
Choosing Answers in ε -Best-Answer Identification for Linear Bandits

Marc Jourdan¹ Rémy Degenne¹

Abstract

In pure-exploration problems, information is gathered sequentially to answer a question on the stochastic environment. While best-arm identification for linear bandits has been extensively studied in recent years, few works have been dedicated to identifying one arm that is ε -close to the best one (and not exactly the best one). In this problem with several correct answers, an identification algorithm should focus on one candidate among those answers and verify that it is correct. We demonstrate that picking the answer with highest mean does not allow an algorithm to reach asymptotic optimality in terms of expected sample complexity. Instead, a *furthest answer* should be identified. Using that insight to choose the candidate answer carefully, we develop a simple procedure to adapt best-arm identification algorithms to tackle ε -best-answer identification in transductive linear stochastic bandits. Finally, we propose an asymptotically optimal algorithm for this setting, which is shown to achieve competitive empirical performance against existing modified best-arm identification algorithms.

1. Introduction

The multi-armed bandit (MAB) setting is an extensively studied problem in sequential decision making (Robbins, 1952; Lattimore & Szepesvári, 2020). The environment is represented by a set of arms, each associated with an unknown reward distribution. The agent interacts with it by sequentially “pulling” arms, that is choosing an arm and observing a sample from its distribution. We focus on the pure-exploration framework in which the objective of the agent is to answer a query as fast and reliably as possible, while disregarding the accumulated cost. We consider noisy linear observations depending on an unknown parameter μ .

¹Univ. Lille, CNRS, Inria, Centrale Lille, UMR 9198-CRISTAL, F-59000 Lille, France. Correspondence to: Marc Jourdan <marc.jourdan@inria.fr>.

The pure-exploration setting for stochastic bandits was first studied in Even-Dar et al. (2002) and Bubeck et al. (2009). It has been studied in two major theoretical frameworks (Audibert et al., 2010; Gabillon et al., 2012; Jamieson & Nowak, 2014; Kaufmann et al., 2016): the *fixed-budget* setting and the *fixed-confidence* setting. In the fixed-budget setting, the objective is to minimize the probability of misidentifying a correct answer given a fixed number of samples. We consider the fixed-confidence setting, where the agent aims at minimizing the number of pulls to identify a correct answer with confidence $1 - \delta$.

The most well known pure exploration setting is best-arm identification (BAI) (Audibert et al., 2010; Chen et al., 2017), in which the goal is to return the arm with largest expected reward. Numerous variants of BAI have been considered in recent years: linear bandits (Soare et al., 2014; Zaki et al., 2020; Degenne et al., 2020a), thresholding bandits (Locatelli et al., 2016; Cheshire et al., 2021), minimum threshold (Kaufmann et al., 2018), combinatorial bandits (Chen et al., 2014; Katz-Samuels et al., 2020; Jourdan et al., 2021), top-m identification (Kalyanakrishnan et al., 2012; Katz-Samuels & Scott, 2019; Réda et al., 2021), matching bandits (Sentenac et al., 2021), logistic bandits (Jun et al., 2021), identifying all ε -optimal answers (Mason et al., 2020; Al Marjani et al., 2022), etc.

When the gap between the best and the second best arm is small, BAI problems are difficult, meaning that an algorithm needs a large number of samples to be correct. To avoid wasteful queries, practitioners might be interested in the easier task of identifying one answer that is ε -close to the best one, but not exactly the best one.

In this work, we consider (ε, δ) -PAC best-answer identification (ε -BAI) for transductive linear bandits. In transductive bandits (Fiez et al., 2019), the set of arms \mathcal{K} that can be pulled by the agent is different from the set of answers \mathcal{Z} on which the identification procedure focuses. In contrast to best-answer identification, our agent aims at identifying one ε -optimal answer (defined below) among the existing ones (Mannor & Tsitsiklis, 2004; Even-Dar et al., 2006; Sabato, 2019). Note that the ranking and selection literature explores this question with a different approach, see Hong et al. (2021) for a review. ε -BAI has been studied for MAB (Garivier & Kaufmann, 2021), spectral bandits (Kocák &

Garivier, 2021), on the unit sphere (Jedra & Proutière, 2020) and as a special case of the multiple correct answer setting (Degenne & Koolen, 2019). As we will show, the asymptotic complexity of ε -BAI is governed by the identification of the *furthest answer*. This is the ε -optimal answer for which it is easiest to verify ε -optimality. Therefore, all ε -optimal answers are not equivalent. When aiming at asymptotic optimality of the expected sample complexity, algorithms should select the candidate ε -optimal answer more carefully than simply using the *greedy answers* defined as $z^*(\mu) \stackrel{\text{def}}{=} \arg \max_{z \in \mathcal{Z}} \langle \mu, z \rangle$ (the answers for which the reward is maximal). In the special case of BAI, the assumption $|z^*(\mu)| = 1$ is made and algorithms use the greedy answer, which is the unique correct answer to identify.

For our algorithm, we adopt a saddle-point (or game) approach between the agent and the *nature*, aiming at iteratively approximating the lower bound on the expected sample complexity. The goal is to design an asymptotically optimal algorithm with competitive empirical performance in finite-time regime.

Contributions (1) We provide an analysis of ε -BAI for transductive linear bandits and highlight a phenomenon which was overlooked by previous work. The choice of the candidate ε -optimal answer is crucial to reach asymptotic optimality in terms of expected sample complexity and one should identify the furthest answer instead of using the greedy answers. (2) By carefully choosing the candidate ε -optimal answer and leaving the sampling rule unchanged, we develop a simple procedure to adapt BAI algorithms to be (ε, δ) -PAC and empirically competitive for ε -BAI in transductive linear stochastic bandits. (3) By leveraging the concept of ε -optimal answer in the sampling rule, we propose an asymptotically optimal algorithm in this setting, which has competitive empirical performance.

