is what we set $\hat{a}_s^\ell(i)$ to on Line 11 (ignoring the range $[a_{0,i}, M_i]$), it follows that $\hat{a}_s^\ell(i)$ is always safe, so we only ever play safe values. A similar argument shows that $\hat{a}_u^\ell(i)$ is always unsafe. Given the accuracy of our estimate of θ_i , it follows that $\hat{a}_u^\ell(i)(\hat{\theta}_{i,\ell} + \epsilon_\ell/\hat{a}_s^{\ell-1}(i))$ is an upper bound on the maximum safe value of coordinate i , while $\hat{a}_s^\ell(i)(\hat{\theta}_{i,\ell} - \epsilon_\ell/\hat{a}_s^{\ell-1}(i))$ is a lower bound. As we only eliminate coordinates that have upper bounds less than the largest lower bound, it follows that we only eliminate suboptimal coordinates on Line 15.

To bound the sample complexity, we first obtain a deterministic lower bound on $\hat{a}_s^\ell(i)$ by solving a recursion (the dependence on $\xi_a(x)$ results from solving this recursion), and then, using this, obtain deterministic upper and lower bounds on $\hat{a}_u^\ell(i)(\hat{\theta}_{i,\ell} + \epsilon_\ell/\hat{a}_s^{\ell-1}(i))$ and $\hat{a}_s^\ell(i)(\hat{\theta}_{i,\ell} - \epsilon_\ell/\hat{a}_s^{\ell-1}(i))$. Noting that these are separated at most by the true gap, we show that if $\hat{a}_s^\ell(i)$ is close enough to the maximum safe value, in order to eliminate coordinate i it suffices to collect roughly $1/\Delta_i^2$ samples. The full proof is given in Appendix A.

6 MONOTONIC RESPONSE

We turn now to our second setting, where we assume that our response is a monotonic function. In order to show that a coordinate is suboptimal, you must obtain an upper bound on the maximum safe function value. Without assuming more than monotonicity and smoothness, we cannot guarantee such an upper bound if we only allow the learner to sample points a such that $g_i(a) \leq \gamma$ —points that are safe. To address this, we allow the learner to sample points a such that $g_i(a) \leq \gamma + \epsilon_{\text{safe}}$, where $\epsilon_{\text{safe}} > 0$ is a parameter that may be specified as desired. However, the goal is still to determine the coordinate $i^* := \arg \max_i f_i(g_i^{-1}(\gamma))$ that achieves the largest safe value.

6.1 Algorithm Description

We propose a modification of SAFEBAI-LINEAR that takes into account the Lipschitz, monotonic nature of the function to guarantee safe learning, SAFEBAI-MONOTONIC. Similar to SAFEBAI-LINEAR, SAFEBAI-MONOTONIC proceeds in epochs, maintaining estimates of the largest verifiably safe values for each coordinate, and refining the tolerance to learn better estimates of the function values.

At a high level, SAFEBAI-MONOTONIC maintains a set of active coordinates at epoch ℓ , \mathcal{X}_ℓ , that it has not yet shown are suboptimal. For each active coordinate i it maintains two values: $\widehat{a}_s^{n_i, \ell}(i)$ and $\widehat{a}_u^{m_i, \ell}(i)$. $\widehat{a}_s^{n_i, \ell}(i)$ is its estimate of the largest safe value, and on Lines 8 to 12 it increases $\widehat{a}_s^{n_i, \ell}(i)$ as close as possible to the safety threshold $g_i^{-1}(\gamma)$, while ensuring $g_i(\widehat{a}_s^{n_i, \ell}(i)) \leq \gamma$. In contrast, $\widehat{a}_u^{m_i, \ell}(i)$ is an estimate of the minimum unsafe value. $\widehat{a}_u^{m_i, \ell}(i)$ is first updated on Lines 13 to 20 until it exceeds the safety threshold, $g_i(\widehat{a}_u^{m_i, \ell}(i)) > \gamma$, while ensuring that $g_i(\widehat{a}_u^{m_i, \ell}(i)) \leq \gamma + \epsilon_{\text{safe}}$. After this threshold is reached, $\widehat{a}_u^{m_i, \ell}(i)$ is then *decreased* on Line 22 to 29, so that $\widehat{a}_u^{m_i, \ell}(i)$ approaches $g_i^{-1}(\gamma)$ from above.

Since f_i is non-increasing, we will always then have that $f_i(\widehat{a}_s^{n_i, \ell}(i))$ is a lower bound on the optimal safe value of coordinate i . Furthermore, once we have reached the condition that $g_i(\widehat{a}_u^{m_i, \ell}(i)) > \gamma$, $f_i(\widehat{a}_u^{m_i, \ell}(i))$ will always be an upper bound on the largest safe value of coordinate i . By estimating these respective values, and refining the estimates of $\widehat{a}_s^{n_i, \ell}(i)$ and $\widehat{a}_u^{m_i, \ell}(i)$ so that they are as close as possible to the safety threshold $g_i^{-1}(\gamma)$, we are able to compare the value of coordinates, and eliminate coordinates whose upper bound is smaller than the lower bound of a different coordinate. We provide additional details on the operation of SAFEBAI-MONOTONIC below.

Safe Value Updates. SAFEBAI-MONOTONIC up-

Algorithm 2 Safe Best-Arm Identification for Monotonic Functions (SAFEBAI-MONOTONIC)

```

1: input: Confidence  $\delta$ , safety gap  $\epsilon_{\text{safe}}$ 
2: initialize:  $\widehat{a}_s^{0,0}(i) \leftarrow a_{0,i}, \widehat{a}_u^{0,0}(i) \leftarrow a_{0,i}, \text{unsafe}(i) \leftarrow 0, \mathcal{X}_0 \leftarrow [d], \ell \leftarrow 1, t \leftarrow 1, \epsilon_\ell \leftarrow 2^{-\ell}, n_i \leftarrow 1, m_i \leftarrow 1$ 
3: while  $|\mathcal{X}_{\ell-1}| > 1$  do
4:    $N_{\ell,t} \leftarrow \lceil 2\sigma^2 \log \frac{8t^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$ 
5:   for  $i \in \mathcal{X}_{\ell-1}$  do
6:      $\widehat{a}_s^{0,\ell}(i) \leftarrow \widehat{a}_s^{n_i-1, \ell-1}(i), n_i \leftarrow 1$ 
7:      $\widehat{a}_u^{0,\ell}(i) \leftarrow \widehat{a}_u^{m_i-1, \ell-1}(i), m_i \leftarrow 1$ 
8:     Estimate $_i(\widehat{a}_s^{0,\ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
9:     while  $\gamma - \widehat{g}_i(\widehat{a}_s^{n_i-1, \ell}(i)) > 2\epsilon_\ell$  do
10:      // increase safe value to  $g_i^{-1}(\gamma)$ 
11:       $\widehat{a}_s^{n_i, \ell}(i) \leftarrow \gamma + \widehat{a}_s^{n_i-1, \ell}(i) - \widehat{g}_i(\widehat{a}_s^{n_i-1, \ell}(i)) - \epsilon_\ell$ 
12:      Estimate $_i(\widehat{a}_s^{n_i, \ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
13:       $n_i \leftarrow n_i + 1$ 
14:     if  $\text{unsafe}(i) = 0$  then
15:      // increase unsafe value to  $g_i^{-1}(\gamma + \epsilon_{\text{safe}})$ 
16:      Estimate $_i(\widehat{a}_u^{0,\ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
17:      while  $\gamma + \epsilon_{\text{safe}} - \widehat{g}_i(\widehat{a}_u^{m_i-1, \ell}(i)) > 2\epsilon_\ell$  do
18:         $\widehat{a}_u^{m_i, \ell}(i) \leftarrow \gamma + \epsilon_{\text{safe}} + \widehat{a}_u^{m_i-1, \ell}(i) - \widehat{g}_i(\widehat{a}_u^{m_i-1, \ell}(i)) - \epsilon_\ell$ 
19:        Estimate $_i(\widehat{a}_u^{m_i, \ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
20:        if  $\widehat{g}_i(\widehat{a}_u^{m_i, \ell}(i)) - \epsilon_\ell \geq \gamma$  then
21:           $\text{unsafe}(i) \leftarrow 1, \text{break}$ 
22:         $m_i \leftarrow m_i + 1$ 
23:      else
24:        // decrease unsafe value to  $g_i^{-1}(\gamma)$ 
25:         $\widehat{a}_u^{1, \ell}(i) \leftarrow \widehat{a}_u^{0, \ell}(i)/2 + \widehat{a}_s^{n_i-1, \ell}(i)/2, m_i \leftarrow 2$ 
26:        Estimate $_i(\widehat{a}_u^{1, \ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
27:        while  $\widehat{g}_i(\widehat{a}_u^{m_i-1, \ell}(i)) - \epsilon_\ell < \gamma$  do
28:           $\widehat{a}_u^{m_i, \ell}(i) \leftarrow \widehat{a}_u^{0, \ell}(i)/2 + \widehat{a}_u^{m_i-1, \ell}(i)/2$ 
29:          if  $\widehat{a}_u^{0, \ell}(i) - \widehat{a}_u^{m_i, \ell}(i) \leq \epsilon_\ell$  then
30:             $\widehat{a}_u^{m_i, \ell}(i) \leftarrow \widehat{a}_u^{0, \ell}(i), \text{break}$ 
31:          Estimate $_i(\widehat{a}_u^{m_i, \ell}(i), N_{\ell,t}), t \leftarrow t + 1$ 
32:           $m_i \leftarrow m_i + 1$ 
33:        // eliminate suboptimal coordinate
34:         $\mathcal{X}_\ell \leftarrow \mathcal{X}_{\ell-1} \setminus \{i \in \mathcal{X}_{\ell-1} : \text{unsafe}(i) = 1, \widehat{f}_i(\widehat{a}_u^{m_i-1, \ell}(i)) + 2\epsilon_\ell \leq \max_j \widehat{f}_j(\widehat{a}_s^{n_j-1, \ell}(j))\}$ 
35:         $\ell \leftarrow \ell + 1$ 
36:      return  $\mathcal{X}_{\ell-1}$ 
37: function Estimate $_i(a, N)$ 
38:   // Pull coordinate  $i$  at value  $a$   $N$  times and form estimates of function values
39:   for  $t = 1, \dots, N$  do
40:     Pull coordinate  $i$  at  $a$ , observe
41:      $y_t = f_i(a) + \eta_t, z_t = g_i(a) + w_t$ 
42:    $\widehat{f}_i(a) \leftarrow N^{-1} \sum_{t=1}^N y_t, \widehat{g}_i(a) \leftarrow N^{-1} \sum_{t=1}^N z_t$ 

```

dates the estimates of the safe values, $\widehat{a}_s^{n_i, \ell}(i)$, on Line 10 as

$$\widehat{a}_s^{n_i, \ell}(i) \leftarrow \gamma + \widehat{a}_s^{n_i-1, \ell}(i) - \widehat{g}_i(\widehat{a}_s^{n_i-1, \ell}(i)) - \epsilon_\ell.$$

At epoch ℓ , we collect enough samples to guarantee that with high probability $|\widehat{g}_i(\widehat{a}_s^{n_i-1, \ell}(i)) - g_i(\widehat{a}_s^{n_i-1, \ell}(i))| \leq \epsilon_\ell$. Thus, using that g_i is 1-Lipschitz,

$$\begin{aligned} g_i(\widehat{a}_s^{n_i, \ell}(i)) &\leq g_i(\widehat{a}_s^{n_i-1, \ell}(i)) + \gamma - \widehat{g}_i(\widehat{a}_s^{n_i-1, \ell}(i)) - \epsilon_\ell \\ &\leq g_i(\widehat{a}_s^{n_i-1, \ell}(i)) + \gamma - g_i(\widehat{a}_s^{n_i-1, \ell}(i)) \\ &= \gamma \end{aligned}$$

so it follows that $\widehat{a}_s^{n_i, \ell}(i)$ is also safe.

Epoch Update Schedule. Note that we can increase the value of $\widehat{a}_s^{n_i, \ell}(i)$ by at most a factor of roughly $\gamma - g_i(\widehat{a}_s^{n_i, \ell}(i))$ at every iteration of the update to $\widehat{a}_s^{n_i, \ell}(i)$. Assume that $\gamma - g_i(\widehat{a}_s^{n_i, \ell}(i)) \gg \epsilon_\ell$. In this regime it follows the dominant term in the update

$$\widehat{a}_s^{n_i+1, \ell}(i) \leftarrow \gamma + \widehat{a}_s^{n_i, \ell}(i) - \widehat{g}_i(\widehat{a}_s^{n_i, \ell}(i)) - \epsilon_\ell$$

is not ϵ_ℓ —our update increment is well above the “noise floor”. As such, increasing ℓ will not help us, we should instead keep incrementing $\widehat{a}_s^{n_i, \ell}(i)$ until we arrive at a value where $\gamma - g_i(\widehat{a}_s^{n_i, \ell}(i)) \approx \epsilon_\ell$. At this point, we are close enough to the noise floor that decreasing ϵ_ℓ will help us further increase $\widehat{a}_s^{n_i, \ell}(i)$. This is precisely the procedure SAFEBAI-MONOTONIC implements, only increasing ℓ once the updates in each coordinate have reached the noise floor.

Unsafe Value Updates. As noted, to guarantee a coordinate is suboptimal in the monotonic setting, we need to find a sample point a such that $g_i(a) \geq \gamma$ and $g_i(a) \leq \gamma + \epsilon_{\text{safe}}$. Once we have found this value, we have an upper bound on the maximum safe value for coordinate i , yet this value may need to be refined, for example if the gap for coordinate i is small. As such, once we find a with $g_i(a) \geq \gamma$, we must decrease the a value to get as close as possible to the threshold $g_i(a) = \gamma$. SAFEBAI-MONOTONIC implements this by updating the “unsafe” estimate, $\widehat{a}_u^{m_i, \ell}(i)$, in two stages. In the first stage, while $\text{unsafe}(i) = 0$, it applies an update analogous to the safe update, but which instead guarantees it will stay below the threshold $\gamma + \epsilon_{\text{safe}}$. Once it can guarantee it has crossed the safety threshold, $g_i(\widehat{a}_u^{m_i, \ell}(i)) \geq \gamma$, it sets $\text{unsafe}(i) = 1$ and decreases $\widehat{a}_u^{m_i, \ell}(i)$ while ensuring $g_i(\widehat{a}_u^{m_i, \ell}(i)) \geq \gamma$ via a binary search procedure.

The elimination criteria of SAFEBAI-MONOTONIC is similar to that of SAFEBAI-LINEAR but differs in that it only eliminates a coordinate after it has crossed the safety threshold, allowing it to obtain an upper bound on that coordinate’s value.

6.2 Sample Complexity

Before stating the sample complexity of SAFEBAI-MONOTONIC, we make an additional assumption.

Assumption 2. For all $x \in [\gamma, \gamma + \epsilon_{\text{safe}}]$ and each $i \in [d]$, there exists some $a \in \mathbb{R}$ such $g_i(a) = x$. As such, the inverse $g_i^{-1}(x)$ is well defined for all $x \in [\gamma, \gamma + \epsilon_{\text{safe}}]$.

This assumption is primarily for technical reasons and allows us to simplify the results somewhat.

Let $\tilde{x}_i := \inf_x x$ s.t. $\exists a, g_i(a) = x$ and define \tilde{g}_i^{-1} as:

$$\tilde{g}_i^{-1}(x) = \begin{cases} a_{0,i} & x \leq \tilde{x}_i \\ g_i^{-1}(x) & \tilde{x}_i < x \leq \gamma + \epsilon_{\text{safe}} \\ g_i^{-1}(\gamma + \epsilon_{\text{safe}}) & x > \gamma + \epsilon_{\text{safe}} \end{cases}.$$

In words, \tilde{g}_i^{-1} extends the inverse of g_i^{-1} to outside its range. Now define the following:

$$\begin{aligned} \bar{n}_{i, \ell} &:= \begin{cases} \frac{\tilde{g}_i^{-1}(\gamma - \epsilon_\ell) - \tilde{g}_i^{-1}(\gamma - 6\epsilon_\ell)}{\epsilon_\ell} & \ell \geq 2 \\ 2(\tilde{g}_i^{-1}(\gamma - \frac{\epsilon_\ell}{2}) - a_{0,i}) & \ell = 1 \end{cases}, \\ \bar{m}_{i, \ell} &:= \begin{cases} \frac{\tilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - \epsilon_\ell) - \tilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 6\epsilon_\ell)}{\epsilon_\ell} & \ell \geq 2 \\ 2(\tilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{\epsilon_\ell}{2}) - a_{0,i}) & \ell = 1 \end{cases}, \end{aligned}$$

$$\begin{aligned} \bar{a}_u^\ell(i) &:= \sum_{s=1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}} \\ &\quad + \left(\ell + \frac{4g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{\epsilon_{\text{safe}}} \right) 2^{-\ell}, \end{aligned}$$

$$\bar{\ell}(i) := \arg \min_{\ell \in \mathbb{N}}$$

$$\text{s.t. } f_i(\bar{a}_u^\ell(i)) + 4\epsilon_\ell \leq f_{i^*}(\tilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell))$$

and let $\bar{\ell}(i^*) := \max_{i \neq i^*} \bar{\ell}(i)$. We the have the following result.

Theorem 3. Under Assumption 1 and 2, with probability $1 - \delta$, Algorithm 2 will terminate and output i^* after collecting at most

$$\tilde{O} \left(\log \frac{1}{\delta} \cdot \sum_{i=1, i \neq i^*}^d \sum_{\ell=1}^{\bar{\ell}(i)} (\bar{m}_{i, \ell} + \bar{n}_{i, \ell} + \ell + 2) 2^{2\ell} \right)$$

samples, and will only pull arms during execution satisfying $g_i(a) \leq \gamma + \epsilon_{\text{safe}}$.

We prove this result in Appendix C. While this is a closed-form and deterministic expression, it is difficult to interpret in general. To obtain a more interpretable expression, we make the following additional assumption.

Assumption 3. For all $i \in [d]$, g_i is differentiable on \mathbb{R} , and $g_i^{-1}(x)$ is well-defined for all $x \in [\gamma - 3/2, \gamma + \epsilon_{\text{safe}}]$. Furthermore, $g_i'(a) \geq L^{-1}$ for all $a \in [g_i^{-1}(\gamma - 3/2), g_i^{-1}(\gamma + \epsilon_{\text{safe}})]$ and some $L > 0$.