1.1. Problem Statement

Transductive Linear Bandits We consider the *transductive linear bandits* setting, where the collection of arms $\mathcal{K} \subseteq \mathbb{R}^d$ and answers $\mathcal{Z} \subseteq \mathbb{R}^d$ are finite sets, potentially different (Fiez et al., 2019), with cardinality $|\mathcal{K}| = K$ and $|\mathcal{Z}| = Z$. Taking $\mathcal{Z} = \mathcal{K}$ yields the *linear bandits* setting. We assume that \mathcal{K} spans \mathbb{R}^d and denote $L_{\mathcal{K}} \stackrel{\text{def}}{=} \max_{a \in \mathcal{K}} \|a\|_2$ where $\|a\|_2$ is the euclidean norm of $a \in \mathcal{K}$. The interaction with the environment goes as follows: in each round $t \geq 1$, the agent chooses an arm $a_t \in \mathcal{K}$ and observes $X_t^{a_t} = \langle \mu, a_t \rangle + \eta_t$ where μ is an unknown *mean parameter* belonging to the set $\mathcal{M} \subseteq \mathbb{R}^d$, which is known to the agent¹. The noise $\eta_t \sim \mathcal{N}(0, \sigma^2)$ is conditionally independent of the past. Without loss of generality,

¹Note that we use $a \in \mathcal{K}$ as a superscript to denote the index of the element in \mathbb{R}^K corresponding to the vector a .

we consider $\sigma^2 = 1$. Prior works have lifted the Gaussian assumption by considering sub-Gaussian distributions. The focus of our work is to highlight a phenomenon which is orthogonal to the distribution, therefore restricting ourselves to Gaussians is an assumption we are willing to make.

(ε, δ) -PAC Best-Answer Identification In ε -BAI ($\varepsilon \geq 0$) for transductive linear bandits, the agent aims at identifying one of the ε -optimal answers by sequentially pulling arms. We address two different notions of ε -optimality: the *additive ε -optimal* answers, $\mathcal{Z}_{\varepsilon}^{\text{add}}(\mu) \stackrel{\text{def}}{=} \{z \in \mathcal{Z} : \langle \mu, z \rangle \geq \max_{z \in \mathcal{Z}} \langle \mu, z \rangle - \varepsilon\}$, and the *multiplicative ε -optimal* answers, $\mathcal{Z}_{\varepsilon}^{\text{mul}}(\mu) \stackrel{\text{def}}{=} \{z \in \mathcal{Z} : \langle \mu, z \rangle \geq (1 - \varepsilon) \max_{z \in \mathcal{Z}} \langle \mu, z \rangle\}$ when $\max_{z \in \mathcal{Z}} \langle \mu, z \rangle > 0$. The notation $\cdot^{\{\text{add}, \text{mul}\}}$ is dropped when the statement holds for both notions of ε -optimality. We will deal with both notions with the same method. Previous works mostly consider the additive ε -optimality (Garivier & Kaufmann, 2021; Kocák & Garivier, 2021). When $\varepsilon = 0$ both notions coincide with BAI, in which there is a unique correct answer. ε -BAI is often seen as a more practical objective than BAI, in cases where getting an answer close to optimal is enough: while a BAI algorithm will spend many samples distinguishing between an ε -optimal answer and the best answer, an ε -BAI algorithm will be able to stop quickly.

Identification Strategy The σ -algebra $\mathcal{F}_t \stackrel{\text{def}}{=} \sigma(a_1, X_1^{a_1}, \dots, a_t, X_t^{a_t})$, called *history*, encompasses all the information available to the agent after t rounds. In the fixed-confidence setting an *identification strategy* is described by three rules: a *sampling rule* $(a_t)_{t \geq 1}$ where $a_t \in \mathcal{K}$ is \mathcal{F}_{t-1} -measurable, a *stopping rule* τ_{δ} which is a stopping time with respect to the filtration $(\mathcal{F}_t)_{t \geq 1}$, also referred as the *sample complexity*, and a *recommendation rule* \hat{z} which is $\mathcal{F}_{\tau_{\delta}}$ -measurable. While the sampling rule could depend on additional internal randomization, the algorithms proposed in this work are deterministic.

In the fixed-confidence setting, the learner is given a confidence parameter $\delta \in (0, 1)$. A strategy is said to be (ε, δ) -PAC if, for all $\mu \in \mathcal{M}$, with probability at most δ it terminates while not recommending an ε -optimal answer, i.e. $\mathbb{P}_{\mu}[\tau_{\delta} < +\infty, \hat{z} \notin \mathcal{Z}_{\varepsilon}(\mu)] \leq \delta$. Among the class of (ε, δ) -PAC algorithms, our goal is to minimize the expected sample complexity $\mathbb{E}_{\mu}[\tau_{\delta}]$.

2. Comparing ε -Optimal Answers

2.1. Lower Bound

For any $w \in (\mathbb{R}^+)^K$, we define the design matrix $V_w \stackrel{\text{def}}{=} \sum_{a \in \mathcal{K}} w^a a a^{\top} \in \mathbb{R}^{d \times d}$, which is symmetric and positive semi-definite, and definite if and only if $\text{Span}(\{a \in \mathcal{K} :$

$w^a \neq 0\}) = \mathbb{R}^d$. For any symmetric positive semi-definite matrix $V \in \mathbb{R}^{d \times d}$, we define the semi-norm $\|x\|_V \stackrel{\text{def}}{=} \sqrt{x^\top V x}$ for $x \in \mathbb{R}^d$, which is a norm if V is positive definite. The probability simplex of dimension $K - 1$ is denoted by Δ_K for all $K \geq 2$.

Alternative to z Given an answer $z \in \mathcal{Z}$, the *alternative to z* is defined as the set of parameters for which z is not an ε -optimal answer, $\neg_\varepsilon z \stackrel{\text{def}}{=} \overline{\{\lambda \in \mathcal{M} : z \notin \mathcal{Z}_\varepsilon(\lambda)\}}$ where \overline{X} denotes the closure of X . Rewriting it for the additive and multiplicative ε -optimality, we obtain: $\neg_\varepsilon^{\text{add}} z = \overline{\{\lambda \in \mathcal{M} : \langle \lambda, z \rangle < \max_{z \in \mathcal{Z}} \langle \lambda, z \rangle - \varepsilon\}}$ and $\neg_\varepsilon^{\text{mul}} z = \overline{\{\lambda \in \mathcal{M} : \langle \lambda, z \rangle < (1 - \varepsilon) \max_{z \in \mathcal{Z}} \langle \lambda, z \rangle\}}$. Identifying an ε -optimal answer $z \in \mathcal{Z}_\varepsilon(\mu)$ is equivalent to rejecting the hypothesis that the unknown mean belongs to the alternative to z , i.e. $\mathcal{H}_0 = \{\mu \in \neg_\varepsilon z\}$. Informally, if we know with enough certainty that μ does not belong to the alternative $\neg_\varepsilon z$, we can safely return the answer z .