We do not assume that the learner knows either the value of L or that g_i satisfies Assumption 3. Under this assumption, we obtain the following corollary.

Corollary 2. *Assume that Assumption 3 holds. Then, with probability $1-\delta$, Algorithm 2 will terminate and output i^* after collecting at most*

$$\tilde{O}\left(\log \frac{1}{\delta} \cdot \sum_{i \neq i^*} \frac{1 + L^3 + (1 + L)g_i^{-1}(\gamma)^2/\epsilon_{\text{safe}}^2}{\Delta_i^2} + \sum_{i=1}^d \left(g_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{1}{2}) + g_i^{-1}(\gamma - \frac{1}{2}) - 2a_{0,i}\right)\right)$$

samples, and will only pull arms during execution satisfying $g_i(a) \leq \gamma + \epsilon_{\text{safe}}$.

Corollary 2 provides the familiar “sum over inverse gap squared” complexity, with the additional polynomial dependence on L , as well as the terms $g_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{1}{2}) + g_i^{-1}(\gamma - \frac{1}{2}) - 2a_{0,i}$ and $g_i^{-1}(\gamma)^2/\epsilon_{\text{safe}}^2$. The former term, $g_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{1}{2}) + g_i^{-1}(\gamma - \frac{1}{2}) - 2a_{0,i}$, results from the initial phase of learning needed to guarantee that we are in the “neighborhood” of the safety boundary, $g_i^{-1}(\gamma)$, and increases as $a_{0,i}$ starts farther from the safety boundary. The latter term results from the complexity needed to learn a value above the safety threshold—if ϵ_{safe} is small, we must explore more conservatively, which will increase our complexity.

7 EXPERIMENTAL RESULTS

In this section, we present experimental results on our algorithms and compare them with two existing models: the safe linear Thompson Sampling (SAFE-LTS) algorithm of Moradipari et al. (2021), and the SAFEOPT algorithm of Sui et al. (2015). We note that the STAGEOPT algorithm proposed in Sui et al. (2018) relies on similar design principles to SAFEOPT, which will cause its failure modes on the instances we present to be similar to that of SAFEOPT. As such, we only plot results for SAFEOPT.

7.1 Algorithm Setup

We demonstrate the superior performance of our algorithm in both a linear function model and a drug response model. We must point out that in both Moradipari et al. (2021) and Sui et al. (2015), the problem setups do not aim to identify a best coordinate with high confidence—Moradipari et al. (2021) seeks to minimize regret, and Sui et al. (2015) is unverifiable, their algorithm does not provide a stopping condition. For effective comparison with our algorithms, we must choose proper stopping criteria for these models. For SAFE-LTS, we adopt a similar stopping cri-

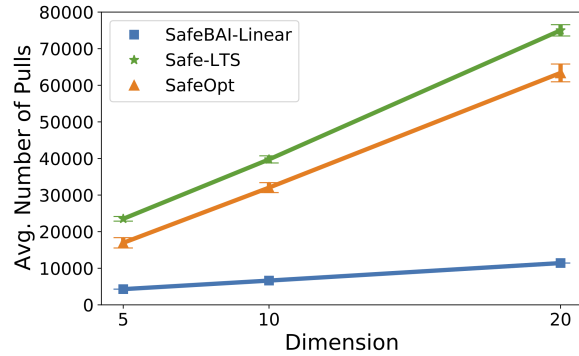


Figure 1: Total arm pulls to termination vs. dimension in the linear model. Error bars give 1 standard deviation.

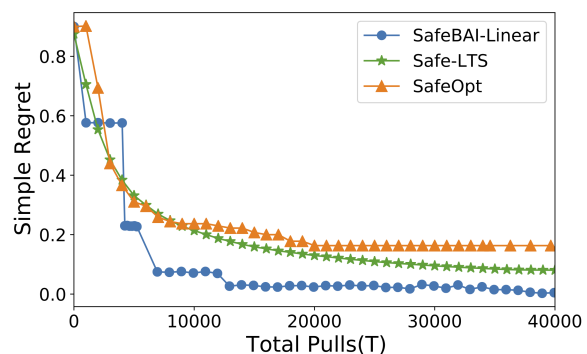


Figure 2: Simple regret in linear model with $d = 10$.

teria as SAFEBAI-LINEAR, with the confidence interval expression following SAFE-LTS’s definition. For SAFEOPT, we discretize the continuous values each coordinate can take into 50 evenly spaced values. We stop when SAFEOPT suggests all possibly optimal safe values have been found and there is a coordinate whose reward lower confidence bound is above all the reward upper confidence bounds of the remaining coordinates. In all our experiments, every algorithm always found the best coordinate, and none pulled any unsafe arms. More details on the experimental setup can be found in Appendix D.

7.2 Linear Response Model

In our first experiment, we consider a setup with $d \in \{5, 10, 20\}$ coordinates. We choose $\theta_1 = 1, \theta_2 = 0.9, \theta_{i>2} = 1$ and $\mu_1 = 1, \mu_2 = 1.5, \mu_{i>2} = 5$, and set the safety threshold γ to 1. With this setting, the minimum gap is $\Delta = 0.4$, and the remaining gaps are all 0.8. The sample observations are perturbed by Gaussian noise with mean 0 and variance $\sigma^2 = 0.5$. The average simple regret values at $d = 10$ are also computed to illustrate how fast the choice of coordinates in each pull is improving. We define the simple regret as $r_t = \theta_{i^*} a_{i^*} - \theta_{\hat{i}_t} \hat{a}_s^t(\hat{i}_t)$, for $\hat{i}_t = \arg \max_i \hat{\theta}_{i,t} \hat{a}_s^t(i)$, where $\hat{\theta}_{i,t}$ is each algorithm’s estimate of θ_i at time t ,

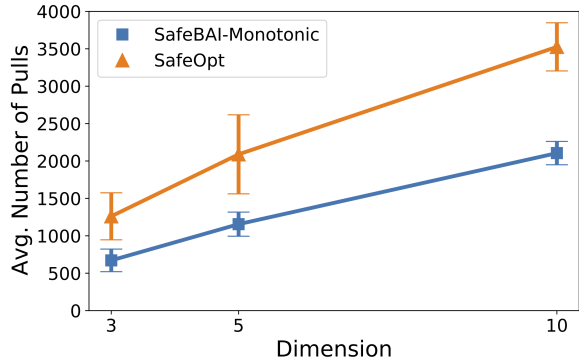


Figure 3: Total arm pulls to termination vs. dimension in drug response model. Error bars give 1 standard deviation.

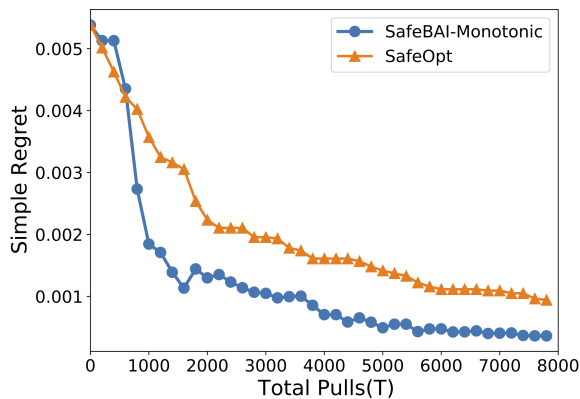


Figure 4: Simple regret in drug response model with $d = 20$.

and $\hat{a}_s^t(i)$ is the largest verifiably safe value for coordinate i at time t . In this setting, we run SAFEOPT with a linear kernel. All data points are the average of 10 trials.

From Figure 1, we observe that the number of arm pulls scales linearly with d , and SAFEBAI-LINEAR outperforms the other two algorithms significantly. Next, as seen from Figure 2, SAFEBAI-LINEAR shows the fastest decrease in simple regret as the number of arm pulls increases, and attains near zero regret after the best coordinate is identified. On the other hand, both SAFE-LTS and SAFEOPT show a slower rate of decrease in simple regret. Intuitively the worse performances of the other two algorithms may be attributed to unnecessary pulls wasted to find the largest safe value for suboptimal coordinates. In contrast, SAFEBAI-LINEAR can quickly identify and remove the suboptimal coordinates before their largest safe values is reached, preventing overexploration.

7.3 Application to Best Drug Identification

In the second experiment, we investigate the performance of SAFEBAI-MONOTONIC in comparison to

SAFEOPT on a nonlinear drug response model. In the dose-finding literature, it is often assumed that higher dosage leads to stronger drug response and efficacy, while also increasing toxicity (Mandrekar et al., 2007; Yuan & Yin, 2009; Cai et al., 2014). A common assumption (see e.g. Thall & Russell (1998) and aforementioned works) has been to rely on a logistic function to model both the efficacy and toxicity: for the i th drug and dose a , the efficacy is modeled as $f_i(a) = \frac{1}{1+e^{-\theta_i \cdot a}}$ and toxicity as $g_i(a) = \frac{1}{1+e^{-\mu_i \cdot a}}$. Note that this choice of f_i and g_i is monotonic and fits within our setting. Applying this model, we can effectively consider a drug selection setup where our goal is to identify the drug with highest utility among a set of candidate drugs, while ensuring none of the dosages tested lead to toxic response, i.e. all drug tests satisfy a safety constraint. We design drug sets with a number of $d \in \{3, 5, 10, 20\}$ drugs, efficacy model parameters $\theta_1 = 0.01, \theta_{i>1} = 10$, toxicity model parameters $\mu_i = 1$. This gives a minimum gap $\Delta \approx 0.497$. The safety threshold γ is set to 0.3. For SAFEOPT, we use an RBF kernel. The sample observations are perturbed by Gaussian noise with mean 0 and variance $\sigma^2 = 0.1$. Similar to the linear case, we will also plot the simple regret values, where here we define the simple regret as $r_t = f_{i^*}(a_{i^*}) - f_{\hat{i}_t}(\hat{a}_s^t(\hat{i}_t))$, for $\hat{i}_t = \arg \max_i \hat{f}_i(\hat{a}_s^t(i))$. We run SAFEOPT with an RBF kernel in this setting. All data points are the average of 20 trials.

The result in Figure 3 suggests SAFEBAI-MONOTONIC is able to much more effectively identify the best safe drug than SAFEOPT. The total number of drug evaluations required for SAFEBAI-MONOTONIC is roughly half of SAFEOPT in all instances. Furthermore, the expensive posterior update in SAFEOPT makes it slow to run with samples > 5000 , while SAFEBAI-MONOTONIC can be efficiently run when many more samples are required, allowing it to easily scale to trials with more drugs. Figure 4 illustrates that SAFEBAI-MONOTONIC also performs much better in terms of the simple regret. Not only is it able to identify and verify the optimal coordinate more quickly, even before verifying it has found the optimal coordinate, its sampling strategy allows it to obtain better intermediate estimates of the best coordinate.

Acknowledgements

The work of AW is supported by an NSF GFRP Fellowship DGE-1762114. The work of KJ is supported in part by grants NSF CCF 2007036 and NSF TRIPODS 2023166.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. *Advances in neural information processing systems*, 24: 2312–2320, 2011.
- Amani, S., Alizadeh, M., and Thrampoulidis, C. Linear stochastic bandits under safety constraints. *arXiv preprint arXiv:1908.05814*, 2019.
- Aziz, M., Kaufmann, E., and Riviere, M.-K. On multi-armed bandit designs for dose-finding clinical trials. *Journal of Machine Learning Research*, 22:1–38, 2021.
- Bechhofer, R. E. A sequential multiple-decision procedure for selecting the best one of several normal populations with a common unknown variance, and its use with various experimental designs. *Biometrics*, 14(3):408–429, 1958.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in multi-armed bandits problems. In *International conference on Algorithmic learning theory*, pp. 23–37. Springer, 2009.
- Cai, C., Yuan, Y., and Ji, Y. A bayesian dose-finding design for oncology clinical trials of combinational biological agents. *Journal of the Royal Statistical Society. Series C, Applied statistics*, 63(1):159, 2014.
- Chen, S., Lin, T., King, I., Lyu, M. R., and Chen, W. Combinatorial pure exploration of multi-armed bandits. *Advances in neural information processing systems*, 27:379–387, 2014.
- Cheshire, J., Ménard, P., and Carpentier, A. The influence of shape constraints on the thresholding bandit problem. In *Conference on Learning Theory*, pp. 1228–1275. PMLR, 2020.
- Even-Dar, E., Mannor, S., and Mansour, Y. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pp. 255–270. Springer, 2002.
- Flaxman, A. D., Kalai, A. T., and McMahan, H. B. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 385–394, 2005.
- Garivier, A., Ménard, P., Rossi, L., and Menard, P. Thresholding bandit for dose-ranging: The impact of monotonicity. *arXiv preprint arXiv:1711.04454*, 2017.
- Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. lil’ucb: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pp. 423–439. PMLR, 2014.
- Jamieson, K. G., Nowak, R. D., and Recht, B. Query complexity of derivative-free optimization. *arXiv preprint arXiv:1209.2434*, 2012.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pp. 655–662, 2012.
- Karnin, Z., Koren, T., and Somekh, O. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pp. 1238–1246. PMLR, 2013.
- Kaufmann, E., Cappé, O., and Garivier, A. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.
- Kazerouni, A., Ghavamzadeh, M., Abbasi-Yadkori, Y., and Van Roy, B. Conservative contextual linear bandits. *arXiv preprint arXiv:1611.06426*, 2016.
- Locatelli, A., Gutzeit, M., and Carpentier, A. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning*, pp. 1690–1698. PMLR, 2016.
- Mandrekar, S. J., Cui, Y., and Sargent, D. J. An adaptive phase i design for identifying a biologically optimal dose for dual agent drug combinations. *Statistics in medicine*, 26(11):2317–2330, 2007.
- Mannor, S. and Tsitsiklis, J. N. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5(Jun): 623–648, 2004.
- Moradipari, A., Thrampoulidis, C., and Alizadeh, M. Stage-wise conservative linear bandits. *Advances in Neural Information Processing Systems*, 33, 2020.
- Moradipari, A., Amani, S., Alizadeh, M., and Thrampoulidis, C. Safe linear thompson sampling with side information. *IEEE Transactions on Signal Processing*, 2021.
- Musuamba, F. T., Manolis, E., Holford, N., Cheung, S. A., Friberg, L. E., Ogungbenro, K., Posch, M., Yates, J., Berry, S., Thomas, N., et al. Advanced methods for dose and regimen finding during drug development: summary of the ema/efpia workshop on dose finding (london 4–5 december 2014). *CPT: pharmacometrics & systems pharmacology*, 6(7):418–429, 2017.
- Pacchiano, A., Ghavamzadeh, M., Bartlett, P., and Jiang, H. Stochastic bandits with linear constraints. In *International Conference on Artificial Intelligence and Statistics*, pp. 2827–2835. PMLR, 2021.
- Paulson, E. A sequential procedure for selecting the population with the largest mean from k normal populations. *The Annals of Mathematical Statistics*, pp. 174–180, 1964.

- Rios, L. M. and Sahinidis, N. V. Derivative-free optimization: a review of algorithms and comparison of software implementations. *Journal of Global Optimization*, 56(3):1247–1293, 2013.
- Riviere, M.-K., Yuan, Y., Jourdan, J.-H., Dubois, F., and Zohar, S. Phase i/ii dose-finding design for molecularly targeted agent: plateau determination using adaptive randomization. *Statistical methods in medical research*, 27(2):466–479, 2018.
- Rogatko, A., Babb, J. S., Tighiouart, M., Khuri, F. R., and Hudes, G. New paradigm in dose-finding trials: patient-specific dosing and beyond phase i. *Clinical cancer research*, 11(15):5342–5346, 2005.
- Soare, M., Lazaric, A., and Munos, R. Best-arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 27:828–836, 2014.
- Sui, Y., Gotovos, A., Burdick, J., and Krause, A. Safe exploration for optimization with gaussian processes. In *International Conference on Machine Learning*, pp. 997–1005. PMLR, 2015.
- Sui, Y., Burdick, J., Yue, Y., et al. Stagewise safe bayesian optimization with gaussian processes. In *International Conference on Machine Learning*, pp. 4781–4789. PMLR, 2018.
- Thall, P. F. and Cook, J. D. Dose-finding based on efficacy–toxicity trade-offs. *Biometrics*, 60(3):684–693, 2004.
- Thall, P. F. and Russell, K. E. A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase i/ii clinical trials. *Biometrics*, pp. 251–264, 1998.
- Valko, M., Carpentier, A., and Munos, R. Stochastic simultaneous optimistic optimization. In *International Conference on Machine Learning*, pp. 19–27. PMLR, 2013.
- Yuan, Y. and Yin, G. Bayesian dose finding by jointly modelling toxicity and efficacy as time-to-event outcomes. *Journal of the Royal Statistical Society: Series C (Applied Statistics)*, 58(5):719–736, 2009.

A Linear Response Functions Proof

A.1 Correctness of SafeBAI-Linear

Lemma 1. For any $i \in [d]$ and $\ell \in \mathbb{N}^+$, define the events

$$\mathcal{E}_{i,\ell}(\theta) = \left\{ |\widehat{\theta}_{i,\ell} - \theta_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right\}, \quad \mathcal{E}_{i,\ell}(\mu) = \left\{ |\widehat{\mu}_{i,\ell} - \mu_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right\} \quad (1)$$

and let

$$\mathcal{E} := \left(\bigcap_{i=1}^d \bigcap_{\ell=1}^{\infty} \mathcal{E}_{i,\ell}(\theta) \right) \cap \left(\bigcap_{i=1}^d \bigcap_{\ell=1}^{\infty} \mathcal{E}_{i,\ell}(\mu) \right).$$

Then it holds that $\mathbb{P}(\mathcal{E}) \geq 1 - \delta$.

Proof of Lemma 1. Given that at ℓ -th iteration, $N_\ell = \lceil 2\sigma^2 \log \frac{8d\ell^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$ and

$$\widehat{a}_s^{\ell-1}(i)\widehat{\theta}_{i,\ell} = \frac{1}{N_\ell} \sum_{t=1}^{n_i} (\widehat{a}_s^{\ell-1}(i)\theta_i + \eta_t) = \widehat{a}_s^{\ell-1}(i)\theta_i + \frac{1}{n_i} \sum_{t=1}^{N_\ell} \eta_t.$$

By Hoeffding's inequality, $|\widehat{a}_s^{\ell-1}(i)\theta_i - \widehat{a}_s^{\ell-1}(i)\widehat{\theta}_{i,\ell}| \leq \sqrt{\frac{2\sigma^2 \log(\frac{8d\ell^2}{\delta})}{N_\ell}} = 2^{-\ell}$ with probability $1 - \frac{\delta}{4d\ell^2}$. Thus we have that $\mathbb{P}(\mathcal{E}_{i,\ell}^c(\theta)) = \mathbb{P}(\mathcal{E}_{i,\ell}^c(\mu)) \leq \frac{\delta}{4d\ell^2}$. Taking union bound over all coordinates and rounds we have:

$$\mathbb{P}\left(\bigcup_{i=1}^d \bigcup_{\ell=1}^{\infty} \mathcal{E}_{i,\ell}^c(\mu)\right) \leq \sum_{i=1}^d \sum_{\ell=1}^{\infty} \frac{\delta}{4d\ell^2} = d \cdot \frac{\delta}{4d} \cdot \frac{\pi^2}{6} \leq \delta/2. \quad (2)$$

Using the same inequality on the union bound we can obtain the inequality result for θ as well, and union bounding over both gives the result. \square

Lemma 2. On the event \mathcal{E} , for all ℓ and $i \in \mathcal{X}_{\ell-1}$, we will have that $\widehat{a}_s^\ell(i) \in \mathcal{S}_i$.

Proof. On the event \mathcal{E} , $|\widehat{\mu}_{i,\ell} - \mu_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$, $|\widehat{\theta}_{i,\ell} - \theta_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$. Then using our choice of $\widehat{a}_s^\ell(i)$, it follows that $\widehat{a}_s^\ell(i)\mu_i \leq \widehat{a}_s^\ell(i)(\widehat{\mu}_{i,\ell} + \frac{\epsilon_\ell}{\widehat{a}_s^{\ell-1}(i)}) = \gamma$. Thus, by construction, $\widehat{a}_s^\ell(i)$ will be safe. \square

Lemma 3. On the event \mathcal{E} , coordinate i is suboptimal if the event

$$\left\{ \exists j \in [n], \ell > 0 : \widehat{a}_u^\ell(i) \left(\widehat{\theta}_{i,\ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) < \widehat{a}_s^\ell(j) \left(\widehat{\theta}_{j,\ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(j)} \right) \right\} \quad (3)$$

holds, so $i^* \in \mathcal{X}_\ell$ for all ℓ .

Proof of Lemma 3. Fix coordinate i and its largest safe value $a_s(i)$. By Lemma 2, we will have that $\widehat{a}_s^\ell(i)$ is safe. Further, applying our choice of $\widehat{a}_u^\ell(i)$, we obtain

$$\widehat{a}_s^\ell(i) = M_i \wedge \gamma / (\widehat{\mu}_{i,\ell} + \frac{\epsilon_\ell}{\widehat{a}_s^{\ell-1}(i)}) \leq a_s(i) = M_i \wedge \gamma / \mu_i \leq M_i \wedge \gamma / (\widehat{\mu}_{i,\ell} - \frac{\epsilon_\ell}{\widehat{a}_s^{\ell-1}(i)}) = \widehat{a}_u^\ell(i).$$

Applying this set of inequalities $\widehat{a}_s^\ell(i) \leq a_s(i) \leq \widehat{a}_u^\ell(i)$, it follows that

$$\begin{aligned} \widehat{a}_u^\ell(i) \left(\widehat{\theta}_{i,\ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(j) \left(\widehat{\theta}_{j,\ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(j)} \right) \\ \implies a_s(i) \left(\widehat{\theta}_{i,\ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(j) \left(\widehat{\theta}_{j,\ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(j)} \right) \\ \implies a_s(i)\theta_i &< a_s(j)\theta_j. \end{aligned}$$

The last inequality suggests i being suboptimal by definition. That $i^* \in \mathcal{X}_\ell$ for all ℓ follows directly since (3) is identical to the condition in Algorithm 1 to remove a coordinate, and so, since (3) implies i is suboptimal, it follows that i^* can never be removed. \square

A.2 Coordinate Elimination Criteria

Lemma 4. Assume $\widehat{a}_s^\ell(i) < M_i$ and $\widehat{a}_s^\ell(i^*) < M_{i^*}$ and that \mathcal{E} holds. Let $\underline{a}(i, \ell)$ be a lower bound for $\widehat{a}_s^\ell(i)$, then any suboptimal coordinate i will be eliminated after the ℓ -th iteration where ℓ satisfies:

$$\frac{\theta_i \underline{a}(i, \ell - 1) + 2^{1-\ell}}{\mu_i \underline{a}(i, \ell - 1) - 2^{1-\ell}} \leq \frac{\theta_{i^*} \underline{a}(i^*, \ell - 1) - 2^{1-\ell}}{\mu_{i^*} \underline{a}(i^*, \ell - 1) + 2^{1-\ell}}. \quad (4)$$

Proof of Lemma 4. Since on \mathcal{E} , $i^* \in \mathcal{X}_\ell$ for all ℓ as shown in Lemma 3, a sufficient condition to eliminate coordinate i is that i^* has a larger estimated maximum safe value, and we can therefore reduce the elimination criteria to simply a comparison between i and i^* . First we transform this inequality via:

$$\frac{\theta_i \underline{a}(i, \ell - 1) + 2^{1-\ell}}{\mu_i \underline{a}(i, \ell - 1) - 2^{1-\ell}} \leq \frac{\theta_{i^*} \underline{a}(i^*, \ell - 1) - 2^{1-\ell}}{\mu_{i^*} \underline{a}(i^*, \ell - 1) + 2^{1-\ell}} \implies \frac{\gamma(\theta_i + \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)})}{\mu_i - \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)}} \leq \frac{\gamma(\theta_{i^*} - \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)})}{\mu_{i^*} + \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)}} \quad (5)$$

Note that, by their definition and on \mathcal{E} , $\widehat{a}_s^\ell(i^*) \geq \frac{\gamma 2^{1-\ell}}{\mu_{i^*} + \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)}}$, $\widehat{a}_u^\ell(i) \leq \frac{\gamma 2^{1-\ell}}{\mu_i - \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)}}$. Using this and that $\underline{a}(i, \ell) \leq \widehat{a}_s^\ell(i)$ and $|\widehat{\theta}_{i, \ell} - \theta_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$, we have that (5) implies:

$$\begin{aligned} \widehat{a}_u^\ell(i) \left(\theta_i + \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)} \right) &< \widehat{a}_s^\ell(i^*) \left(\theta_{i^*} - \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)} \right) \\ \implies \widehat{a}_u^\ell(i) \left(\theta_i + 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(i^*) \left(\theta_{i^*} - 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)} \right) \\ \implies \widehat{a}_u^\ell(i) \left(\widehat{\theta}_{i, \ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(i^*) \left(\widehat{\theta}_{i^*, \ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)} \right) \end{aligned}$$

From the last inequality and Lemma 3, we have that i stops getting sampled at such ℓ . \square

Lemma 5. On \mathcal{E} , for any saturated coordinate i (i.e. $a_s(i) = M_i$), when i^* is not saturated, we will eliminate i after the ℓ -th iteration when ℓ satisfies:

$$M_i \theta_i + 2^{1-\ell} < \gamma \cdot \frac{\theta_{i^*} \underline{a}(i^*, \ell - 1) - 2^{1-\ell}}{\mu_{i^*} \underline{a}(i^*, \ell - 1) + 2^{1-\ell}} \quad (6)$$

Proof of Lemma 5. Since on \mathcal{E} , $i^* \in \mathcal{X}_\ell$ for all ℓ as shown in Lemma 3, a sufficient condition to eliminate coordinate i is that i^* has a larger estimated maximum safe value, and we can therefore reduce the elimination criteria to simply a comparison between i and i^* . First we transform this inequality via:

$$M_i \theta_i + 2^{1-\ell} < \gamma \cdot \frac{\theta_{i^*} \underline{a}(i^*, \ell - 1) - 2^{1-\ell}}{\mu_{i^*} \underline{a}(i^*, \ell - 1) + 2^{1-\ell}} \implies M_i \left(\theta_i + \frac{2^{1-\ell}}{M_i} \right) < \frac{\gamma \cdot (\theta_{i^*} - \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)})}{\mu_{i^*} + \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)}} \quad (7)$$

Note that $\widehat{a}_s^\ell(i) = M_i = \widehat{a}_s^{\ell-1}(i)$, $\widehat{a}_s^\ell(i^*) \geq \frac{\gamma 2^{1-\ell}}{\mu_{i^*} + \frac{2^{1-\ell}}{\underline{a}(i^*, \ell - 1)}}$. Using $\underline{a}(i, \ell) \leq \widehat{a}_s^\ell(i)$ and $|\widehat{\theta}_{i, \ell} - \theta_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$, we have that (7) implies:

$$\begin{aligned} M_i \left(\theta_i + \frac{2^{1-\ell}}{M_i} \right) &< \widehat{a}_s^\ell(i^*) \left(\theta_{i^*} - 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)} \right) \\ \implies \widehat{a}_u^\ell(i) \left(\theta_i + 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(i^*) \left(\theta_{i^*} - 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)} \right) \\ \implies \widehat{a}_u^\ell(i) \left(\widehat{\theta}_{i, \ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)} \right) &< \widehat{a}_s^\ell(i^*) \left(\widehat{\theta}_{i^*, \ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)} \right). \end{aligned}$$

From the last inequality and Lemma 3, we have that i stops getting sampled at such ℓ . \square

Lemma 6. *On \mathcal{E} , when i^* is a saturated coordinate (i.e. $a_s(i^*) = M_{i^*}$), then for any unsaturated suboptimal coordinate i , we will eliminate i after the ℓ -th iteration where ℓ satisfies:*

$$\gamma \cdot \frac{\theta_i \underline{a}(i, \ell - 1) + 2^{1-\ell}}{\mu_i \underline{a}(i, \ell - 1) - 2^{1-\ell}} < M_{i^*} \theta_{i^*} - 2^{1-\ell} \quad (8)$$

Proof of Lemma 6. Since on \mathcal{E} , $i^* \in \mathcal{X}_\ell$ for all ℓ as shown in Lemma 3, a sufficient condition to eliminate coordinate i is that i^* has a larger estimated maximum safe value, and we can therefore reduce the elimination criteria to simply a comparison between i and i^* . First we transform this inequality via:

$$\gamma \cdot \frac{\theta_i \underline{a}(i, \ell - 1) + 2^{1-\ell}}{\mu_i \underline{a}(i, \ell - 1) - 2^{1-\ell}} < M_{i^*} \theta_{i^*} - 2^{1-\ell} \implies \frac{\gamma \cdot (\theta_i + \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)})}{\mu_i - \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)}} < M_{i^*} (\theta_{i^*} - \frac{2^{1-\ell}}{M_{i^*}}) \quad (9)$$

Note that $\widehat{a}_s^\ell(i^*) = M_{i^*} = \widehat{a}_s^{\ell-1}(i^*)$, $\widehat{a}_u^\ell(i) \leq \frac{\gamma}{\mu_i - \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)}}$. Using $\underline{a}(i, \ell) \leq \widehat{a}_s^\ell(i)$ and $|\widehat{\theta}_{i, \ell} - \theta_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$, we have (9) implies:

$$\begin{aligned} \widehat{a}_u^\ell(i) (\theta_i + \frac{2^{1-\ell}}{\underline{a}(i, \ell - 1)}) &< M_{i^*} (\theta_{i^*} - \frac{2^{1-\ell}}{M_{i^*}}) \\ \implies \widehat{a}_u^\ell(i) (\theta_i + 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}) &< \widehat{a}_s^\ell(i^*) (\theta_{i^*} - 2 \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)}) \\ \implies \widehat{a}_u^\ell(i) (\widehat{\theta}_{i, \ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}) &< \widehat{a}_s^\ell(i^*) (\widehat{\theta}_{i^*, \ell} - \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i^*)}) \end{aligned}$$

From the last inequality and Lemma 3, we have that i stops getting sampled at such ℓ . \square

A.3 Lower Bounding Safe Value Estimate

Lemma 7. *Define*

$$\underline{a}_s^\ell(i) := \frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}}.$$

Let $\ell_0 := \arg \min_{\ell} \ell$ s.t. $\underline{a}_s^\ell(i) \geq M_i$. Then, on \mathcal{E} :

$$\widehat{a}_s^\ell(i) \geq \begin{cases} \min\{\max\{\underline{a}_s^\ell(i), a_{0,i}\}, M_i\} & \ell < \ell_0 \\ M_i & \ell \geq \ell_0 \end{cases}.$$

Proof of Lemma 7. Recall our choice of $\widehat{a}_s^\ell(i) = \min\{\max\{\gamma/(\widehat{\mu}_{i, \ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}), a_{0,i}\}, M_i\}$. Let $\ell_0 := \arg \min_{\ell} \ell$ s.t. $\underline{a}_s^\ell(i) \geq M_i$. We will first show that for $\ell < \ell_0$, we can lower bound $\widehat{a}_s^\ell(i) \geq \max\{\underline{a}_s^\ell(i), a_{0,i}\}$. We prove this lemma via induction.

In the base case, since on \mathcal{E} , $|\widehat{\mu}_{i, \ell} - \mu_i| \leq \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}$, we have

$$\widehat{a}_s^1(i) = \max\left\{\frac{\gamma}{\widehat{\mu}_{i,1} + \frac{2^{-1}}{a_{0,i}}}, a_{0,i}\right\} \geq \max\left\{\frac{\gamma^1}{\mu_i + 2 \frac{2^{-1}}{a_{0,i}}}, a_{0,i}\right\} = \max\left\{\frac{\gamma^1 \cdot 2}{2\mu_i + \frac{2}{a_{0,i}}}, a_{0,i}\right\} = \max\{\underline{a}_s^1(i), a_{0,i}\},$$

hence the base case holds. Next, suppose the inequality holds for $\ell - 1$, i.e. $\widehat{a}_s^{\ell-1}(i) \geq \max\{\underline{a}_s^{\ell-1}(i), a_{0,i}\}$, then we obtain, on \mathcal{E} ,

$$\begin{aligned} \frac{\gamma}{\widehat{\mu}_{i, \ell} + \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}} &\geq \frac{\gamma}{\mu_i + 2 \cdot \frac{2^{-\ell}}{\widehat{a}_s^{\ell-1}(i)}} \\ &= \frac{\gamma}{\mu_i + 2 \cdot \frac{2^{-\ell}}{\max\{\gamma/(\widehat{\mu}_{i, \ell-1} + \frac{2^{1-\ell}}{\widehat{a}_s^{\ell-2}(i)}), a_{0,i}\}}} \end{aligned}$$

$$\begin{aligned}
 &\geq \frac{\gamma}{\mu_i + 2 \cdot \frac{2^{-\ell}}{\gamma/(\widehat{\mu}_{i,\ell-1} + \frac{2^{1-\ell}}{a_s^{\ell-2}(i)})}} \\
 &\geq \frac{\gamma}{\mu_i + 2 \cdot 2^{-\ell} \frac{\mu_i \sum_{k=1}^{\ell-1} 2^{\ell-1-k} \gamma^{k-1} 2^{\sum_{s=1}^k s} + \frac{2^{\ell-1}}{a_{0,i}}}{\gamma^{\ell-1} 2^{\sum_{k=1}^{\ell-1} k}}} \\
 &= \frac{\gamma \cdot \gamma^{\ell-1} 2^{\sum_{k=1}^{\ell-1} k}}{\mu_i \cdot \gamma^{\ell-1} 2^{\sum_{k=1}^{\ell-1} k} + 2 \cdot 2^{-\ell} (\mu_i \sum_{k=1}^{\ell-1} 2^{\ell-1-k} \gamma^{k-1} 2^{\sum_{s=1}^k s} + \frac{2^{\ell-1}}{a_{0,i}})} \\
 &= \frac{\gamma^\ell 2^{\sum_{k=1}^{\ell-1} k} 2^\ell}{\mu_i \cdot \gamma^{\ell-1} 2^{\sum_{k=1}^{\ell-1} k} 2^\ell + 2 \mu_i \sum_{k=0}^{\ell-1} 2^{\ell-1-k} \gamma^{k-1} 2^{\sum_{s=1}^k s} + \frac{2^\ell}{a_{0,i}}} \\
 &= \frac{\gamma^\ell 2^{\sum_{k=1}^{\ell-1} k}}{\mu_i \sum_{k=1}^{\ell-1} 2^{\ell-k} \gamma^{k-1} 2^{\sum_{s=1}^k s} + \frac{2^\ell}{a_{0,i}}}.
 \end{aligned}$$

Since $\sum_{s=1}^k s = s(s+1)/2$, it follows then that

$$\widehat{a}_s^\ell(i) = \min \left\{ \max \left\{ \frac{\gamma}{\widehat{\mu}_{i,\ell} + \frac{2^{1-\ell}}{a_s^{\ell-1}(i)}}, a_{0,i} \right\}, M_i \right\} \geq \min \{ \max \{ \underline{a}_s^\ell(i), a_{0,i} \}, M_i \}.$$

Since $\ell < \ell_0$, by definition $\underline{a}_s^\ell(i) < M_i$, which implies that $\max \{ \underline{a}_s^\ell(i), a_{0,i} \} \leq M_i$ (since $a_{0,i} \leq M_i$ by assumption), so $\widehat{a}_s^\ell(i) \geq \max \{ \underline{a}_s^\ell(i), a_{0,i} \}$, proving the inductive step. We conclude that for any $\ell < \ell_0$, we can lower bound $\widehat{a}_s^\ell(i) \geq \max \{ \underline{a}_s^\ell(i), a_{0,i} \} = \min \{ \max \{ \underline{a}_s^\ell(i), a_{0,i} \}, M_i \}$.