Asymptotic Lower Bound Theorem 2.1 gives an asymptotic lower bound on the expected sample complexity of any (ε, δ) -PAC strategy for both additive and multiplicative ε -optimality. This is a corollary of Theorem 1 in Degenne & Koolen (2019), which holds for any multiple answer instance and sub-Gaussian distributions (Appendix C.2).

Theorem 2.1 (Theorem 1 in Degenne & Koolen (2019)). *For all (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,*

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_\varepsilon(\mu)$$

where the inverse of the characteristic time is

$$T_\varepsilon(\mu)^{-1} \stackrel{\text{def}}{=} \max_{z \in \mathcal{Z}_\varepsilon(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2. \quad (1)$$

An (ε, δ) -PAC algorithm is said to be *asymptotically optimal* if the bound is tight: for all $\mu \in \mathcal{M}$, $\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \leq T_\varepsilon(\mu)$. The first lower bound for BAI was proved in Garivier & Kaufmann (2016).

As noted by Chernoff (1959), the complexity $T_\varepsilon(\mu)^{-1}$ is the value of a zero-sum game between two players. The agent chooses an ε -optimal answer and a pulling proportion over arms, $(z, w) \in \mathcal{Z}_\varepsilon(\mu) \times \Delta_K$. The nature plays the most confusing alternative $\lambda \in \neg_\varepsilon z$ with respect to a reweighted Kullback-Leibler divergence ($\|\cdot\|_{V_w}^2$ for Gaussians) in order to fool the agent into rejecting this answer. Our algorithm, named **L ε BAI** (**L**inear **ε -BAI**), is based on this formulation. Even for known μ , computing $T_\varepsilon(\mu)^{-1}$ is in general intractable due to the non-convexity of $\neg_\varepsilon z$ and the additional maximization over $\mathcal{Z}_\varepsilon(\mu)$. When ε is large enough to have $\mathcal{Z}_\varepsilon(\lambda) = \mathcal{Z}$ for all $\lambda \in \mathcal{M}$, then $T_\varepsilon(\mu) = 0$, i.e. it is *so easy* that no sample is needed.

While lower bounds for BAI have been derived in the non-asymptotic regime, it remains unclear whether equivalent lower bounds hold for ε -BAI (Garivier & Kaufmann, 2021).

Comparison with BAI Since $T_\varepsilon(\mu) \leq T_0(\mu)$ for all $\mu \in \mathcal{M}$ (because $\neg_\varepsilon z \subseteq \neg_0 z$), ε -BAI is easier than BAI. There exists arbitrarily hard BAI instances that can be solved if seen as an ε -BAI problem, e.g. when the gap between the best and the second best arm is arbitrarily small.

2.2. Furthest Answer

Our contributions are linked with the concept the furthest answer: it should be leveraged in the recommendation-stopping pair (Section 3) and in the sampling rule (Section 4). In a nutshell, to reach asymptotic optimality in terms of sample complexity one should identify that furthest answer instead of simply using the greedy answers: *all ε -optimal answers are not equivalent*.

The furthest answer $z_F(\mu)$ is the ε -optimal answer for which it is easiest to verify that μ does not belong to its alternative (when using an optimal allocation over arms $w_F(\mu) \in \Delta_K$). Introduced in Degenne & Koolen (2019) and Garivier & Kaufmann (2021), it is defined as

$$(z_F(\mu), w_F(\mu)) \stackrel{\text{def}}{=} \arg \max_{(z, w) \in \mathcal{Z}_\varepsilon(\mu) \times \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2. \quad (2)$$

$z_F(\mu)$ belongs to the ε -optimal answers $\mathcal{Z}_\varepsilon(\mu)$, as does the greedy answers $z^*(\mu) = \arg \max_{z \in \mathcal{Z}} \langle \mu, z \rangle$, but these answers may differ. In BAI with a unique best arm the set $\mathcal{Z}_\varepsilon(\mu)$ is a singleton, hence those two notions coincide.

We assume there is a unique furthest answer for the unknown μ , i.e. $|z_F(\mu)| = 1$. When $|z_F(\mu)| > 1$, some function of μ has to have exactly the same value for all answers of the set. This happens with probability 0 if μ arises from an absolutely continuous distribution. Almost all BAI algorithms make the assumption that $|z^*(\mu)| = 1$, which implies $|z_F(\mu)| = 1$ in the BAI case. Since the furthest answer is assumed unique, we abuse notation and denote by $z_F(\mu)$ both that answer, and the singleton containing it as in (2). $z^*(\mu)$ denotes a set as we don't assume $|z^*(\mu)| = 1$. The dependence of $z_F(\mu)$ and $w_F(\mu)$ on ε is omitted.

Asymptotic Sub-optimality of $z^*(\mu)$ An (ε, δ) -PAC strategy is said to be *asymptotically greedy* if the only ε -optimal answers for which the algorithm will stop asymptotically are the greedy answers $z^*(\mu)$, i.e. for all $\mu \in \mathcal{M}$,

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] = 0. \quad (3)$$

Lemma 2.2 shows that any asymptotically greedy (ε, δ) -PAC strategy is asymptotically sub-optimal whenever

$z_F(\mu) \notin z^*(\mu)$, i.e. it can only reach $T_{g,\varepsilon}(\mu)$ which is strictly higher than $T_\varepsilon(\mu)$.

Lemma 2.2. *For all asymptotically greedy (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,*

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_{g,\varepsilon}(\mu)$$

where the inverse of the greedy characteristic time is

$$T_{g,\varepsilon}(\mu)^{-1} \stackrel{\text{def}}{=} \max_{z \in z^*(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2, \quad (4)$$

and $T_{g,\varepsilon}(\mu) > T_\varepsilon(\mu)$ if and only if $z_F(\mu) \notin z^*(\mu)$.

Lemma 2.3 shows that any (ε, δ) -PAC strategy recommending any greedy answers $\hat{z} \in z^*(\mu_{\tau_\delta})$ which succeeds in identifying $z^*(\mu)$ is asymptotically greedy. Since $\mu \mapsto z^*(\mu)$ is continuous and \mathcal{Z} is finite, it is sufficient to have a sampling rule ensuring that $\lim_{t \rightarrow +\infty} \mu_t = \mu$.

Lemma 2.3. *Any (ε, δ) -PAC strategy recommending $\hat{z} \in z^*(\mu_{\tau_\delta})$ is asymptotically greedy if the sampling rule ensures that $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, z^*(\mu_{\tau_\delta}) = z^*(\mu)] = 1$.*

Asymptotic Optimality of $z_F(\mu)$ The furthest answer has by definition a central role in the characteristic time. Among the oracles that first choose an answer $z \in \mathcal{Z}_\varepsilon(\mu)$ and then sample according to the optimal proportions to verify that $\mu \notin \neg_\varepsilon z$, the only one achieving asymptotic optimality is the one picking $z_F(\mu)$. To be asymptotically optimal, an ε -BAI algorithm has to implicitly identify $z_F(\mu)$.

The definition of $z_F(\mu)$ comes from an asymptotic lower bound, and no finite time lower bounds are available for ε -BAI. It could be that for larger δ (hence small stopping times), identifying $z_F(\mu)$ among $\mathcal{Z}_\varepsilon(\mu)$ is too costly to be done before stopping. In that regime, it could be that an algorithm cannot do better than picking any ε -optimal answer. This is an interesting open question for future work. Strong moderate confidence terms (independent of δ) affecting the sample complexity have been shown in different settings (Katz-Samuels & Jamieson, 2020; Mason et al., 2020).




Figure 1. Influence of ε on (a) the proportion of draws where $z_F(\mu) \notin z^*(\mu)$, (b) the median (and first/third quartile) of $\frac{T_\varepsilon^{\text{mul}}(\mu)}{T_{g,\varepsilon}^{\text{mul}}(\mu)}$, when $z_F(\mu) \notin z^*(\mu)$.

Numerical Simulations We compare the furthest and the greedy answers for the multiplicative ε -optimality (see Figure 8 in Appendix F.2.2 for additive ε -optimality).

We consider $d = 2$, $\mathcal{M} = \mathbb{R}^2$ and $\mathcal{Z} = \mathcal{K}$ with $K = 4$. We use $\mu = (1, 0)$ and generate 25000 random instances. In each one of them, we consider $z_1 = \mu$ and draw uniformly at random $z_2 \in \{(\cos(\theta), \sin(\theta)) : \theta \in [-\theta_\varepsilon, \theta_\varepsilon]\}$ and $z_3, z_4 \in \{(\cos(\theta), \sin(\theta)) : \theta \in (-\pi, -\theta_\varepsilon) \cup (\theta_\varepsilon, \pi)\}$, where $\theta_\varepsilon \stackrel{\text{def}}{=} \arccos(1 - \varepsilon)$. This yields $z_1 = z^*(\mu)$, $z_2 \in \mathcal{Z}_\varepsilon(\mu)$ and $z_3, z_4 \in \mathcal{Z} \setminus \mathcal{Z}_\varepsilon(\mu)$. To approximate $(T_\varepsilon(\mu), z_F(\mu))$, we discretize Δ_4 with 10000 vectors. This is repeated for several values of ε . We never observed $|z^*(\mu)| > 1$ or $|z_F(\mu)| > 1$.

Figure 1(a) reveals that the proportion of draws where $z^*(\mu) \neq z_F(\mu)$ is not negligible ($\approx 14\%$). On those instances, Figure 1(b) shows that $\frac{T_\varepsilon^{\text{mul}}(\mu)}{T_{g,\varepsilon}^{\text{mul}}(\mu)}$ is on average 0.9. Therefore, when they are different, the furthest answer has a 10% lower characteristic time than greedy answers.

3. From BAI to ε -BAI Algorithms

We propose a simple procedure to convert any BAI algorithm into an (ε, δ) -PAC algorithm. While leaving the original sampling rule unchanged, the stopping-recommendation rule are carefully chosen thanks to the concept of furthest answer.

Structure Since ε -BAI is easier than BAI, the stopping rule of BAI algorithms has to be modified for ε -BAI. Instead of stopping whenever a single best arm is identified, it is enough to stop when we know that an arm is ε -best. In most (ε -)BAI algorithms, the stopping-recommendation pair and the sampling rule can be thought as two independent blocks. There exists stopping-recommendation pairs that guarantee the strategy to be (ε, δ) -PAC regardless of the sampling rule (e.g. see Lemma 3.1). Therefore, we can take the sampling rule from a BAI algorithm and couple it with a stopping-recommendation pair with this property.

We will now describe such a stopping-recommendation pair for ε -BAI in transductive linear Gaussian bandits. Due to its generality, this procedure can be readily adapted to tackle general distributions (e.g. sub-Gaussian) and different structures (e.g. spectral bandits) by simply adapting the stopping rule and its associated threshold.

3.1. Stopping-Recommendation Pairs

Estimator Let $N_{t-1} \in (\mathbb{R}^+)^K$ denotes the counts of pulled arms at the start of round t , $N_{t-1}^a = \sum_{s=1}^{t-1} \mathbf{1}_{\{a_s=a\}}$. We denote the Ordinary Least Square (OLS) estimator by $\mu_{t-1} = V_{N_{t-1}}^{-1} \sum_{s=1}^{t-1} X_s^{a_s} a_s$. When $\mu_{t-1} \in \mathcal{M}$, this is also the Maximum Likelihood Estimator (MLE).