Now take $\ell = \ell_0$. In this case we can use that $\widehat{a}_s^{\ell_0-1}(i) \geq \max \{ \underline{a}_s^{\ell_0-1}(i), a_{0,i} \}$ and repeat the above calculation to get that

$$\frac{\gamma}{\widehat{\mu}_{i,\ell_0} + \frac{2^{1-\ell_0}}{a_s^{\ell_0-1}(i)}} \geq \underline{a}_s^{\ell_0}(i).$$

Since by definition $\underline{a}_s^{\ell_0}(i) \geq M_i$, we conclude

$$\widehat{a}_s^{\ell_0}(i) = \min \left\{ \max \left\{ \frac{\gamma}{\widehat{\mu}_{i,\ell_0} + \frac{2^{1-\ell_0}}{a_s^{\ell_0-1}(i)}}, a_{0,i} \right\}, M_i \right\} \geq \min \{ \max \{ \underline{a}_s^{\ell_0}(i), a_{0,i} \}, M_i \} \geq \min \{ \max \{ M_i, a_{0,i} \}, M_i \} = M_i.$$

Since we set $\widehat{a}_s^\ell(i)$ to M_i for all epochs after the first epoch for which $\widehat{a}_s^\ell(i) = M_i$, it follows that $\widehat{a}_s^\ell(i) = M_i$ for $\ell \geq \ell_0$. \square

Lemma 8. *We will have that*

$$\underline{a}_s^\ell(i) := \frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}} \geq \frac{\gamma}{(1+\alpha)\mu_i}$$

as long as

$$\ell \geq \Phi_i(\alpha) := \max \left\{ 8, \sqrt{8 \log_2 \frac{2\gamma}{\alpha}}, \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i}\mu_i\alpha}}, 4 \log_2 \frac{2^{1/2}}{\gamma}, \log_2 \frac{4}{\gamma\alpha} + 2 \log_2 \left(\log_2 \frac{4}{\gamma\alpha} \right) \right\}.$$

Proof. We first want to find a lower bound on ℓ that will guarantee

$$2^\ell / a_{0,i} \leq \frac{\alpha}{2} \cdot \mu_i \gamma^{\ell-1} 2^{\ell(\ell+1)/2} \tag{10}$$

since in this case we can lower bound

$$\frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}} \geq \frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{\alpha}{2} \cdot \mu_i \gamma^{\ell-1} 2^{\ell(\ell+1)/2}}$$

$$= \frac{\gamma}{\mu_i} \cdot \frac{1}{\sum_{k=1}^{\ell} \frac{2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}}{\gamma^{\ell-1} 2^{\ell(\ell+1)/2}} + \alpha/2}.$$

To show (10), we can rearrange it as

$$(2^{1/2}/\gamma)^\ell \leq \frac{a_{0,i} \mu_i \alpha}{2\gamma} 2^{\ell^2/2}.$$

Note that $(2^{1/2}/\gamma)^\ell \leq 2^{\ell^2/4}$ for

$$\ell \geq 4 \log_2(2^{1/2}/\gamma).$$

So for ℓ satisfying this, we will have that $(2^{1/2}/\gamma)^\ell \leq \frac{a_{0,i} \mu_i \alpha}{2\gamma} 2^{\ell^2/2}$ as long as

$$2^{\ell^2/4} \leq \frac{a_{0,i} \mu_i \alpha}{2\gamma} 2^{\ell^2/2} \implies \frac{2\gamma}{a_{0,i} \mu_i \alpha} \leq 2^{\ell^2/4} \implies \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i} \mu_i \alpha}} \leq \ell$$

so a sufficient condition to meet (10) is

$$\ell \geq \max \left\{ \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i} \mu_i \alpha}}, 4 \log_2 \frac{2^{1/2}}{\gamma} \right\}.$$

Our goal now is to find ℓ such that

$$1 + \frac{\alpha}{2} \geq \sum_{k=1}^{\ell} \frac{2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}}{\gamma^{\ell-1} 2^{\ell(\ell+1)/2}} = 1 + \sum_{k=1}^{\ell-1} \frac{2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}}{\gamma^{\ell-1} 2^{\ell(\ell+1)/2}}.$$

If we can find such an ℓ , the desired conclusion will follow. Note that

$$2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} \leq \max\{2^\ell, 2\gamma^{\ell-2} 2^{\frac{(\ell-1)\ell}{2}}\}, \quad \forall k \leq \ell-1. \quad (11)$$

To see this, consider

$$\log_2(2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}) = \ell - k + (k-1) \log_2 \gamma + k^2/2 + k/2.$$

Taking the derivative of this with respect to k we get

$$-1/2 + \log_2 \gamma + k.$$

As this is negative for $k < 1/2 - \log_2 \gamma$ and otherwise positive, it follows that the maximum of $2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}$ over the range $k \in [1, \ell-1]$ must either occur at $k=1$, or $k=\ell-1$.

It follows that

$$\sum_{k=1}^{\ell-1} \frac{2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2}}{\gamma^{\ell-1} 2^{\ell(\ell+1)/2}} \leq \frac{(\ell-1) \max\{2^\ell, 2\gamma^{\ell-2} 2^{\frac{(\ell-1)\ell}{2}}\}}{\gamma^{\ell-1} 2^{\ell(\ell+1)/2}} = (\ell-1) \max\{\gamma^{1-\ell} 2^{\ell/2 - \ell^2/2}, \gamma^{-1} 2^{1-\ell}\}.$$

Now

$$(\ell-1) \gamma^{-1} 2^{1-\ell} \leq \alpha/2 \iff \log_2(\ell-1) - \log_2 \frac{\gamma^\alpha}{4} \leq \ell.$$

Assume that $-\log_2 \frac{\gamma^\alpha}{4} > 0$, as otherwise we are trivially done. We claim that the above inequality is satisfied as long as

$$\ell \geq \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right).$$

To see this, note that with $\ell = \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\})$, we have

$$\begin{aligned} \log_2(\ell - 1) - \log_2 \frac{\gamma^\alpha}{4} &= \log_2 \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right) - 1 \right) - \log_2 \frac{\gamma^\alpha}{4} \\ &\leq \log_2 \left(3 \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right) - \log_2 \frac{\gamma^\alpha}{4} \\ &\leq \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right) \\ &= \ell \end{aligned}$$

where the last inequality follows since

$$\begin{aligned} \log_2 \left(3 \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right) &= \log_2 3 + \log_2 \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right) \\ &\leq 2 \log_2 \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right). \end{aligned}$$

To guarantee that

$$(\ell - 1)\gamma^{1-\ell} 2^{\ell/2 - \ell^2/2} \leq \alpha/2 \iff \log_2(\ell - 1) + \ell \log_2(2^{1/2}/\gamma) - \log_2 \frac{\alpha}{2\gamma} \leq \ell^2/2$$

it suffices to take

$$\ell \geq \max\left\{8, \sqrt{-8 \log_2 \frac{\alpha}{2\gamma}}, 4 \log_2 \frac{2^{1/2}}{\gamma}\right\}.$$

Putting all of this together, we have shown that

$$\frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}} \geq \frac{\gamma}{(1+\alpha)\mu_i}$$

as long as

$$\ell \geq \max\left\{8, \sqrt{-8 \log_2 \frac{\alpha}{2\gamma}}, \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i}\mu_i\alpha}}, 4 \log_2 \frac{2^{1/2}}{\gamma}, \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right)\right\}.$$

Note that if $\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} = 3$, then the condition

$$\ell \geq \max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} + 2 \log \left(\max\{-\log_2 \frac{\gamma^\alpha}{4}, 3\} \right)$$

is implied by $\ell \geq 8$. The stated conclusion follows. \square

A.4 Sample Complexity of SafeBAI-Linear

We will define the following:

$$\begin{aligned} \Psi_i^1(x) &:= \log_2 \frac{4(2 + \frac{\theta_{i^*}}{\mu_{i^*}} + \frac{\theta_i}{\mu_i})}{x} \\ \Psi_i^2(x) &:= \max \left\{ \log_2 \frac{4(2 + M_i \theta_i / \gamma)}{x}, \log_2 \frac{4}{\gamma} \right\} \\ \Psi_i^3(x) &:= \log_2 \frac{4}{x} \\ \alpha_i &:= \frac{\gamma}{\mu_i M_i} - 1 \end{aligned}$$

Lemma 9. *On \mathcal{E} , when running Algorithm 1:*

1. If $M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$, then we will have eliminated coordinate i once

$$\ell \geq \max \left\{ \Psi_i^1(\Delta_i), \Phi_i(1), \Phi_{i^*}(1) \right\}.$$

2. If $M_i \in \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$, then we will have eliminated coordinate i once

$$\ell \geq \max \left\{ \Psi_i^2(\Delta_i), \Phi_i(\alpha_i), \Phi_{i^*}(1) \right\} \quad \text{or} \quad \ell \geq \max \left\{ \Psi_i^1 \left(\frac{\gamma\theta_{i^*}}{\mu_{i^*}} - \frac{\gamma\theta_i}{\mu_i} \right), \Phi_i(1), \Phi_{i^*}(1) \right\}.$$

3. If $M_i \notin \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$, then we will have eliminated coordinate i once

$$\ell \geq \max \left\{ \Psi_{i^*}^2(\Delta_i), \Phi_i(1), \Phi_{i^*}(\alpha_{i^*}) \right\} \quad \text{or} \quad \ell \geq \max \left\{ \Psi_i^1 \left(\frac{\gamma\theta_{i^*}}{\mu_{i^*}} - \frac{\gamma\theta_i}{\mu_i} \right), \Phi_i(1), \Phi_{i^*}(1) \right\}.$$

4. If $M_i \in \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$, then we will have eliminated coordinate i once any of the following conditions is met

$$\begin{aligned} \ell \geq \max \left\{ \Psi_i^3(\Delta_i), \Phi_i(\alpha_i), \Phi_{i^*}(\alpha_{i^*}) \right\}, \quad \ell \geq \max \left\{ \Psi_i^1 \left(\frac{\gamma\theta_{i^*}}{\mu_{i^*}} - \frac{\gamma\theta_i}{\mu_i} \right), \Phi_i(1), \Phi_{i^*}(1) \right\} \\ \ell \geq \max \left\{ \Psi_i^2 \left(\frac{\gamma\theta_{i^*}}{\mu_{i^*}} - M_i\theta_i \right), \Phi_i(\alpha_i), \Phi_{i^*}(1) \right\}, \quad \ell \geq \max \left\{ \Psi_{i^*}^2 \left(M_{i^*}\theta_{i^*} - \frac{\gamma\theta_i}{\mu_i} \right), \Phi_i(1), \Phi_{i^*}(\alpha_{i^*}) \right\}. \end{aligned}$$

Proof. We prove each case individually.

Case 1 ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$). By Lemma 7, on \mathcal{E} ,

$$\widehat{a}_s^\ell(i) \geq \max \left\{ \frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}}, a_{0,i} \right\}$$

since by Lemma 2 we know that $\widehat{a}_s^\ell(i) \in \mathcal{S}_i$, and $M_i \notin \mathcal{S}_i$ by assumption, so we are in the first case given in Lemma 7. By Lemma 8, we can therefore lower bound

$$\widehat{a}_s^\ell(i) \geq \frac{\gamma^\ell 2^{\ell(\ell+1)/2}}{\mu_i \sum_{k=1}^{\ell} 2^{\ell-k} \gamma^{k-1} 2^{k(k+1)/2} + \frac{2^\ell}{a_{0,i}}} \geq \frac{\gamma}{2\mu_i}$$

as long as $\ell \geq \Phi_i(1)$, and similarly

$$\widehat{a}_s^\ell(i^*) \geq \frac{\gamma}{2\mu_{i^*}}$$

as long as $\ell \geq \Phi_{i^*}(1)$.

By Lemma 4, we will then have eliminated coordinate i once

$$\begin{aligned} \frac{\theta_i \frac{\gamma}{2\mu_i} + 2^{1-\ell}}{\mu_i \frac{\gamma}{2\mu_i} - 2^{1-\ell}} \leq \frac{\theta_{i^*} \frac{\gamma}{2\mu_{i^*}} - 2^{1-\ell}}{\mu_{i^*} \frac{\gamma}{2\mu_{i^*}} + 2^{1-\ell}} &\iff 2 \cdot 2^{1-\ell} \left(2 + \frac{\theta_{i^*}}{\mu_{i^*}} + \frac{\theta_i}{\mu_i} \right) \leq \gamma \left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i} \right) \\ &\iff 2^{-\ell} \leq \frac{\Delta_i}{4 \left(2 + \frac{\theta_{i^*}}{\mu_{i^*}} + \frac{\theta_i}{\mu_i} \right)} \end{aligned}$$

where the last expression follows since we have assumed $M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$, so $\Delta_i = \gamma \left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i} \right)$.

Case 2 ($M_i \in \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$). Note that $M_i = \frac{\gamma}{(1+\alpha_i)\mu_i}$ for $\alpha_i = \frac{\gamma}{\mu_i M_i} - 1$. Since $M_i \in \mathcal{S}_i$ by assumption, we will have $\alpha_i \geq 0$. By Lemma 7 and Lemma 8, we will have that $\widehat{a}_s^\ell(i) = M_i$ once $\ell \geq \Phi_i(\alpha_i)$. Also assume that $\ell \geq \Phi_{i^*}(1)$ so that $\widehat{a}_s^\ell(i^*) \geq \gamma/(2\mu_{i^*})$. We can then apply Lemma 5 to get that we will eliminate coordinate i once:

$$M_i\theta_i + 2^{1-\ell} < \gamma \cdot \frac{\theta_{i^*} \frac{\gamma}{2\mu_{i^*}} - 2^{1-\ell}}{\mu_{i^*} \frac{\gamma}{2\mu_{i^*}} + 2^{1-\ell}} \iff 2^{2-\ell} \left(\frac{3}{2} + \frac{M_i\theta_i}{\gamma} + \frac{2^{1-\ell}}{\gamma} \right) < \left(\gamma \frac{\theta_{i^*}}{\mu_{i^*}} - M_i\theta_i \right) = \Delta_i$$

A sufficient condition to meet this is that

$$\ell \geq \max \left\{ \log_2 \frac{4(2 + M_i\theta_i/\gamma)}{\Delta_i}, \log_2 \frac{4}{\gamma} \right\}.$$

While we are guaranteed to terminate for this choice of ℓ , it is possible that we will also terminate and eliminate coordinate i for an ℓ value where we have not yet guaranteed that M_i is safe. In that case we can revert to the first case, with the only difference being that $\gamma(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}) \neq \Delta_i$, so we replace the Δ_i in the complexity derived in the first case by $\gamma(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i})$. Nonetheless, our sufficient condition above still works.

Case 3 ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$). This case is nearly identical to the previous case. To guarantee that $\hat{a}_s^\ell(i^*) = M_{i^*}$, we can take $\ell \geq \Phi_{i^*}(\alpha_{i^*})$, where α_{i^*} is defined as in the previous case. Assume also that $\ell \geq \Phi_i(1)$ so that we can guarantee $\hat{a}_s^\ell(i) \geq \gamma/(2\mu_i)$. By Lemma 6 we will then eliminate coordinate i once

$$\gamma \cdot \frac{\frac{\theta_i}{2\mu_i} + 2^{1-\ell}}{\frac{\gamma}{2\mu_i} - 2^{1-\ell}} < M_{i^*}\theta_{i^*} - 2^{1-\ell} \iff 2^{2-\ell} \left(\frac{3}{2} + \frac{M_{i^*}\theta_{i^*}}{\gamma} - \frac{2^{1-\ell}}{\gamma} \right) < M_{i^*}\theta_{i^*} - \gamma \frac{\theta_i}{\mu_i} = \Delta_i$$

and a sufficient condition to meet this is

$$\ell \geq \log_2 \frac{4(3/2 + M_{i^*}\theta_{i^*}/\gamma)}{\Delta_i}.$$

As in the previous case, we may reach the termination criteria before identifying that M_{i^*} is safe, so we can also derive a sufficient condition analogous to case 1. This time, we have $\Delta_i < \gamma(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i})$, so the case 1 condition holds naturally without replacing Δ_i .

Case 4 ($M_i \in \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$). In this case, if ℓ is large enough that $\ell \geq \Phi_{i^*}(\alpha_{i^*})$ and $\ell \geq \Phi_i(\alpha_i)$ hold, on \mathcal{E} , a sufficient condition to eliminate i is that

$$M_i\theta_i + 2^{1-\ell} < M_{i^*}\theta_{i^*} - 2^{1-\ell} \iff 2^{-\ell} < 4(M_{i^*}\theta_{i^*} - M_i\theta_i) = \Delta_i.$$

The result then follows from this and by noting that we might terminate earlier than this by meeting any of the conditions in Case 1-3. \square

Lemma 10. *With probability at least $1 - \delta$, Algorithm 1 will output i^* , only pull safe arms, and terminate after collecting at most*

$$8\sigma^2 \sum_{i \neq i^*} \log \frac{8d\mathcal{L}_i^2}{\delta} \cdot 2^{2\mathcal{L}_i} + 2 \sum_{i \neq i^*} \mathcal{L}_i$$

samples, where \mathcal{L}_i denotes the minimum feasible ℓ value given in Lemma 9 for the case coordinate i falls in.

Proof of Lemma 10. Lemma 1 implies that $\mathbb{P}[\mathcal{E}] \geq 1 - \delta$. We will assume that \mathcal{E} holds for the remainder of the proof.

By Lemma 3, on \mathcal{E} , $i^* \in \mathcal{X}_\ell$ for all ℓ . Since Algorithm 1 will only terminate when $|\mathcal{X}_\ell| = 1$, we will return i^* . By Lemma 2, we will only pull safe arms on \mathcal{E} .