GLR-based Stopping Rule Given a candidate answer $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$ and the history \mathcal{F}_t , the algorithm stops as soon as the Generalized Likelihood Ratio (GLR, Appendix C.1) exceeds a stopping threshold $\beta(t-1, \delta)$

$$\inf_{\lambda \in \neg_\varepsilon z_t} \|\mu_{t-1} - \lambda\|_{V_{N_{t-1}}}^2 > 2\beta(t-1, \delta). \quad (5)$$

In Lemma 3.1, we show that combining a recommendation rule such that $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$ and this stopping rule is sufficient to obtain a (ε, δ) -PAC strategy regardless of the sampling rule. This holds even when the stopping criterion is checked only on an infinite subset of \mathbb{N} . The proof (Appendix D) leverages the concentration inequalities of Kaufmann & Koolen (2021).

Lemma 3.1. *Let $\mathcal{T} \subseteq \mathbb{N}$ with $|\mathcal{T}| = \infty$. Given any sampling and recommendation rules such that $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$ for all $t \in \mathcal{T}$, then evaluating the stopping criterion (5) at each time $t \in \mathcal{T}$ with the threshold*

$$\beta(t, \delta) = 2K \ln(4 + \ln(t/K)) + KC_G \left(\frac{\ln(1/\delta)}{K} \right) \quad (6)$$

yields an (ε, δ) -PAC strategy. $\mathcal{C}_G(x) \approx x + \ln(x)$ as in (9).

Since this result holds for any sampling and recommendation rules satisfying one mild requirement, $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$ for all $t \in \mathcal{T}$, this leaves open the question on how to design those two rules to stop as early as possible. Algorithms that are agnostic to the choice of the candidate ε -optimal answer might have a higher expected sample complexity than the ones aiming at identifying the furthest answer.

Recommendation Rule Taking a greedy answer $z_t \in z^*(\mu_{t-1})$ is a direct choice. Thanks to its efficient implementation, using a greedy answer is the only computationally feasible recommendation rule for combinatorial or continuous answers sets. Unfortunately, when $z_F(\mu) \notin z^*(\mu)$, this approach leads to sub-optimal algorithms in terms of asymptotic sample complexity (Lemmas 2.2-2.3).

When Z is not too large or when we disregard the computational cost, a more careful choice than the greedy one alleviates this sub-optimality. The ε -optimal answers for which the GLR (l.h.s of (5)) is maximized are the *instantaneous furthest answers*

$$z_F(\mu_{t-1}, N_{t-1}) \stackrel{\text{def}}{=} \arg \max_{z \in \mathcal{Z}_\varepsilon(\mu_{t-1})} \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{t-1} - \lambda\|_{V_{N_{t-1}}}^2.$$

By definition, $z_F(\mu_{t-1}, N_{t-1})$ are the ε -optimal answers for which we have the most evidence against $\mathcal{H}_0 = \{\mu \in \neg_\varepsilon z\}$ at time t . At a lower computational cost than using a furthest answer for the current estimator $z_t \in z_F(\mu_{t-1})$, we will see that using an instantaneous furthest answer enjoys similar empirical performance (sample complexity). For all the above sets of candidate answers, the ties are broken arbitrarily. Empirically, we only observed singletons.

Dependence in K In linear bandits, when K is large, dependencies in K can be replaced by d (Lattimore & Szepesvári, 2020). The focus of our work is to highlight the importance of carefully choosing answers, therefore having K instead of d is a price we are willing to pay for simpler arguments. Prior works removed the K dependency in the analysis of game-based algorithms (Degenne et al., 2020a; Tirinzoni et al., 2020; Réda et al., 2021).

3.2. Modified BAI Algorithms

Modification Procedure Given any BAI algorithm for transductive linear Gaussian bandits, we modify it to use (5) as stopping rule while leaving the sampling rule unchanged. By Lemma 3.1, the resulting algorithm is an (ε, δ) -PAC strategy. For the recommendation rule, theory (Lemmas 2.2-2.3) and experiments (Figure 2) both suggest to use $z_t \in z_F(\mu_{t-1}, N_{t-1})$ instead of $z_t \in z^*(\mu_{t-1})$. We do not prove any theoretical guarantees on the sample complexity of the modified algorithms since such results depend heavily on each sampling rule.

BAI Benchmarks Lots of algorithms have been designed to tackle the BAI setting and we mention below the ones used in the experiments as benchmarks. Soare et al. (2014) proposed a static allocation design \mathcal{XY} -Static and its elimination-based improvement \mathcal{XY} -Adaptive, which are linked to a G -optimal design. In Xu et al. (2018), LinGapE was introduced as the first gap-based BAI algorithm. All the above BAI algorithms are not shown to be asymptotically optimal and depend on δ (except \mathcal{XY} -Static). Algorithm such as DKM (Degenne et al., 2019) and LinGame (Degenne et al., 2020a) are asymptotically optimal and their sampling rule does not depend on δ .

Other Stopping Rules For all BAI algorithm using a GLR-based stopping rule the ε -BAI stopping rule (5) is a natural modification. Some other non-GLR-based stopping rule also have a direct extension to ε -BAI. This is the case for the gap-based stopping rule for additive ε -optimality employed by LinGapE, where we can stop when the gap is smaller than ε instead of stopping when it is negative.

3.3. Experiments

We perform experiments to highlight the empirical performance of the modified BAI algorithms on additive ε -BAI problems. Moreover, we show that using $z_t \in z_F(\mu_{t-1}, N_{t-1})$ in (5) achieves lower empirical stopping time compared to $z_t \in z^*(\mu_{t-1})$, and outperforms the ε -gap stopping rule with $z_t \in z^*(\mu_{t-1})$. We consider linear bandits, i.e. $\mathcal{K} = \mathcal{Z}$, with $\mathcal{M} = \mathbb{R}^d$ and $(\varepsilon, \delta) = (0.05, 0.01)$, and perform 5000 runs. The stopping-recommendation pair is updated at each time t .