We can then bound the number of samples taken by arm i using Lemma 9. Every i with i^* will fall into one of the four cases given in Lemma 9. We let \mathcal{L}_i denote the minimum sufficient ℓ value in the relevant case for coordinate i . Then we can bound the sample complexity as (since we will pull i^* as many times as the next best arm)

$$2 \sum_{i \neq i^*} \sum_{\ell=1}^{\mathcal{L}_i} (2\sigma^2 \log \frac{8d\ell^2}{\delta} 2^{2\ell} + 1) \leq 8\sigma^2 \sum_{i \neq i^*} \log \frac{8d\mathcal{L}_i^2}{\delta} \cdot 2^{2\mathcal{L}_i} + 2 \sum_{i \neq i^*} \mathcal{L}_i.$$

\square

Lemma 11. *We can upper bound*

$$2^{2\Phi_i(\alpha)} \leq \tilde{\mathcal{O}} \left(\max \left\{ 2^{16}, \xi_{\sqrt{32}}(2\gamma)\xi_{\sqrt{32}}(1/\alpha), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right)\xi_4(2\gamma)\xi_4(1/\alpha), \frac{16}{\gamma^8}, \frac{16}{\gamma^2\alpha^2} \right\} \right).$$

Proof. Recall that

$$\Phi_i(\alpha) = \max \left\{ 8, \sqrt{8 \log_2 \frac{2\gamma}{\alpha}}, \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i}\mu_i\alpha}}, 4 \log_2 \frac{2^{1/2}}{\gamma}, \log_2 \frac{4}{\gamma\alpha} + 2 \log_2 \left(\log_2 \frac{4}{\gamma\alpha} \right) \right\}.$$

Using that $\sqrt{a+b} \leq \sqrt{a} + \sqrt{b}$, we can then upper bound

$$\begin{aligned} 2^2 \sqrt{8 \log_2 \frac{2\gamma}{\alpha}} &\leq 2^2 \sqrt{8 \log_2 2\gamma} 2^2 \sqrt{8 \log_2 1/\alpha} \leq \xi_{\sqrt{32}}(2\gamma) \xi_{\sqrt{32}}(1/\alpha) \\ 2^2 \sqrt{4 \log_2 \frac{2\gamma}{a_{0,i}\mu_i\alpha}} &\leq 2^4 \sqrt{\log_2 \frac{1}{a_{0,i}\mu_i}} 2^4 \sqrt{\log_2(2\gamma)} 2^4 \sqrt{\log_2(1/\alpha)} \leq \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \xi_4(2\gamma) \xi_4(1/\alpha) \\ 2^{2 \cdot 4 \log_2 \frac{2^{3/2}}{\gamma}} &= \frac{16}{\gamma^8} \\ 2^{2(\log_2 \frac{4}{\gamma\alpha} + 2 \log_2(\log_2 \frac{4}{\gamma\alpha}))} &\leq \tilde{\mathcal{O}}\left(\frac{16}{\gamma^2 \alpha^2}\right). \end{aligned}$$

The result follows from these. □

Theorem 4 (Full version of Theorem 2). *We define the following values:*

1. *Case 1* ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$):

$$N_{1,i} := \frac{1 + \theta_{i^*}^2/\mu_{i^*}^2 + \theta_i^2/\mu_i^2}{\Delta_i^2} + \max \left\{ \frac{1}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \xi_4(2\gamma), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \xi_4(2\gamma) \right\}.$$

2. *Case 2* ($M_i \in \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$):

$$\begin{aligned} N_{2,i} := \frac{1 + M_i^2 \theta_i^2 / \gamma^2}{\Delta_i^2} + \max \left\{ \frac{1}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_{\sqrt{32}}(2\gamma) \xi_{\sqrt{32}}\left(\frac{M_i \mu_i}{\gamma - M_i \mu_i}\right), \right. \\ \left. \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \xi_4(2\gamma) \xi_4\left(\frac{M_i \mu_i}{\gamma - M_i \mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \xi_4(2\gamma), \frac{M_i^2 \mu_i^2}{\gamma^2 (\gamma - M_i \mu_i)^2} \right\}. \end{aligned}$$

3. *Case 3* ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$):

$$\begin{aligned} N_{3,i} := \frac{1 + M_{i^*}^2 \theta_{i^*}^2 / \gamma^2}{\Delta_i^2} + \max \left\{ \frac{1}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_{\sqrt{32}}(2\gamma) \xi_{\sqrt{32}}\left(\frac{M_{i^*} \mu_{i^*}}{\gamma - M_{i^*} \mu_{i^*}}\right), \right. \\ \left. \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \xi_4(2\gamma) \xi_4\left(\frac{M_{i^*} \mu_{i^*}}{\gamma - M_{i^*} \mu_{i^*}}\right), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \xi_4(2\gamma), \frac{16 M_{i^*}^2 \mu_{i^*}^2}{\gamma^2 (\gamma - M_{i^*} \mu_{i^*})^2} \right\}. \end{aligned}$$

4. *Case 4* ($M_i \in \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$):

$$\begin{aligned} N_{4,i} := \frac{1}{\Delta_i^2} + \max \left\{ \frac{1}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma) \xi_{\sqrt{32}}\left(\frac{M_i \mu_i}{\gamma - M_i \mu_i}\right), \xi_{\sqrt{32}}(2\gamma) \xi_{\sqrt{32}}\left(\frac{M_{i^*} \mu_{i^*}}{\gamma - M_{i^*} \mu_{i^*}}\right), \right. \\ \left. \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right) \xi_4(2\gamma) \xi_4\left(\frac{M_i \mu_i}{\gamma - M_i \mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right) \xi_4(2\gamma) \xi_4\left(\frac{M_{i^*} \mu_{i^*}}{\gamma - M_{i^*} \mu_{i^*}}\right), \right. \\ \left. \frac{M_{i^*}^2 \mu_{i^*}^2}{\gamma^2 (\gamma - M_{i^*} \mu_{i^*})^2}, \frac{M_i^2 \mu_i^2}{\gamma^2 (\gamma - M_i \mu_i)^2} \right\}. \end{aligned}$$

Let $c(i)$ denote the case coordinate i falls in. Then, with probability at least $1 - \delta$, Algorithm 1 will output i^* , only pull safe arms, and terminate after collecting at most

$$\tilde{\mathcal{O}}\left(\log \frac{d}{\delta} \cdot \sum_{i \neq i^*} N_{c(i),i}\right)$$

samples.

Proof. We apply Lemma 11 to simplify each of the cases. While for cases 2-4 there exist multiple termination criteria, as we are concerned with interpretability, we consider only the first.

Case 1 ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$). Here we need $\ell \geq \max\{\Psi_i^1(\Delta_i), \Phi_i(1), \Phi_{i^*}(a)\}$. By Lemma 11, and noting that $\xi_a(1) = 1$, we can upper bound

$$\begin{aligned} & \max\{2^{2\Psi_i^1(\Delta_i)}, 2^{2\Phi_i(1)}, 2^{2\Phi_{i^*}(1)}\} \\ & \leq \tilde{\mathcal{O}}\left(\max\left\{\frac{1 + \theta_{i^*}^2/\mu_{i^*}^2 + \theta_i^2/\mu_i^2}{\Delta_i^2}, 2^{16}, \frac{16}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right)\xi_4(2\gamma), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right)\xi_4(2\gamma)\right\}\right). \end{aligned}$$

Case 2 ($M_i \in \mathcal{S}_i$ and $M_{i^*} \notin \mathcal{S}_{i^*}$). We now need $\ell \geq \max\{\Psi_i^2(\Delta_i), \Phi_i(\alpha_i), \Phi_{i^*}(1)\}$. Applying the same simplifications, we can bound

$$\begin{aligned} & \max\{2^{2\Psi_i^2(\Delta_i)}, 2^{2\Phi_i(\alpha_i)}, 2^{2\Phi_{i^*}(1)}\} \\ & \leq \tilde{\mathcal{O}}\left(\max\left\{\frac{1 + M_i^2\theta_i^2/\gamma^2}{\Delta_i^2}, 2^{16}, \frac{16}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_{\sqrt{32}}(2\gamma)\xi_{\sqrt{32}}\left(\frac{M_i\mu_i}{\gamma - M_i\mu_i}\right), \right. \right. \\ & \quad \left. \left. \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right)\xi_4(2\gamma)\xi_4\left(\frac{M_i\mu_i}{\gamma - M_i\mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right)\xi_4(2\gamma), \frac{16M_i^2\mu_i^2}{\gamma^2(\gamma - M_i\mu_i)^2}\right\}\right). \end{aligned}$$

Case 3 ($M_i \notin \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$). We now need $\ell \geq \max\{\Psi_{i^*}^2(\Delta_i), \Phi_i(1), \Phi_{i^*}(\alpha_{i^*})\}$. We can bound

$$\begin{aligned} & \max\{2^{2\Psi_{i^*}^2(\Delta_i)}, 2^{2\Phi_i(1)}, 2^{2\Phi_{i^*}(\alpha_{i^*})}\} \\ & \leq \tilde{\mathcal{O}}\left(\max\left\{\frac{1 + M_{i^*}^2\theta_{i^*}^2/\gamma^2}{\Delta_i^2}, 2^{16}, \frac{16}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma), \xi_{\sqrt{32}}(2\gamma)\xi_{\sqrt{32}}\left(\frac{M_{i^*}\mu_{i^*}}{\gamma - M_{i^*}\mu_{i^*}}\right), \right. \right. \\ & \quad \left. \left. \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right)\xi_4(2\gamma)\xi_4\left(\frac{M_{i^*}\mu_{i^*}}{\gamma - M_{i^*}\mu_{i^*}}\right), \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right)\xi_4(2\gamma), \frac{16M_{i^*}^2\mu_{i^*}^2}{\gamma^2(\gamma - M_{i^*}\mu_{i^*})^2}\right\}\right). \end{aligned}$$

Case 4 ($M_i \in \mathcal{S}_i$ and $M_{i^*} \in \mathcal{S}_{i^*}$). In the final case, we need $\ell \geq \max\{\Psi_i^3(\Delta_i), \Phi_i(\alpha_i), \Phi_{i^*}(\alpha_{i^*})\}$. We can then upper bound

$$\begin{aligned} & \max\{2^{2\Psi_i^3(\Delta_i)}, 2^{2\Phi_i(\alpha_i)}, 2^{2\Phi_{i^*}(\alpha_{i^*})}\} \\ & \leq \tilde{\mathcal{O}}\left(\max\left\{\frac{1}{\Delta_i^2}, 2^{16}, \frac{16}{\gamma^8}, \xi_{\sqrt{32}}(2\gamma)\xi_{\sqrt{32}}\left(\frac{M_i\mu_i}{\gamma - M_i\mu_i}\right), \xi_{\sqrt{32}}(2\gamma)\xi_{\sqrt{32}}\left(\frac{M_{i^*}\mu_{i^*}}{\gamma - M_{i^*}\mu_{i^*}}\right), \right. \right. \\ & \quad \xi_4\left(\frac{1}{a_{0,i}\mu_i}\right)\xi_4(2\gamma)\xi_4\left(\frac{M_i\mu_i}{\gamma - M_i\mu_i}\right), \xi_4\left(\frac{1}{a_{0,i^*}\mu_{i^*}}\right)\xi_4(2\gamma)\xi_4\left(\frac{M_{i^*}\mu_{i^*}}{\gamma - M_{i^*}\mu_{i^*}}\right), \\ & \quad \left. \frac{16M_{i^*}^2\mu_{i^*}^2}{\gamma^2(\gamma - M_{i^*}\mu_{i^*})^2}, \frac{16M_i^2\mu_i^2}{\gamma^2(\gamma - M_i\mu_i)^2}\right\}\right). \end{aligned}$$

Obtaining a final complexity. To obtain the final result, we apply Lemma 10 with the upper bounds given here. For the simplified case presented in the main text, we note that $\xi_a(\gamma)$ can be thought of as an absolute constant for a and γ not too large. \square

Proof of Proposition 1. Let $\tilde{x} := \max\{2, x\}$. Note that, for $z > 0$,

$$2^{a\sqrt{\log_2 \tilde{x}}} \leq \tilde{x}^z \iff a\sqrt{\log_2 \tilde{x}} \leq z \log_2 \tilde{x} \iff \frac{a}{z} \leq \sqrt{\log_2 \tilde{x}} \iff 2^{a^2/z^2} \leq \tilde{x}.$$

As $2^{a\sqrt{\log_2 \tilde{x}}}$ is increasing in \tilde{x} , we then have

$$2^{a\sqrt{\log_2 \tilde{x}}} \leq \begin{cases} 2^{a\sqrt{\log_2 \tilde{x}}} & \tilde{x} < 2^{a^2/z^2} \\ \tilde{x}^z & 2^{a^2/z^2} \leq \tilde{x} \end{cases} \leq \begin{cases} 2^{a\sqrt{\log_2 2^{a^2/z^2}}} & \tilde{x} < 2^{a^2/z^2} \\ \tilde{x}^z & 2^{a^2/z^2} \leq \tilde{x} \end{cases} \leq \tilde{x}^z + 2^{a\sqrt{\log_2 2^{a^2/z^2}}} = \tilde{x}^z + 2^{a^2/z}.$$

As this holds for all $z > 0$, we can take the minimum over z . \square

B Lower Bound Proof

Throughout this section we assume that $\sigma^2 = 1$.

Let ν denote a distribution with

$$\nu_i = \mathcal{N}\left(\begin{bmatrix} a_i \theta_i \\ a_i \mu_i \end{bmatrix}, I\right)$$

for some $\theta_i > 0, \mu_i > 0$, and let

$$i^* = \arg \max_i \theta_i / \mu_i.$$

We assume that i^* is the unique maximum of θ_i / μ_i .

Lemma 12. *Let τ denote the stopping time for any δ -PAC algorithm. Then,*

$$\mathbb{E}_\nu[\tau] \geq \sum_{i \neq i^*} \frac{\frac{2}{3a_i^2} \left(\frac{1}{\mu_i^2} + \left(\frac{\theta_{i^*}}{\mu_i \mu_{i^*}} \right)^2 + \frac{\theta_i^2}{\mu_i^4} \right)}{\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i} \right)^2} \cdot \log \frac{1}{2.4\delta}.$$

Proof. We rely on the change-of-measure technique of Kaufmann et al. (2016). Note that while Kaufmann et al. (2016) assumes the distributions are 1-dimensional, their proof goes through identically for 2-dimensional distributions. In particular, their Lemma 1 holds as

$$\sum_{i=1}^d \mathbb{E}_\nu[N_i(\tau)] \text{KL}(\nu_i, \nu'_i) \geq \sup_{\mathcal{E} \in \mathcal{F}_\tau} d(\mathbb{P}_\nu(\mathcal{E}), \mathbb{P}_{\nu'}(\mathcal{E})).$$

where here ν_i, ν'_i are 2-dimensional distributions, $N_i(\tau)$ is the number of pulls of arm i up to stopping time τ , and \mathcal{F}_τ is the filtration up to τ .

To prove the result, for every coordinate i we construct three alternate instances—one which involves perturbing θ_i , and two which involve perturbing μ_i . We then apply the above inequality with each alternate instance to show a lower bound on the number of times we must pull arm i .

Alternate Instance 1. Fix $i \neq i^*$ and consider the alternate instance ν' where $\nu'_j = \nu_j$ for $j \neq i$, and

$$\nu_i = \mathcal{N}\left(\begin{bmatrix} a_i \theta'_i \\ a_i \mu_i \end{bmatrix}, I\right)$$

for $\theta'_i = \mu_i \theta_{i^*} / \mu_{i^*} + \alpha, \alpha > 0$. Let $\mathcal{E} = \{\widehat{i} = i^*\}$ (where here i^* denotes the optimal arm on ν). Note that on ν' , we will have that $\theta'_i / \mu_i = \theta_{i^*} / \mu_{i^*} + \alpha / \mu_i > \theta_{i^*} / \mu_{i^*}$. It follows that arm i is the optimal arm under instance ν' , so for any δ -PAC algorithm, $\mathbb{P}_\nu[\mathcal{E}] \geq 1 - \delta, \mathbb{P}_{\nu'}[\mathcal{E}] \leq \delta$. In addition, we have that

$$\text{KL}(\nu_i, \nu'_i) = \frac{a_i^2}{2} (\theta_i - \theta'_i)^2 = \frac{a_i^2}{2} (\theta_i - \mu_i \theta_{i^*} / \mu_{i^*} - \alpha)^2 = \frac{a_i^2 \mu_i^2}{2} (\theta_i / \mu_i - \theta_{i^*} / \mu_{i^*} - \alpha / \mu_i)^2.$$

Lower bounding $d(\mathbb{P}_\nu(\mathcal{E}), \mathbb{P}_{\nu'}(\mathcal{E})) \geq \log \frac{1}{2.4\delta}$ as in Kaufmann et al. (2016) and letting $\alpha \rightarrow 0$, we conclude that

$$\mathbb{E}_\nu[N_i(\tau)] \geq \frac{2}{a_i^2 \mu_i^2 \left(\frac{\theta_i}{\mu_i} - \frac{\theta_{i^*}}{\mu_{i^*}} \right)^2} \cdot \log \frac{1}{2.4\delta}.$$

Alternate Instance 2. Next, consider the alternate instance ν' where $\mu'_i = \mu_{i^*} \theta_i / \theta_{i^*} - \alpha$, and leaving all other parameters the same as ν . We now have that

$$\frac{\theta_i}{\mu'_i} = \frac{\theta_i}{\frac{\mu_{i^*} \theta_i}{\theta_{i^*}} - \alpha} = \frac{1}{\frac{\mu_{i^*}}{\theta_{i^*}} - \frac{\alpha}{\theta_i}} > \frac{\theta_{i^*}}{\mu_{i^*}}$$

where the last inequality holds for small enough α . Thus, i is the optimal coordinate for ν' . Using that

$$\text{KL}(\nu_i, \nu'_i) = \frac{a_i^2}{2}(\mu_i - \mu'_i)^2 = \frac{a_i^2}{2}\left(\mu_i - \frac{\mu_{i^*}\theta_i}{\theta_{i^*}} + \alpha\right)^2 = \frac{a_i^2}{2}\left(\frac{\mu_i\mu_{i^*}}{\theta_{i^*}}\right)^2\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i} + \frac{\theta_{i^*}\alpha}{\mu_i\mu_{i^*}}\right)^2.$$

Letting α tend to 0 and using the same event \mathcal{E} , we conclude that

$$\mathbb{E}_\nu[N_i(\tau)] \geq \frac{2}{a_i^2\left(\frac{\mu_i\mu_{i^*}}{\theta_{i^*}}\right)^2\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} \cdot \log \frac{1}{2.4\delta}.$$

Alternate Instance 3. We next consider the alternate instance ν' where $\mu'_i = \mu_i - \frac{\mu_i^2}{\theta_i}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)$ and all other parameters the same as ν . Now,

$$\frac{\mu'_i}{\theta_i} = \frac{\mu_i}{\theta_i} - \frac{\mu_i^2}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right) = \frac{\mu_{i^*}}{\theta_{i^*}} + \left(\frac{\mu_i}{\theta_i} - \frac{\mu_{i^*}}{\theta_{i^*}}\right) - \frac{\mu_i^2}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right).$$

Note that

$$\frac{\mu_i}{\theta_i} - \frac{\mu_{i^*}}{\theta_{i^*}} = \frac{\mu_i\mu_{i^*}}{\theta_i\theta_{i^*}}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right) < \frac{\mu_i^2}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)$$

where the inequality follows since i^* is the unique maximizer of θ_i/μ_i , so $(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}) > 0$ and $\mu_i/\theta_i > \mu_{i^*}/\theta_{i^*}$. Thus,

$$\frac{\mu_{i^*}}{\theta_{i^*}} + \left(\frac{\mu_i}{\theta_i} - \frac{\mu_{i^*}}{\theta_{i^*}}\right) - \frac{\mu_i^2}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right) < \frac{\mu_{i^*}}{\theta_{i^*}}$$

so $\theta_i/\mu'_i > \mu_{i^*}/\theta_{i^*}$, and i is therefore the optimal coordinate on instance ν' . Note also that

$$\text{KL}(\nu_i, \nu'_i) = \frac{a_i^2}{2}(\mu_i - \mu'_i)^2 = \frac{a_i^2}{2}\left(\frac{\mu_i^2}{\theta_i}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)\right)^2 = \frac{a_i^2\mu_i^4}{2\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2.$$