Hard Instance We adapt the usual hard instance studied in BAI for linear bandits to enforce the existence of multiple correct answers, i.e. $|\mathcal{Z}_\varepsilon(\mu)| > 1$. Taking $\mu = e_1$ with $e_a = (\mathbf{1}_{(a'=a)})_{a' \in [d]}$, the answers set is defined as $\mathcal{Z} = \{e_1, \dots, e_d, a_{d+1}, a_{d+2}\}$ where $a_{d+1} = \cos(\phi_1)e_1 + \sin(\phi_1)e_2 \in \mathcal{Z}_\varepsilon(\mu)$ and $a_{d+2} = \cos(\phi_2)e_1 + \sin(\phi_2)e_2 \notin \mathcal{Z}_\varepsilon(\mu)$. Considering $d = 2$, we use $\phi_1 = r_\varepsilon\theta_\varepsilon$ and $\phi_2 = (1 + r_\varepsilon)\theta_\varepsilon$ with $\theta_\varepsilon = \arccos(1 - \varepsilon)$ and $r_\varepsilon = 0.1$.

On this instance, the BAI algorithms without modification require on average 545 times more samples than compared to their modified version (Table 6 in Appendix F.2.1). The discrepancy is particularly striking since the hard instance for ε -BAI is even harder for BAI.




Figure 2. Empirical stopping time of the modified BAI algorithms with $z_t \in z_F(\mu_{t-1}, N_{t-1})$ on the hard instance (star is mean). “-G” denotes $z_t \in z^*(\mu_{t-1})$. “-O” denotes the ε -gap stopping rule for $z_t \in z^*(\mu_{t-1})$.

Figure 2 reveals that, for all modified BAI, considering an instantaneous furthest answer instead of a greedy answer leads to lower empirical stopping time. Their ratio is 0.92 on average. This matches the asymptotic observations in Figure 8(b) (Appendix F.2.2, equivalent of Figure 1(b) for the additive setting). The modified LinGapE using $z_t \in z_F(\mu_{t-1}, N_{t-1})$ outperforms the ε -gap extension of the original stopping rule, which is equivalent to using (5) with $z_t \in z^*(\mu_{t-1})$. While guided by the asymptotic regime, using $z_F(\mu_{t-1}, N_{t-1})$ instead of $z^*(\mu_{t-1})$ for the stopping-recommendation pair has practical utility in the moderate confidence regime with a 10% speed-up in terms of sample complexity.

4. $L\varepsilon$ BAI Algorithm

Leveraging the concept of furthest answer in the sampling rule, we present $L\varepsilon$ BAI (Linear ε -BAI), an asymptotically optimal algorithm for (ε, δ) -PAC best-answer identification in transductive linear bandits. It deals with both the multiplicative and the additive ε -optimality. Similarly to works

on linear bandits (Abbasi-yadkori et al., 2011; Soare et al., 2014), we assume that the set of parameters is bounded, i.e. there exists $M > 0$ such that for all $\mu \in \mathcal{M}$, $\|\mu\|_2 \leq M$.

Algorithm 1 $L\varepsilon$ BAI

Input: History \mathcal{F}_t , \mathcal{Z} -oracle $\mathcal{L}^\mathcal{Z}$ and learner $\mathcal{L}^\mathcal{K}$.

Output: Candidate ε -optimal answer \hat{z} .

Pull once each arm $a \in \mathcal{K}$, set $n_0 = K$ and $W_{n_0} = \mathbf{1}_K$;

for $t = n_0 + 1, \dots$ **do**

 Get $z_t \in z_F(\mu_{t-1}, N_{t-1})$;

 If (5) holds for z_t then return z_t ;

 Get $(\tilde{z}_t, w_t^{\mathcal{L}^\mathcal{K}})$ from $\mathcal{L}^\mathcal{Z} \times \mathcal{L}^\mathcal{K}$;

 Let $w_t = \frac{1}{tK} + (1 - \frac{1}{t})w_t^{\mathcal{L}^\mathcal{K}}$ and $W_t = W_{t-1} + w_t$;

 Closest alternative:

$\lambda_t \in \arg \min_{\lambda \in \neg_\varepsilon \tilde{z}_t} \|\mu_{t-1} - \lambda\|_{V_{w_t}}^2$;

 Optimistic gains: $\forall a \in \mathcal{K}$,

$U_t^a = (\|\mu_{t-1} - \lambda_t\|_{aa^\top} + \sqrt{c_{t-1}^a})^2$;

 Feed $\mathcal{L}^\mathcal{K}$ with gain $g_t(w) = (1 - \frac{1}{t})\langle w, U_t \rangle$;

 Pull $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - W_t^a$, observe $X_{t-1}^{a_t}$;

end for

Structure After pulling each arm once, at each round $t \geq n_0 + 1$, if the stopping condition (5) for the candidate answer $z_t \in z_F(\mu_{t-1}, N_{t-1})$, we return z_t ; else, the sampling rule returns an arm a_t to pull. Then, the statistics are updated based on this new observation.

Sampling Rule The algorithmic ingredients used in the sampling rule of $L\varepsilon$ BAI build upon the ones in LinGame (Degenne et al., 2020a). It is a saddle-point algorithm approximating a two-player zero-sum game. At each round $t \geq n_0 + 1$, if the algorithm hasn’t stopped yet, the agent chooses an ε -optimal answer and a pulling proportion over arms $(\tilde{z}_t, w_t^{\mathcal{L}^\mathcal{K}}) \in \mathcal{Z}_\varepsilon(\mu_{t-1}) \times \Delta_K$, where \tilde{z}_t can be different from z_t . A mild logarithmic forced exploration is added, i.e. $w_t = \frac{1}{tK}\mathbf{1}_K + (1 - \frac{1}{t})w_t^{\mathcal{L}^\mathcal{K}}$. The agent will play by combining a no-regret learner on Δ_K (e.g. AdaHedge of De Rooij et al. (2014)), denoted by $\mathcal{L}^\mathcal{K}$, and a \mathcal{Z} -oracle, denoted by $\mathcal{L}^\mathcal{Z}$. While Theorem 4.1 was proven for $\tilde{z}_s \in z_F(\mu_{s-1})$, we obtain similar empirical performance with the heuristic $\tilde{z}_s \in z_F(\mu_{s-1}, \mu_{s-1})$ at a much lower computational cost.