Using the same event \mathcal{E} as above, we conclude that

$$\mathbb{E}_\nu[N_i(\tau)] \geq \frac{2}{\frac{a_i^2\mu_i^4}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} \cdot \log \frac{1}{2.4\delta}.$$

Full Lower Bound. Putting these three lower bound together, it follows that

$$\begin{aligned} \mathbb{E}_\nu[N_i(\tau)] &\geq \max \left\{ \frac{2}{a_i^2\mu_i^2\left(\frac{\theta_i}{\mu_i} - \frac{\theta_{i^*}}{\mu_{i^*}}\right)^2}, \frac{2}{a_i^2\left(\frac{\mu_i\mu_{i^*}}{\theta_{i^*}}\right)^2\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2}, \frac{2}{\frac{a_i^2\mu_i^4}{\theta_i^2}\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} \right\} \cdot \log \frac{1}{2.4\delta} \\ &\geq \frac{2}{3a_i^2\left(\frac{1}{\mu_i^2} + \left(\frac{\theta_{i^*}}{\mu_i\mu_{i^*}}\right)^2 + \frac{\theta_i^2}{\mu_i^4}\right)} \cdot \log \frac{1}{2.4\delta}. \end{aligned}$$

Repeating this argument for each $i \neq i^*$ and noting that $\tau = \sum_{i=1}^d N_i(\tau)$ gives the result. \square

Proof of Theorem 1. For any coordinate i , note that any $(\delta, \{\mathcal{X}_i\}_{i \in [d]})$ -PAC algorithm can pull all arms $a_i \leq \gamma/\mu_i$. Assume that we pull $a_i = \gamma/\mu_i$ for all time, then the observation model and hypothesis test are identical to that used in Lemma 12. It therefore follows from Lemma 12 that

$$\mathbb{E}_\nu[\tau] \geq \sum_{i \neq i^*} \frac{\frac{2}{3a_i^2}\left(\frac{1}{\mu_i^2} + \left(\frac{\theta_{i^*}}{\mu_i\mu_{i^*}}\right)^2 + \frac{\theta_i^2}{\mu_i^4}\right)}{\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} \cdot \log \frac{1}{2.4\delta}.$$

Plugging in $a_i = \gamma/\mu_i$ (note that this value minimizes the above, subject to the safety constraint) gives

$$\frac{\frac{2}{3a_i^2}\left(\frac{1}{\mu_i^2} + \left(\frac{\theta_{i^*}}{\mu_i\mu_{i^*}}\right)^2 + \frac{\theta_i^2}{\mu_i^4}\right)}{\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} = \frac{\frac{2}{3}\mu_i^2\left(\frac{1}{\mu_i^2} + \left(\frac{\theta_{i^*}}{\mu_i\mu_{i^*}}\right)^2 + \frac{\theta_i^2}{\mu_i^4}\right)}{\gamma^2\left(\frac{\theta_{i^*}}{\mu_{i^*}} - \frac{\theta_i}{\mu_i}\right)^2} = \frac{\frac{2}{3}\left(1 + \frac{\theta_{i^*}^2}{\mu_{i^*}^2} + \frac{\theta_i^2}{\mu_i^2}\right)}{\Delta_i^2}$$

which proves the result. \square

C Monotonic Response Functions Proof

The following is a useful fact that we will employ in several places.

Proposition 2. *Assume that f is nondecreasing and is 1-Lipschitz. Then for $x \geq y$, we will have*

$$y - f(y) \leq x - f(x).$$

Proof. This trivially follows from the assumptions on f . Since the function is 1-Lipschitz, we have $|f(x) - f(y)| \leq |x - y|$, and since the function is nondecreasing and $x \geq y$, we have $|f(x) - f(y)| = f(x) - f(y)$ and $|x - y| = x - y$. Rearranging these expressions gives the result. \square

Notation. Let $\hat{n}_{i,\ell}$ denote the first value of $n_i - 1$ in epoch ℓ when $\gamma - \hat{g}_i(\hat{a}_s^{n_i-1,\ell}(i)) \leq 2\epsilon_\ell$. Similarly, while $\text{unsafe}(i) = 0$, let $\hat{m}_{i,\ell}$ denote the first value of $m_i - 1$ in epoch ℓ when $\gamma + \epsilon_{\text{safe}} - \hat{g}_i(\hat{a}_u^{m_i-1,\ell}(i)) \leq 2\epsilon_\ell$, and when $\text{unsafe}(i) = 1$, let $\hat{m}_{i,\ell}$ denote the first value of $m_i - 1$ for which the condition on Line 24 or Line 26 is met. Let $\hat{\ell}_{\text{unsafe}}(i)$ denote the epoch on which we set $\text{unsafe}(i) = 1$.

Lemma 13. *Let \mathcal{E}_{fun} be the event on which, for all ℓ , $i \in \mathcal{X}_\ell$, n_i , and m_i ,*

$$\begin{aligned} |\hat{g}_i(\hat{a}_s^{n_i,\ell}(i)) - g_i(\hat{a}_s^{n_i,\ell}(i))| &\leq \epsilon_\ell, & |\hat{f}_i(\hat{a}_s^{n_i,\ell}(i)) - f_i(\hat{a}_s^{n_i,\ell}(i))| &\leq \epsilon_\ell \\ |\hat{g}_i(\hat{a}_u^{m_i,\ell}(i)) - g_i(\hat{a}_u^{m_i,\ell}(i))| &\leq \epsilon_\ell, & |\hat{f}_i(\hat{a}_u^{m_i,\ell}(i)) - f_i(\hat{a}_u^{m_i,\ell}(i))| &\leq \epsilon_\ell. \end{aligned}$$

Then $\mathbb{P}[\mathcal{E}_{\text{fun}}] \geq 1 - \delta$.

Proof. Our estimate $\hat{g}_i(\hat{a}_s^{n_i,\ell}(i))$ will have mean $g_i(\hat{a}_s^{n_i,\ell}(i))$ and variance $1/N_{\ell,t}$, and similarly $\hat{f}_i(\hat{a}_s^{n_i,\ell}(i))$ will have mean $f_i(\hat{a}_s^{n_i,\ell}(i))$, and variance $1/N_{\ell,t}$. By the concentration of sub-Gaussian random variables, it follows that, with probability at least $1 - \delta/(2t^2)$,

$$|\hat{g}_i(\hat{a}_s^{n_i,\ell}(i)) - g_i(\hat{a}_s^{n_i,\ell}(i))| \leq \sqrt{\frac{2\sigma^2 \log \frac{8t^2}{\delta}}{N_{\ell,t}}} = \epsilon_\ell, \quad |\hat{f}_i(\hat{a}_s^{n_i,\ell}(i)) - f_i(\hat{a}_s^{n_i,\ell}(i))| \leq \sqrt{\frac{2\sigma^2 \log \frac{8t^2}{\delta}}{N_{\ell,t}}} = \epsilon_\ell.$$

The same is true of $\hat{g}_i(\hat{a}_u^{m_i,\ell}(i))$ and $\hat{f}_i(\hat{a}_u^{m_i,\ell}(i))$. As we increment t each time we collect samples, we can union bound over the total number of times samples are collected, and will get that the above events hold every time with probability at least

$$1 - \sum_{t=1}^{\infty} \frac{\delta}{2t^2} = 1 - \frac{\delta}{2} \cdot \frac{\pi^2}{6} \geq 1 - \delta.$$

\square

C.1 Controlling Safe and Unsafe Value Estimates

Lemma 14. *For all k , ℓ , and i , we will have that $\hat{a}_s^{k,\ell}(i) \geq a_{0,i}$ and $\hat{a}_u^{k,\ell}(i) \geq a_{0,i}$.*

Proof. We initialize both $\hat{a}_s^{0,0}(i) = \hat{a}_u^{0,0}(i) = a_{0,i}$. The only point at which we change $\hat{a}_s^{k,\ell}(i)$ is Line 10, and we update it as

$$\hat{a}_s^{n_i,\ell}(i) \leftarrow \gamma + \hat{a}_s^{n_i-1,\ell}(i) - \hat{g}_i(\hat{a}_s^{n_i-1,\ell}(i)) - \epsilon_\ell.$$

Note that, by the condition of the preceding while loop, to run this line we must have that $\gamma - \hat{g}_i(\hat{a}_s^{n_i-1,\ell}(i)) > 2\epsilon_\ell$. It follows that

$$\begin{aligned} \hat{a}_s^{n_i,\ell}(i) &= \gamma + \hat{a}_s^{n_i-1,\ell}(i) - \hat{g}_i(\hat{a}_s^{n_i-1,\ell}(i)) - \epsilon_\ell \\ &\geq \hat{a}_s^{n_i-1,\ell}(i) + \epsilon_\ell \\ &\geq \hat{a}_s^{n_i-1,\ell}(i). \end{aligned}$$

Thus, we only increase the value of $\hat{a}_s^{n_i-1,\ell}(i)$, so since $\hat{a}_s^{0,0}(i) = a_{0,i}$, the result follows. A similar argument shows that while $\text{unsafe}(i) = 0$, we will also have that $\hat{a}_u^{m_i,\ell}(i) \geq a_{0,i}$. Once $\text{unsafe}(i) = 1$, note that we will always have that $\hat{a}_u^{m_i,\ell}(i) \geq \hat{a}_s^{n_i-1,\ell}(i)$, so it follows here that $\hat{a}_u^{m_i,\ell}(i) \geq a_{0,i}$ as well. \square

Lemma 15. For all k, ℓ and i , on \mathcal{E}_{fun} we will have that $\widehat{a}_s^{k,\ell}(i) \leq g_i^{-1}(\gamma)$ and $\widehat{a}_u^{k,\ell}(i) \leq g_i^{-1}(\gamma + \epsilon_{\text{safe}})$.

Proof. We prove this result by induction. By assumption $\widehat{a}_s^{0,0}(i) \leq g_i^{-1}(\gamma)$. Assume that $\widehat{a}_s^{k-1,\ell}(i) \leq g_i^{-1}(\gamma)$, then, on \mathcal{E}_{fun} ,

$$\begin{aligned} \widehat{a}_s^{k,\ell}(i) &= \gamma + \widehat{a}_s^{k-1,\ell}(i) - \widehat{g}_i(\widehat{a}_s^{k-1,\ell}(i)) - \epsilon_\ell \\ &\leq \gamma + \widehat{a}_s^{k-1,\ell}(i) - g_i(\widehat{a}_s^{k-1,\ell}(i)) \\ &\leq \gamma + g_i^{-1}(\gamma) - g_i(g_i^{-1}(\gamma)) \\ &= g_i^{-1}(\gamma) \end{aligned}$$

where the last inequality follows from Proposition 2. As we always set $\widehat{a}_s^{0,\ell}(i) = \widehat{a}_s^{\widehat{n}_i, \ell-1, \ell-1}(i)$, the upper bound on $\widehat{a}_s^{k,\ell}(i)$ follows.

While $\text{unsafe}(i) = 0$, the update of $\widehat{a}_u^{k,\ell}(i)$ is analogous to the update of $\widehat{a}_s^{k,\ell}(i)$ except we replace γ with $\gamma + \epsilon_{\text{safe}}$. Thus, in this regime an identical argument to the above shows that $\widehat{a}_u^{k,\ell}(i) \leq g_i^{-1}(\gamma + \epsilon_{\text{safe}})$. When $\text{unsafe}(i) = 1$, note that we can only decrease $\widehat{a}_u^{k,\ell}(i)$ so the result also holds there. \square

Lemma 16. On \mathcal{E}_{fun} , $\widehat{a}_s^{\widehat{n}_i, \ell, \ell}(i) \geq \widetilde{g}_i^{-1}(\gamma - 3\epsilon_\ell)$.

Proof. By construction $\gamma - \widehat{g}_i(\widehat{a}_s^{\widehat{n}_i, \ell, \ell}(i)) \leq 2\epsilon_\ell$. This implies, on \mathcal{E}_{fun} ,

$$g_i(\widehat{a}_s^{\widehat{n}_i, \ell, \ell}(i)) \geq \gamma - 3\epsilon_\ell \implies \widehat{a}_s^{\widehat{n}_i, \ell, \ell}(i) \geq \widetilde{g}_i^{-1}(\gamma - 3\epsilon_\ell)$$

where we use \widetilde{g}^{-1} since, by Lemma 14, we know that $\widehat{a}_s^{\widehat{n}_i, \ell, \ell}(i) \geq a_{0,i}$. \square

Lemma 17. On \mathcal{E}_{fun} , while $\text{unsafe}(i) = 0$, $\widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i) \geq \widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_\ell)$.

Proof. Note that when $\text{unsafe}(i) = 0$, the sequence $\widehat{a}_u^{m_i, \ell}(i)$ is updated analogously to $\widehat{a}_s^{n_i, \ell}(i)$, except with γ replaced by $\gamma + \epsilon_{\text{safe}}$. We can therefore use the proof from Lemma 16 to get the result. \square

Lemma 18. On \mathcal{E}_{fun} , for all i and $\ell \geq \widehat{\ell}_{\text{unsafe}}(i)$, we will have that $\widehat{a}_u^{\widehat{m}_i, \ell, \ell} \geq g_i^{-1}(\gamma)$.

Proof. We prove this by induction. Note that to set $\text{unsafe}(i) = 1$, we must have that $\widehat{g}_i(\widehat{a}_u^{m_i, \ell}(i)) - \epsilon_\ell \geq \gamma$. On \mathcal{E}_{fun} , this implies that $g_i(\widehat{a}_u^{m_i, \ell}(i)) \geq \gamma$, so $g_i(\widehat{a}_u^{\widehat{m}_i, \widehat{\ell}_{\text{unsafe}}(i), \widehat{\ell}_{\text{unsafe}}(i)}(i)) \geq \gamma$.

Now assume that for some $\ell \geq \widehat{\ell}_{\text{unsafe}}(i)$, $g_i(\widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i)) \geq \gamma$. At round $\ell + 1$, the while loop on Line 24 either terminates when $\widehat{g}_i(\widehat{a}_u^{m_i-1, \ell+1}(i)) - \epsilon_{\ell+1} \geq \gamma$ and $\widehat{a}_u^{\widehat{m}_i, \ell+1, \ell+1}(i) = \widehat{a}_u^{m_i-1, \ell+1}(i)$, or it terminates when the if statement on Line 26 is true, and we set $\widehat{a}_u^{\widehat{m}_i, \ell+1, \ell+1}(i) = \widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i)$. In the latter case we are done since we have assumed that $g_i(\widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i)) \geq \gamma$. In the former case, on \mathcal{E}_{fun} we will have that $\widehat{g}_i(\widehat{a}_u^{m_i-1, \ell+1}(i)) - \epsilon_{\ell+1} \geq \gamma$ implies $g_i(\widehat{a}_u^{m_i-1, \ell+1}(i)) \geq \gamma$, so $g_i(\widehat{a}_u^{\widehat{m}_i, \ell+1, \ell+1}(i)) \geq \gamma$. This proves the inductive hypothesis so the result follows. \square

Lemma 19. On \mathcal{E}_{fun} , once $\text{unsafe}(i) = 1$, we can bound

$$\widehat{a}_u^{0, \ell+1}(i) = \widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i) \leq \sum_{s=1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}} + \left(\ell + \frac{4g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{\epsilon_{\text{safe}}} \right) 2^{-\ell} =: \bar{a}_u^\ell(i).$$

Proof. Assume that for epoch ℓ , the if statement on Line 26 is never true. Then, by definition, we will have that $\widehat{g}_i(\widehat{a}_u^{\widehat{m}_i, \ell-1, \ell}(i)) - \epsilon_\ell < \gamma$ and $\widehat{g}_i(\widehat{a}_u^{\widehat{m}_i, \ell, \ell}(i)) - \epsilon_\ell \geq \gamma$. It follows that we can bound, on \mathcal{E}_{fun} ,

$$g_i(\widehat{a}_u^{\widehat{m}_i, \ell-1, \ell}(i)) < \gamma + 2\epsilon_\ell.$$

By Lemma 15, we can also bound $g_i(\widehat{a}_u^{\widehat{m}_i, \ell-1, \ell}(i)) \leq \gamma + \epsilon_{\text{safe}}$. Putting this together we get that

$$\widehat{a}_u^{\widehat{m}_i, \ell-1, \ell}(i) \leq g_i^{-1}(\min\{\gamma + 2\epsilon_\ell, \gamma + \epsilon_{\text{safe}}\}).$$

Note that this is well-defined since by assumption $g_i^{-1}(\gamma + \epsilon_{\text{safe}})$ is well-defined. By definition, $\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) = \frac{\widehat{a}_u^{0,\ell}(i)}{2} + \frac{\widehat{a}_u^{\widehat{m}_{i,\ell-1},\ell}(i)}{2}$ so it follows

$$\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) \leq \frac{\widehat{a}_u^{0,\ell}(i)}{2} + \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_\ell, \gamma + \epsilon_{\text{safe}}\})}{2}.$$

If we instead terminate with the if statement on Line 26, we will have that $\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) = \widehat{a}_u^{0,\ell}(i)$. Furthermore, the termination criteria is only met if $\widehat{a}_u^{0,\ell}(i) - \widehat{a}_u^{m_{i,\ell},\ell}(i) \leq \epsilon_\ell$ and $\widehat{g}_i(\widehat{a}_u^{m_{i,\ell},\ell}(i)) - \epsilon_\ell < \gamma$. On \mathcal{E}_{fun} , this implies that $\widehat{a}_u^{m_{i,\ell},\ell}(i) < g_i^{-1}(\min\{\gamma + 2\epsilon_\ell, \gamma + \epsilon_{\text{safe}}\})$, which implies that

$$\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) = \widehat{a}_u^{0,\ell}(i) = \frac{\widehat{a}_u^{0,\ell}(i)}{2} + \frac{\widehat{a}_u^{0,\ell}(i)}{2} \leq \frac{\widehat{a}_u^{0,\ell}(i)}{2} + \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_\ell, \gamma + \epsilon_{\text{safe}}\}) + \epsilon_\ell}{2}.$$