Given (\tilde{z}_t, w_t) from $\mathcal{L}^\mathcal{Z} \times \mathcal{L}^\mathcal{K}$, the nature plays the most confusing alternative parameter $\lambda_t \in \arg \min_{\lambda \in \neg_\varepsilon \tilde{z}_t} \|\mu_{t-1} - \lambda\|_{V_{w_t}}^2$. To update $\mathcal{L}^\mathcal{K}$, the agent uses gains $g_t(w) = (1 - \frac{1}{t})\langle w, U_t \rangle$ where the optimistic gains are defined for all $a \in \mathcal{K}$ as $U_t^a = (\|\mu_{t-1} - \lambda_t\|_{aa^\top} + \sqrt{c_{t-1}^a})^2$ with $c_{t-1}^a = \min \left\{ 2\beta(s^2, s^{2/3}) \|a\|_{V_{N_s}^{-1}}^2, 4M^2 L_{\mathcal{K}}^2 \right\}$. Under a good event, the quantity $\langle w, U_t \rangle$ is an upper bound on the unknown $\inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2$ (Lemma E.18).

Finally, a_t is obtained deterministically by *tracking*, i.e. $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - W_t^a$.

The computational cost is discussed in Appendix F.1. To obtain efficient implementations for combinatorial or large arms sets K , $\mathbf{L}\varepsilon\text{BAI}$ should be modified by using existing improvements for game-based algorithms (Tirinzone et al., 2020; Réda et al., 2021; Jourdan et al., 2021).

When $\varepsilon = 0$, $\mathbf{L}\varepsilon\text{BAI}$ is close to LinGame , but uses one learner instead of Z learners. Other differences are that LinGame uses regularization in the estimator and a stopping threshold featuring d .

4.1. Upper Bound

For both the multiplicative and the additive ε -optimality, Theorem 4.1 shows that $\mathbf{L}\varepsilon\text{BAI}$ yields an (ε, δ) -PAC asymptotically optimal algorithm. The proof sketch of Theorem 4.1 is inspired by the one of LinGame (Degenne et al., 2020a), hence we will only highlight the novel technical difficulties that had to be addressed.

Theorem 4.1. *Let $\mathcal{L}^{\mathcal{K}}$ with sub-linear regret (e.g. AdaHedge) and $\mathcal{L}^{\mathcal{Z}}$ returning $\tilde{z}_t \in z_F(\mu_{t-1})$. Using (6) as stopping threshold $\beta(t, \delta)$, $\mathbf{L}\varepsilon\text{BAI}$ yields an (ε, δ) -PAC algorithm and, for all $\mu \in \mathcal{M}$ such that $|z_F(\mu)| = 1$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_{\mu} [\tau_{\delta}]}{\ln(1/\delta)} \leq T_{\varepsilon}(\mu).$$

Technical Difficulties In BAI, we have $|z^*(\mu)| = 1$. The key property used in BAI proofs which does not hold in ε -BAI is that for all $z \neq z^*(\mu)$, μ belongs to the alternative $\neg_0 z$. The consequence of this is that whenever the answer used by the sampling rule \tilde{z}_t is wrong, the correct parameter belongs to $\neg_0 \tilde{z}_t$, hence the algorithm will sample in order to try and exclude that true parameter, which cannot succeed and will at some point correct the mistake. In ε -BAI we can have $\tilde{z}_t \neq z_F(\mu)$ while having $\mu \notin \neg_{\varepsilon} \tilde{z}_t$ and there is a priori no such self-correction mechanism to enforce that $\tilde{z}_t = z_F(\mu)$ after a while.

Our analysis reveals that a similar self-correction mechanism can be obtained for $\mathbf{L}\varepsilon\text{BAI}$. Let $\neg_F z$ be the *furthest alternative* to z , i.e. the set of parameters for which z is not the unique furthest answer. Intuitively, as it uses $\tilde{z}_t \in z_F(\mu_{t-1})$, $\mathbf{L}\varepsilon\text{BAI}$ samples to asymptotically exclude $\neg_F \tilde{z}_t$. Leveraging the logarithmic forced exploration, this cannot succeed when $\tilde{z}_t \neq z_F(\mu)$. Those two choices yield a self-correction mechanism for ε -BAI. More formally, we show that, under a good concentration event, the event $\tilde{z}_s \neq z_F(\mu)$ only happens a sub-linear number of times.

5. Related Work

Track-And-Stop (TaS) First introduced in Garivier & Kaufmann (2016) to solve BAI in MAB, TaS computes at each time step the optimal allocation, and then tracks it (with added forced exploration). When no close form solutions are available, i.e. when additional structure is considered, TaS-based algorithms suffer from intractable computational cost. TaS-based algorithms are asymptotically optimal for BAI and efficient to compute (Jedra & Proutière, 2020). Building on Frank-Wolfe algorithm, the computational efficient FWS has recently been introduced (Wang et al., 2021).

ε -BAI Algorithms Tackling ε -BAI in MAB for additive ε -optimality, ε -TaS (Garivier & Kaufmann, 2021) recommends $z_F(\mu_t, N_t)$ and uses the associated GLRT as stopping rule. The sampling rule computes $w_F(\mu_t)$ and then tracks it with added forced exploration. Addressing additive spectral bandits, SpectralTaS (Kocák & Garivier, 2021) recommends $z^*(\mu_t)$ and uses the GLRT associated with $z_F(\mu_t, N_t)$ for the stopping rule. For the sampling rule, a mirror ascent algorithm is run based on a super-gradient of a function depending on any ε -optimal answer. While the choice of the answer is not discussed, it is our understanding that a greedy answer is used (matching their candidate answer). When considering ε -BAI on the unit sphere, Jedra & Proutière (2020) recommend $z^*(\mu_t)$ and use the associated GLRT, however their sampling rule is uniform over a spanner.

Designed for the multiple-correct answer setting, Sticky TaS (Degenne & Koolen, 2019) is a modified TaS algorithm: at round t , they compute $\bigcup_{\mu' \in \mathcal{C}_t} z_F(\mu')$ where \mathcal{C}_t is a continuous confidence region around μ_t , and stick to one of those (given an arbitrary order). For some identification problems (e.g. Any Half-Space), it rewrites as computing $\bigcup_{\mu \in \mathcal{D}_t} z_F(\mu')$, where \mathcal{D}_t is discrete. There is no such rewriting for ε -BAI. While Sticky TaS can be implemented for \mathcal{D}_t , it is not feasible for \mathcal{C}_t . Experiments suggest that it performs on par with ε -TaS at a higher computational cost, i.e. solving the same optimization for each parameter in a confidence region.