By construction $\widehat{a}_u^{0,\ell+1}(i) = \widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i)$, so we have shown that for all $\ell > \widehat{\ell}_{\text{unsafe}}(i)$,

$$\begin{aligned} \widehat{a}_u^{0,\ell+1}(i) &\leq \frac{\widehat{a}_u^{0,\ell}(i)}{2} + \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_\ell, \gamma + \epsilon_{\text{safe}}\}) + \epsilon_\ell}{2} \\ &\leq \sum_{s=\widehat{\ell}_{\text{unsafe}}(i)+1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\}) + \epsilon_s}{2^{\ell-s+1}} + \frac{\widehat{a}_u^{0,\widehat{\ell}_{\text{unsafe}}(i)}(i)}{2^{\ell-\widehat{\ell}_{\text{unsafe}}(i)+1}}. \end{aligned}$$

We would like to obtain a deterministic bound that does not depend on $\widehat{\ell}_{\text{unsafe}}(i)$. To this end, we bound

$$\begin{aligned} \sum_{s=\widehat{\ell}_{\text{unsafe}}(i)+1}^{\ell} \frac{\epsilon_s}{2^{\ell-s+1}} &\leq \sum_{s=1}^{\ell} \frac{2^{-s}}{2^{\ell-s+1}} = \ell 2^{-\ell}, \\ \sum_{s=\widehat{\ell}_{\text{unsafe}}(i)+1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}} &\leq \sum_{s=1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}}. \end{aligned}$$

Finally, by Lemma 15, we can bound $\widehat{a}_u^{0,\widehat{\ell}_{\text{unsafe}}(i)}(i) \leq g_i^{-1}(\gamma + \epsilon_{\text{safe}})$, and by Lemma 22 we can bound $\widehat{\ell}_{\text{unsafe}}(i) \leq \log_2 \frac{8}{\epsilon_{\text{safe}}}$, to get that

$$\frac{\widehat{a}_u^{0,\widehat{\ell}_{\text{unsafe}}(i)}(i)}{2^{\ell-\widehat{\ell}_{\text{unsafe}}(i)+1}} \leq \frac{g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{2^{\ell-\log_2 \frac{8}{\epsilon_{\text{safe}}}+1}} = \frac{4g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{\epsilon_{\text{safe}} 2^\ell}.$$

Putting these bounds together gives the result. \square

C.2 Bounding Number of Epochs

Lemma 20. *On \mathcal{E}_{fun} , $\widehat{n}_{i,\ell} \leq \bar{n}_{i,\ell}$ where*

$$\bar{n}_{i,\ell} := \begin{cases} \widetilde{g}_i^{-1}(\gamma - \epsilon_\ell) - \widetilde{g}_i^{-1}(\gamma - 6\epsilon_\ell) & \ell \geq 2 \\ 2(\widetilde{g}_i^{-1}(\gamma - \frac{\epsilon_\ell}{2}) - a_{0,i}) & \ell = 1 \end{cases}.$$

Proof. Throughout this proof, we make use of Lemma 14 to note that $\widehat{a}_s^{n_i,\ell}(i) \geq a_{0,i}$ always. This allows us to lower bound terms with a \widetilde{g}^{-1} .

By definition, $\widehat{n}_{i,\ell}$ is the smallest value of $n_i - 1$ such that $\gamma - \widehat{g}_i(\widehat{a}_s^{n_i-1,\ell}(i)) \leq 2\epsilon_\ell$. Assume we are in the regime where $n_i - 1 < \widehat{n}_{i,\ell}$, so that $\gamma - \widehat{g}_i(\widehat{a}_s^{n_i-1,\ell}(i)) > 2\epsilon_\ell$. By construction,

$$\widehat{a}_s^{n_i,\ell}(i) = \gamma + \widehat{a}_s^{n_i-1,\ell}(i) - \widehat{g}_i(\widehat{a}_s^{n_i-1,\ell}(i)) - \epsilon_\ell \geq \widehat{a}_s^{n_i-1,\ell}(i) + \epsilon_\ell.$$

It follows that while $n_i - 1 < \widehat{n}_{i,\ell}$, we can lower bound

$$\widehat{a}_s^{n_i,\ell}(i) \geq \widehat{a}_s^{0,\ell}(i) + n_i \epsilon_\ell.$$

On \mathcal{E}_{fun} , we can guarantee that $\gamma - \widehat{g}_i(\widehat{a}_s^{n_i-1,\ell}(i)) \leq 2\epsilon_\ell$ once

$$\gamma - g_i(\widehat{a}_s^{n_i-1,\ell}(i)) \leq \epsilon_\ell \iff \widehat{a}_s^{n_i-1,\ell}(i) \geq \widetilde{g}_i^{-1}(\gamma - \epsilon_\ell).$$

It follows that we will have reached the exit condition once

$$\widehat{a}_s^{0,\ell}(i) + n_i\epsilon_\ell \geq \widetilde{g}_i^{-1}(\gamma - \epsilon_\ell) \iff n_i \geq \frac{\widetilde{g}_i^{-1}(\gamma - \epsilon_\ell) - \widehat{a}_s^{0,\ell}(i)}{\epsilon_\ell}. \quad (12)$$

Note also that, for $\ell \geq 2$, on \mathcal{E}_{fun} we must have that

$$\gamma - g_i(\widehat{a}_s^{0,\ell}(i)) \leq 6\epsilon_\ell$$

since we know that the termination criteria $\gamma - \widehat{g}_i(\widehat{a}_s^{0,\ell}(i)) \leq 2\epsilon_{\ell-1} = 4\epsilon_\ell$ was reached for $\widehat{a}_s^{0,\ell}(i) = \widehat{a}_s^{\widehat{n}_{i,\ell-1},\ell-1}(i)$ at epoch $\ell-1$, and since we sample up to tolerance ϵ_ℓ at round ℓ (note that this holds even if the termination criteria was reached for $n_i - 1 = 0$ at round $\ell-1$, since we collect data on Line 8 to ensure we have an accurate enough sample, i.e. we obtain an $\epsilon_{\ell-1}$ -accurate sample at $\widehat{g}_i(\widehat{a}_s^{0,\ell-1}(i))$ instead of relying on our previous $\epsilon_{\ell-2}$ -accurate estimate of this). It follows that $\widehat{a}_s^{0,\ell}(i) \geq \widetilde{g}_i^{-1}(\gamma - 6\epsilon_\ell)$, so a sufficient condition to meet (12) is that

$$n_i \geq \frac{\widetilde{g}_i^{-1}(\gamma - \epsilon_\ell) - \widetilde{g}_i^{-1}(\gamma - 6\epsilon_\ell)}{\epsilon_\ell}.$$

In the case where $\ell = 1$, we can simply replace $\widehat{a}_s^{0,\ell}(i)$ in (12) with $a_{0,i}$.

Finally, if we have that $\widehat{n}_{i,\ell} = 0$, i.e. we terminate on the first iteration, we note that this is still a valid upper bound since $\widetilde{g}_i^{-1}(\gamma - \epsilon_\ell) \geq \widetilde{g}_i^{-1}(\gamma - 6\epsilon_\ell)$. \square

Lemma 21. On \mathcal{E}_{fun} , $\widehat{m}_{i,\ell} \leq \bar{m}_{i,\ell}$ where

$$\bar{m}_{i,\ell} := \begin{cases} \frac{\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - \epsilon_\ell) - \widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 6\epsilon_\ell)}{\epsilon_\ell} & \ell \geq 2 \\ 2(\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{1}{2}) - a_{0,i}) & \ell = 1 \end{cases}.$$

Proof. The proof of this follows identically to the proof of Lemma 20, since while $\text{unsafe}(i) = 0$, the sequences $\widehat{a}_s^{k,\ell}(i)$ and $\widehat{a}_u^{k,\ell}(i)$ are updated in analogous ways, with the γ in the $\widehat{a}_s^{k,\ell}(i)$ update replaced by $\gamma + \epsilon_{\text{safe}}$ in the $\widehat{a}_u^{k,\ell}(i)$ update. \square

Lemma 22. On \mathcal{E}_{fun} , $\widehat{\ell}_{\text{unsafe}}(i) \leq \lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil$.

Proof. Since we only terminate the while loop on Line 15 once $\gamma + \epsilon_{\text{safe}} - \widehat{g}_i(\widehat{a}_u^{m_i-1,\ell}(i)) \leq 2\epsilon_\ell$, we will have that

$$\begin{aligned} \widehat{a}_u^{m_i,\ell}(i) &= \gamma + \epsilon_{\text{safe}} + \widehat{a}_u^{m_i-1,\ell}(i) - \widehat{g}_i(\widehat{a}_u^{m_i-1,\ell}(i)) - \epsilon_\ell \\ &\geq \widehat{a}_u^{m_i-1,\ell}(i) + \epsilon_\ell \\ &\geq \widehat{a}_u^{m_i-1,\ell}(i). \end{aligned}$$

It follows that $\widehat{a}_u^{m_i-1,\ell}(i)$ is increasing in m_i if the termination criteria of the while loop has not been met.

Let $\widehat{m}_{\text{unsafe}}$ denote the m_i value where we set $\text{unsafe}(i) \leftarrow 1$. As the while loop on Line 15 has not yet terminated, it follows that $\widehat{a}_u^{\widehat{m}_{\text{unsafe}},\widehat{\ell}_{\text{unsafe}}(i)}(i) \geq \widehat{a}_u^{\widehat{m}_{i,\widehat{\ell}_{\text{unsafe}}(i)}-1,\widehat{\ell}_{\text{unsafe}}(i)-1}(i) \geq \widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1})$, where the last inequality follows by Lemma 17.

By definition, we will have that $\widehat{g}_i(\widehat{a}_u^{\widehat{m}_{\text{unsafe}},\widehat{\ell}_{\text{unsafe}}(i)}(i)) - \epsilon_{\widehat{\ell}_{\text{unsafe}}(i)} \geq \gamma$ which, on \mathcal{E}_{fun} , is implied by $g_i(\widehat{a}_u^{\widehat{m}_{\text{unsafe}},\widehat{\ell}_{\text{unsafe}}(i)}(i)) - 2\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)} \geq \gamma$. The above implies $g_i(\widehat{a}_u^{\widehat{m}_{\text{unsafe}},\widehat{\ell}_{\text{unsafe}}(i)}(i)) \geq g_i(\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1}))$, so a sufficient condition is that

$$g_i(\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1})) \geq \gamma + 2\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)}$$

Assume that $\widehat{\ell}_{\text{unsafe}}(i)$ is large enough that $\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1}) = g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1})$, then the above is equivalent to

$$\epsilon_{\text{safe}} \geq 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1} + 2\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)} = 8\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)}.$$

Using that $\epsilon_\ell = 2^{-\ell}$ and rearranging this gives that $\widehat{\ell}_{\text{unsafe}}(i) \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}$.

As we assume that $g_i^{-1}(\gamma)$ is well-defined, we will have that $\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1}) = g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1})$ once $\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1} \geq \gamma \iff \epsilon_{\text{safe}} \geq 6\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)} \iff \widehat{\ell}_{\text{unsafe}}(i) \geq \log_2 \frac{6}{\epsilon_{\text{safe}}}$. It follows that once $\ell \geq \log_2 \frac{6}{\epsilon_{\text{safe}}}$, we will be in the regime where $\widetilde{g}_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1}) = g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 3\epsilon_{\widehat{\ell}_{\text{unsafe}}(i)-1})$, and where a sufficient condition for termination will be $\widehat{\ell}_{\text{unsafe}}(i) \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}$. Putting these together we conclude that we will have set $\text{unsafe}(i) = 1$ for $\ell \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}$, so $\widehat{\ell}_{\text{unsafe}}(i) \leq \lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil$. \square

C.3 Sample Complexity of SafeBAI-Monotonic

Lemma 23. *On \mathcal{E}_{fun} , $i^* \in \mathcal{X}_\ell$ for all ℓ .*

Proof. We will only eliminate i^* if $\text{unsafe}(i^*) = 1$ and there exists some $i \neq i^*$ such that

$$\widehat{f}_{i^*}(\widehat{a}_u^{\widehat{m}_{i^*, \ell}}(i^*)) + \epsilon_\ell < \widehat{f}_i(\widehat{a}_s^{\widehat{n}_{i, \ell}}(i)) - \epsilon_\ell.$$

On \mathcal{E}_{fun} , this implies that

$$f_{i^*}(\widehat{a}_u^{\widehat{m}_{i^*, \ell}}(i^*)) < f_i(\widehat{a}_s^{\widehat{n}_{i, \ell}}(i)).$$

By Lemma 15, we will have that $g_{i^*}(\widehat{a}_s^{\widehat{n}_{i, \ell}}(i)) \leq \gamma$, so $\widehat{a}_s^{\widehat{n}_{i, \ell}}(i)$ is safe. By Lemma 18 we will have that $g_{i^*}(\widehat{a}_u^{\widehat{m}_{i^*, \ell}}(i^*)) \geq \gamma$ for $\ell \geq \widehat{\ell}_{\text{unsafe}}(i^*)$, so $\widehat{a}_u^{\widehat{m}_{i^*, \ell}}(i^*)$ is unsafe. However, this is a contradiction because i^* by definition satisfies

$$f_{i^*}(g_{i^*}^{-1}(\gamma)) \geq \max_{i \neq i^*} f_i(g_i^{-1}(\gamma))$$

and since f_i and f_{i^*} are monotonic. Thus, we must have that $i^* \in \mathcal{X}_\ell$ for all ℓ . \square

Lemma 24. *On \mathcal{E}_{fun} , we will have eliminated coordinate $i \neq i^*$ once*

$$\ell \geq \max \left\{ \left\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \right\rceil, \bar{\ell}(i) \right\}$$

where

$$\bar{\ell}(i) := \arg \min_{\ell \in \mathbb{N}} \ell \quad \text{s.t.} \quad f_i(\bar{a}_u^\ell(i)) + 4\epsilon_\ell \leq f_{i^*}(\widetilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell)).$$

Proof. By Lemma 23, we will always have that $i^* \in \mathcal{X}_\ell$. It follows that a sufficient condition to eliminate coordinate i is that $\text{unsafe}(i) = 1$ and

$$\widehat{f}_i(\widehat{a}_u^{\widehat{m}_{i, \ell}}(i)) + \epsilon_\ell \leq \widehat{f}_{i^*}(\widehat{a}_s^{\widehat{n}_{i^*, \ell}}(i^*)) - \epsilon_\ell. \quad (13)$$

By Lemma 22, we will have that $\text{unsafe}(i) = 1$ once $\ell \geq \lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil$. Assume ℓ meets this constraint so that $\text{unsafe}(i) = 1$. On \mathcal{E}_{fun} , (13) is implied by

$$f_i(\widehat{a}_u^{\widehat{m}_{i, \ell}}(i)) + 2\epsilon_\ell \leq f_{i^*}(\widehat{a}_s^{\widehat{n}_{i^*, \ell}}(i^*)) - 2\epsilon_\ell.$$

By Lemma 19, since we have that $\text{unsafe}(i) = 1$, we can upper bound $\widehat{a}_u^{\widehat{m}_{i, \ell}}(i) \leq \bar{a}_u^\ell(i)$, which implies $f_i(\widehat{a}_u^{\widehat{m}_{i, \ell}}(i)) \leq f_i(\bar{a}_u^\ell(i))$. Furthermore, by Lemma 16 we can lower bound $\widehat{a}_s^{\widehat{n}_{i^*, \ell}}(i^*) \geq \widetilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell)$ so $f_{i^*}(\widehat{a}_s^{\widehat{n}_{i^*, \ell}}(i^*)) \geq f_{i^*}(\widetilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell))$. It follows that the above expression is implied by

$$f_i(\bar{a}_u^\ell(i)) + 2\epsilon_\ell \leq f_{i^*}(\widetilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell)) - 2\epsilon_\ell.$$

Combining these gives the result. \square

Theorem 5 (Full version of Theorem 3). *With probability $1 - \delta$, Algorithm 2 will terminate and output i^* after collecting at most*

$$\sum_{i=1}^d \left(\sum_{\ell=1}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (\bar{m}_{i,\ell} + \bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (\bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} \right)$$

samples, and will only pull safe arms during execution, where we define $\bar{\ell}(i^) := \max_{i \neq i^*} \bar{\ell}(i)$.*

Proof. First, by Lemma 13, we will have that $\mathbb{P}[\mathcal{E}_{\text{fun}}] \geq 1 - \delta$. We assume that \mathcal{E}_{fun} holds for the remainder of the proof.

By Lemma 23, on \mathcal{E}_{fun} $i^* \in \mathcal{X}_\ell$ for all ℓ . Since Algorithm 2 only terminates when $|\mathcal{X}_\ell| = 1$, it follows that it will output i^* .

That Algorithm 2 only pulls safe arms on \mathcal{E}_{fun} is given by Lemma 15.

For some suboptimal coordinate $i \neq i^*$, Lemma 24 gives that on \mathcal{E}_{fun} we will eliminate i after

$$\ell \geq \max \left\{ \left\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \right\rceil, \bar{\ell}(i) \right\}.$$

Furthermore, by Lemma 20 and Lemma 21, we can bound $\hat{n}_{i,\ell} \leq \bar{n}_{i,\ell}$ and $\hat{m}_{i,\ell} \leq \bar{m}_{i,\ell}$. Furthermore, once $\text{unsafe}(i) = 1$, we can bound the number of iterations of the while loop on Line 24 by $\log_2(1/\epsilon_\ell) = \ell$, since after this many iterations the if statement on Line 26 will have been met. Putting this together, we can upper bound

$$t \leq \bar{t} := \sum_{i=1}^d \sum_{\ell=1}^{\bar{\ell}(i)} (\bar{m}_{i,\ell} + \bar{n}_{i,\ell} + \ell + 2)$$

where the 2 arises since we always collect at least one set of samples for both safe and unsafe at every epoch. By this same reasoning, we can bound the total number of samples collected by Algorithm 2 from pulling suboptimal arms by

$$\sum_{i=1, i \neq i^*}^d \left(\sum_{\ell=1}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (\bar{m}_{i,\ell} + \bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (\bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} \right)$$

where we have used Lemma 22 to upper bound the number of iterations where $\text{unsafe}(i) = 0$, and thus where we have to pay $\bar{m}_{i,\ell}$. Furthermore, we can bound the total number of pulls to i^* by

$$\max_{i \neq i^*} \left(\sum_{\ell=1}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (\bar{m}_{i^*,\ell} + \bar{n}_{i^*,\ell} + \ell + 2) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (\bar{n}_{i^*,\ell} + \ell + 2) N_{\ell, \bar{t}} \right)$$

where we have used the fact that we will only pull i^* for as many epochs as it takes to eliminate the second-to-last remaining arm. \square

C.4 Interpreting the Complexity

Throughout this section, we will assume that Assumption 3 holds. The following result shows that this implies that $g_i^{-1}(\cdot)$ is L -Lipschitz, for each i .

Proposition 3. *For any function g_i satisfying Assumption 3, we will have that g_i^{-1} exists and satisfies:*

$$|g_i^{-1}(x) - g_i^{-1}(y)| \leq L|x - y|$$

for all $x, y \in [\gamma - 3/2, \gamma + \epsilon_{\text{safe}}]$.

Proof. The following is a standard analysis result

Fact 1. Assume that $f : I \rightarrow \mathbb{R}$ is strictly monotone and continuous on the interval I , and that its derivative f' is defined at some $c \in I$. Then its inverse $f^{-1} : f(I) \rightarrow \mathbb{R}$ is differentiable at $f(c)$ and

$$(f^{-1})'(f(c)) = \frac{1}{f'(c)}.$$

Applying Fact 1 to our setting, we have that

$$(g_i^{-1})'(g_i(x)) = \frac{1}{g_i'(x)} \leq L$$

for any $x \in [g_i^{-1}(\gamma - 3/2), g_i^{-1}(\gamma + \epsilon_{\text{safe}})]$. As the Lipschitz constant of a function over an interval is bounded by its maximum derivative over that interval, the result follows. \square

Lemma 25. Under Assumption 3, we can bound

$$\bar{n}_{i,\ell} \leq \begin{cases} 5L & \ell \geq 2 \\ 2(g_i^{-1}(\gamma - \frac{1}{2}) - a_{0,i}) & \ell = 1 \end{cases}, \quad \bar{m}_{i,\ell} \leq \begin{cases} 5L & \ell \geq 2 \\ 2(g_i^{-1}(\gamma + \epsilon_{\text{safe}} - \frac{1}{2}) - a_{0,i}) & \ell = 1 \end{cases}.$$

Proof. Recall that:

$$\bar{n}_{i,\ell} := \begin{cases} \frac{\tilde{g}_i^{-1}(\gamma - \epsilon_\ell) - \tilde{g}_i^{-1}(\gamma - 6\epsilon_\ell)}{\epsilon_\ell} & \ell \geq 2 \\ 2(\tilde{g}_i^{-1}(\gamma - \frac{1}{2}) - a_{0,i}) & \ell = 1 \end{cases}.$$

As we have assumed that $g_i^{-1}(\gamma - 3/2)$ is well-defined, we can replace all \tilde{g}_i^{-1} in the above with g_i^{-1} . Then using that the inverse of a monotonic function is also monotonic, we can bound

$$\frac{g_i^{-1}(\gamma - \epsilon_\ell) - g_i^{-1}(\gamma - 6\epsilon_\ell)}{\epsilon_\ell} = \frac{|g_i^{-1}(\gamma - \epsilon_\ell) - g_i^{-1}(\gamma - 6\epsilon_\ell)|}{\epsilon_\ell} \leq \frac{5L\epsilon_\ell}{\epsilon_\ell} = 5L.$$

This gives the upper bound on $\bar{n}_{i,\ell}$, and the identical argument can be used to upper bound $\bar{m}_{i,\ell}$. \square

Lemma 26. We can bound

$$\bar{a}_u^\ell(i) \leq g_i^{-1}(\gamma) + \left((L+1)\ell + 6L + \frac{6g_i^{-1}(\gamma)}{\epsilon_{\text{safe}}} \right) 2^{-\ell}.$$

Proof. By definition,

$$\bar{a}_u^\ell(i) = \sum_{s=1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}} + \left(\ell + \frac{4g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{\epsilon_{\text{safe}}} \right) 2^{-\ell}.$$

Our primary concern will be with upper bounding the first term. Let ℓ_0 denote the minimum value of ℓ such that $2\epsilon_s \leq \epsilon_{\text{safe}}$: $\ell_0 = \lceil \log_2 \frac{2}{\epsilon_{\text{safe}}} \rceil$. Then,

$$\begin{aligned} \sum_{s=1}^{\ell} \frac{g_i^{-1}(\min\{\gamma + 2\epsilon_s, \gamma + \epsilon_{\text{safe}}\})}{2^{\ell-s+1}} &= \sum_{s=1}^{\ell_0-1} \frac{g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{2^{\ell-s+1}} + \sum_{s=\ell_0}^{\ell} \frac{g_i^{-1}(\gamma + 2\epsilon_s)}{2^{\ell-s+1}} \\ &\leq g_i^{-1}(\gamma + \epsilon_{\text{safe}}) 2^{\ell_0-\ell-1} + \sum_{s=\ell_0}^{\ell} \frac{g_i^{-1}(\gamma) + 2L\epsilon_s}{2^{\ell-s+1}} \\ &\leq g_i^{-1}(\gamma + \epsilon_{\text{safe}}) 2^{\ell_0-\ell-1} + \sum_{s=1}^{\ell} \frac{g_i^{-1}(\gamma) + 2L\epsilon_s}{2^{\ell-s+1}} \\ &\leq g_i^{-1}(\gamma + \epsilon_{\text{safe}}) 2^{\ell_0-\ell-1} + g_i^{-1}(\gamma) + L\ell 2^{-\ell}. \end{aligned}$$

Furthermore, since $\ell_0 = \lceil \log_2 \frac{\epsilon_{\text{safe}}}{2} \rceil$, we have that

$$g_i^{-1}(\gamma + \epsilon_{\text{safe}}) 2^{\ell_0-\ell-1} \leq g_i^{-1}(\gamma + \epsilon_{\text{safe}}) 2^{\log_2 \frac{2}{\epsilon_{\text{safe}}} - \ell} = \frac{2g_i^{-1}(\gamma + \epsilon_{\text{safe}})}{\epsilon_{\text{safe}}} 2^{-\ell}$$

Finally, we upper bound $g_i^{-1}(\gamma + \epsilon_{\text{safe}}) \leq g_i^{-1}(\gamma) + L\epsilon_{\text{safe}}$. The conclusion then follows by combining these bounds. \square

Lemma 27. *We can bound*

$$\bar{\ell}(i) \leq \max \left\{ 3 + \log_2 \frac{2(L+1)}{\Delta_i} + 2 \log_2 \left(\log_2 \frac{2(L+1)}{\Delta_i} + 3 \right), \log_2 \frac{14L + 12 + 12g_i^{-1}(\gamma)/\epsilon_{\text{safe}}}{\Delta_i}, \log_2 \frac{8}{\epsilon_{\text{safe}}} + 1 \right\}.$$

Proof. Recall that

$$\bar{\ell}(i) = \arg \min_{\ell \in \mathbb{N}} \ell \quad \text{s.t.} \quad f_i(\bar{a}_u^\ell(i)) + 4\epsilon_\ell \leq f_{i^*}(\tilde{g}_{i^*}^{-1}(\gamma - 3\epsilon_\ell)).$$

We first upper bound this by

$$\arg \min_{\ell \in \mathbb{N}, \ell \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}} + 1\ell \quad \text{s.t.} \quad f_i(\bar{a}_u^\ell(i)) + 4\epsilon_\ell \leq f_{i^*}(g_{i^*}^{-1}(\gamma - 3\epsilon_\ell))$$

which, by Lemma 22, implies that we are in the regime where $\ell \geq \widehat{\ell}_{\text{unsafe}}(i)$. As we have assumed $g_{i^*}^{-1}$ is L -Lipschitz, and since the inverse of a monotonic function is monotonic, we have that

$$|g_{i^*}^{-1}(\gamma) - g_{i^*}^{-1}(\gamma - 3\epsilon_\ell)| = g_{i^*}^{-1}(\gamma) - g_{i^*}^{-1}(\gamma - 3\epsilon_\ell) \leq 3L\epsilon_\ell \implies g_{i^*}^{-1}(\gamma - 3\epsilon_\ell) \geq g_{i^*}^{-1}(\gamma) - 3L\epsilon_\ell.$$

As f_{i^*} is 1-Lipschitz, this then implies that

$$f_{i^*}(g_{i^*}^{-1}(\gamma)) - f_{i^*}(g_{i^*}^{-1}(\gamma - 3\epsilon_\ell)) \leq g_{i^*}^{-1}(\gamma) - g_{i^*}^{-1}(\gamma - 3\epsilon_\ell) \leq 3L\epsilon_\ell \implies f_{i^*}(g_{i^*}^{-1}(\gamma)) - 3L\epsilon_\ell \leq f_{i^*}(g_{i^*}^{-1}(\gamma - 3\epsilon_\ell)).$$

Similarly, using that f_i is 1-Lipschitz, we can bound (using that, by Lemma 19, $\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) \leq \bar{a}_u^\ell(i)$ and by Lemma 18 $\widehat{a}_u^{\widehat{m}_{i,\ell},\ell}(i) \geq g_i^{-1}(\gamma)$, which implies that $\bar{a}_u^\ell(i) \geq g_i^{-1}(\gamma)$):

$$f_i(\bar{a}_u^\ell(i)) - f_i(g_i^{-1}(\gamma)) \leq \bar{a}_u^\ell(i) - g_i^{-1}(\gamma) \implies f_i(\bar{a}_u^\ell(i)) \leq f_i(g_i^{-1}(\gamma)) + |\bar{a}_u^\ell(i) - g_i^{-1}(\gamma)|.$$

It follows that we can upper bound

$$\begin{aligned} \bar{\ell}(i) &\leq \arg \min_{\ell \in \mathbb{N}, \ell \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}} \ell \quad \text{s.t.} \quad f_i(g_i^{-1}(\gamma)) + |\bar{a}_u^\ell(i) - g_i^{-1}(\gamma)| + 4\epsilon_\ell \leq f_{i^*}(g_{i^*}^{-1}(\gamma)) - 3L\epsilon_\ell \\ &= \arg \min_{\ell \in \mathbb{N}, \ell \geq \log_2 \frac{8}{\epsilon_{\text{safe}}}} \ell \quad \text{s.t.} \quad |\bar{a}_u^\ell(i) - g_i^{-1}(\gamma)| + (4 + 3L)\epsilon_\ell \leq \Delta_i. \end{aligned}$$

By Lemma 26, we can bound

$$|\bar{a}_u^\ell(i) - g_i^{-1}(\gamma)| \leq \left((L+1)\ell + 6L + \frac{6g_i^{-1}(\gamma)}{\epsilon_{\text{safe}}} \right) 2^{-\ell},$$

so it follows that an upper bound on $\bar{\ell}(i)$ is any ℓ satisfying

$$\begin{aligned} 2^\ell &\geq \max \left\{ \frac{2(L+1)\ell}{\Delta_i}, 2 \frac{7L + 6 + 6g_i^{-1}(\gamma)/\epsilon_{\text{safe}}}{\Delta_i}, \log_2 \frac{8}{\epsilon_{\text{safe}}} + 1 \right\} \\ &\geq \max \left\{ \frac{(L+1)\ell + 7L + 6 + 6g_i^{-1}(\gamma)/\epsilon_{\text{safe}}}{\Delta_i}, \log_2 \frac{8}{\epsilon_{\text{safe}}} + 1 \right\}. \end{aligned}$$

To ensure that $2^\ell \geq \frac{2(L+1)\ell}{\Delta_i} \iff \ell \geq \log_2 \ell + \log_2(2(L+1)/\Delta_i)$ it suffices to take

$$\ell \geq \max\{\log_2(2(L+1)/\Delta_i), 3\} + 2 \log_2(\max\{\log_2(2(L+1)/\Delta_i), 3\})$$

since, in this case,

$$\begin{aligned} \log_2 \ell + \log_2(2(L+1)/\Delta_i) &\leq \log_2(\max\{\log_2(2(L+1)/\Delta_i), 3\} + 2 \log_2(\max\{\log_2(2(L+1)/\Delta_i), 3\})) \\ &\quad + \log_2(2(L+1)/\Delta_i) \end{aligned}$$

$$\begin{aligned}
 &\leq \log_2(3 \max\{\log_2(2(L+1)/\Delta_i), 3\}) + \log_2(2(L+1)/\Delta_i) \\
 &= \log_2 3 + \log_2(\max\{\log_2(2(L+1)/\Delta_i), 3\}) + \log_2(2(L+1)/\Delta_i) \\
 &\leq 2 \log_2(\max\{\log_2(2(L+1)/\Delta_i), 3\}) + \max\{\log_2(2(L+1)/\Delta_i), 3\} \\
 &= \ell.
 \end{aligned}$$

We conclude that

$$\bar{\ell}(i) \leq \max \left\{ 3 + \log_2 \frac{2(L+1)}{\Delta_i} + 2 \log_2 \left(\log_2 \frac{2(L+1)}{\Delta_i} + 3 \right), \log_2 \frac{14L+12+12g_i^{-1}(\gamma)/\epsilon_{\text{safe}}}{\Delta_i}, \log_2 \frac{8}{\epsilon_{\text{safe}}} + 1 \right\}.$$

□

Proof of Corollary 2. This result follows from Theorem 3 and the upper bounds given in Lemma 25 and Lemma 27. By Theorem 3, we can bound the sample complexity as

$$\sum_{i=1}^d \left(\sum_{\ell=1}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (\bar{m}_{i,\ell} + \bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (\bar{n}_{i,\ell} + \ell + 2) N_{\ell, \bar{t}} \right).$$

By Lemma 25, we can upper bound $\bar{n}_{i,\ell} \leq 5L$ for $\ell \geq 2$ and $\bar{n}_{i,1} \leq 2(g_i^{-1}(\gamma - 1/2) - a_{0,i})$ and similarly $\bar{m}_{i,\ell} \leq 5L$ and $\bar{m}_{i,1} \leq 2(g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 1/2) - a_{0,i})$. By Lemma 27 we can bound

$$\bar{\ell}(i) \leq \max \left\{ 3 + \log_2 \frac{2(L+1)}{\Delta_i} + 2 \log_2 \left(\log_2 \frac{2(L+1)}{\Delta_i} + 3 \right), \log_2 \frac{14L+12+12g_i^{-1}(\gamma)/\epsilon_{\text{safe}}}{\Delta_i}, \log_2 \frac{8}{\epsilon_{\text{safe}}} + 1 \right\}.$$

Recall that $N_{\bar{t},\ell} = 2 \log \frac{8\bar{t}^2}{\delta} \cdot 2^{2\ell}$. For $i \neq i^*$, letting \lesssim denote inequality

$$\begin{aligned}
 &\sum_{\ell=1}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (\bar{m}_{i,\ell} + \bar{n}_{i,\ell} + \ell) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (\bar{n}_{i,\ell} + \ell) N_{\ell, \bar{t}} \\
 &\leq \left(3 + g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 1/2) + g_i^{-1}(\gamma - 1/2) - 2a_{0,i} \right) N_{1, \bar{t}} \\
 &\quad + \sum_{\ell=2}^{\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil} (2 + 10L + \bar{\ell}(i)) N_{\ell, \bar{t}} + \sum_{\ell=\lceil \log_2 \frac{8}{\epsilon_{\text{safe}}} \rceil + 1}^{\bar{\ell}(i)} (1 + 5L + \bar{\ell}(i)) N_{\ell, \bar{t}} \\
 &\leq 8 \left(3 + g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 1/2) + g_i^{-1}(\gamma - 1/2) - 2a_{0,i} \right) \log \frac{8\bar{t}^2}{\delta} \\
 &\quad + 4(2 + 10L + \bar{\ell}(i)) \log \frac{8\bar{t}^2}{\delta} \cdot 2^{2\bar{\ell}(i)} \\
 &\leq 8 \left(3 + g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 1/2) + g_i^{-1}(\gamma - 1/2) - 2a_{0,i} \right) \log \frac{8\bar{t}^2}{\delta} \\
 &\quad \tilde{\mathcal{O}} \left((2 + 10L + \bar{\ell}(i)) \log \frac{8\bar{t}^2}{\delta} \cdot \frac{1 + L^2 + g_i^{-1}(\gamma)^2/\epsilon_{\text{safe}}^2}{\Delta_i^2} \right)
 \end{aligned}$$

where $\tilde{\mathcal{O}}(\cdot)$ hides logarithmic terms and absolute constants. Noting that $\bar{\ell}(i)$ is only logarithmic in problem parameters, we can bound all of this by

$$\tilde{\mathcal{O}} \left(\left(1 + g_i^{-1}(\gamma + \epsilon_{\text{safe}} - 1/2) + g_i^{-1}(\gamma - 1/2) - 2a_{0,i} \right) \log \frac{1}{\delta} + \frac{(1+L)(1+L^2+g_i^{-1}(\gamma)^2/\epsilon_{\text{safe}}^2)}{\Delta_i^2} \cdot \log \frac{1}{\delta} \right).$$

A similar upper bound can be used to bound the contribution of i^* . Combining these gives the result. □

D Additional Experiments and details

In the following, we elaborate some additional details of experiments in Section 7. First we note that in our instances, we need to choose the initial points $a_{0,i}$ and boundary values M_i . In the linear response model, we choose $a_{0,i} = 0.1$ and $M_i = 1.5$ for all i , whereas in the nonlinear drug response model, we set initial safe dosage value $a_{0,i} = -3$ and $M_i = -0.5$ for all i .

We try to make our implementation of SAFEOPT more realistic and efficient. In the linear response model, we choose a linear kernel with a β_t value selected from Abbasi-Yadkori et al. (2011):

$$\beta_t(\delta) = R \sqrt{\log \left(\frac{d(1 + \frac{(t-1)L^2}{\lambda})}{\delta} \right)} + \sqrt{\gamma} S$$

with $\gamma = 1$ (note that this only scales logarithmically in d since it is assumed the learner knows the coordinates are orthogonal). In the drug response model, we choose a RBF kernel $k_{rbf}(x, x') = \sigma_{rbf}^2 \exp \left(-\frac{(x-x')^2}{2l^2} \right)$ with signal variance $\sigma_{rbf}^2 = 1$ and length scale $l = 1$.

To simplify the algorithms' executions in the drug response model, we made some minor changes to the algorithms in our experiments. In SAFEBAI-MONOTONIC we replaced the theoretical expression for $N_{\ell,t}$ from $\lceil 2 \log \frac{8t^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$ to $\lceil 2 \log \frac{8t^2}{\delta} \cdot \epsilon_\ell^{-2} \rceil$. In SAFEOPT, we replaced the theoretical β_t value with a fixed constant 3. These modifications slightly reduced the amount of pulls needed. Nonetheless, the empirical performances were satisfactory, as no error in result was found.