6. Experiments

We show that $\mathbf{L}\varepsilon\text{BAI}$ has competitive empirical performance compared to existing ε -BAI algorithms, which are computationally expensive, and that using an instantaneous furthest answer is efficient both in terms of computational cost and sample complexity. Moreover, $\mathbf{L}\varepsilon\text{BAI}$ performs on par with the modified BAI algorithms, which are not asymptotically optimal, on hard and random instances.

As heuristic with lower computational cost (not supported by Theorem 4.1), the \mathcal{Z} -oracle in $\mathbf{L}\varepsilon\text{BAI}$ returns an instanta-

neous furthest answer, i.e. $\tilde{z}_t \in z_F(\mu_{t-1}, N_{t-1})$. The experiments below are considering the multiplicative ε -optimality (Appendix F.2.1-F.2.2 for supplementary ones). We use the same experimental setup as in Section 3.3. On the 5000 runs, we report the standard deviation of means by using sub-samples of 100 runs.

Table 1. Empirical stopping time ($\pm \sigma$) on the hard instance with $\mathcal{K} = \{e_1, e_2\}$.

	$z^*(\mu_{t-1})$	$z_F(\mu_{t-1})$	$z_F(\mu_{t-1}, N_{t-1})$
$L\varepsilon$ BAI	416 (± 13)	383 (± 16)	381 (± 17)
ε -TaS	400 (± 14)	371 (± 15)	371 (± 15)
Fixed	401 (± 14)	374 (± 14)	374 (± 14)
Uniform	492 (± 16)	450 (± 17)	449 (± 17)

ε -BAI and Candidate Answer Even when K is small, algorithms based on solving the optimization problem $(z_F(\mu), w_F(\mu))$ are intractable, i.e. ε -TaS or recommending the furthest answer. We evaluate their performance empirically on the hard instance with $\mathcal{K} = \{e_1, e_2\}$, and discretize uniformly Δ_2 with 500 vectors.

In Table 1, we combine and compare four ε -BAI sampling rules with three candidate answers for the stopping rule (5). Comparing the rows of Table 1 reveals that $L\varepsilon$ BAI performs on par with ε -TaS and the ‘‘oracle’’ *fixed* algorithm, which tracks the unknown optimal allocation $w_F(\mu)$. It also consistently outperforms uniform sampling ($\approx 85\%$).

Comparing the columns of Table 1, we see that using a greedy answer is consistently worse than a (instantaneous) furthest answer, with a ratio of stopping time being on average 0.92 (coherent with Figure 1(b)). Moreover, it highlights that using an instantaneous furthest answer achieves similar performance as a furthest answer at a lower computational cost. In the following experiments, the stopping-recommendation pair is (5) combined with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

Modified BAI Algorithms Figure 3 compares $L\varepsilon$ BAI with the modified BAI algorithms, all using the same stopping-recommendation pair. We see that $L\varepsilon$ BAI slightly outperforms the modified LinGapE and \mathcal{XY} -Adaptive, performs better than the modified LinGame and \mathcal{XY} -Static and is on par with the ‘‘oracle’’ *fixed* algorithm. Uniform sampling and the modified DKM perform poorly.

Random Instances To assess the impact of higher dimensions, random instances are considered (one per run). For the answer set, 19 vectors $(a_k)_{k \in [19]}$ are uniformly drawn from $\mathbb{S}^{d-1} \stackrel{\text{def}}{=} \{a \in \mathbb{R}^d : \|a\|_2 = 1\}$ and set $\mu = a_1$. To enforce multiple correct answers, a modification of the greedy




Figure 3. Empirical stopping time on the hard instance ($\mathcal{K} = \mathcal{Z}$). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

answer is added such that $a_{20,i} = a_{1,i}$ for $i \neq i_0$ and $a_{20,i_0} = \frac{1 - \|\mu\|_2^2 + \mu_{i_0}^2 - r_\varepsilon \varepsilon}{\mu_{i_0}}$ where $i_0 = \arg \min_{i \in [d]} \mu_i$ and $r_\varepsilon = 0.1$. Those instances are motivated by a practical BAI example where a modified/corrupted version of the unique correct answer exists. Seeing the problem as an ε -BAI one allows to return an ε -optimal answer, while avoiding wasteful queries required by BAI algorithms (Table 4).

In Figure 4, $L\varepsilon$ BAI shows similar empirical performance with modified BAI algorithms. Even though it is outperformed by the modified LinGapE, $L\varepsilon$ BAI is almost twice as fast as the modified \mathcal{XY} -Adaptive and appears to be slightly more robust than the modified LinGame to increasing dimension.

7. Conclusion

In (ε, δ) -PAC best-answer identification for transductive linear bandits, we have shown that the choice of the candidate ε -optimal answer is important for the sample complexity. Using an instantaneous furthest answer as candidate answer, we proposed a simple procedure to adapt existing BAI algorithms for ε -BAI problems. Leveraging it in the sampling rule as well, we introduced $L\varepsilon$ BAI which is asymptotically optimal and has competitive empirical performance.

Computing the furthest answer requires solving the closest alternative sub-problem $|\mathcal{Z}_\varepsilon(\mu_{t-1})|$ times. While that number is small (in particular much less than Z) in the examples we considered, that computation can become an issue if many different answers are close to each other. If we extend the setting to continuous answers, the computation of the furthest answer by iterating becomes unfeasible. The question of finding an ε -optimal point of a reward function in a non-finite set is the general question of optimization, which is central to many areas of machine learning. Extending the problem-dependent approach of the bandit framework to




Figure 4. Empirical stopping time on random instances ($\mathcal{K} = \mathcal{Z}$) for $d \in \{6, 8, 10, 12\}$ (from left to right). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

that setting is an interesting research direction.

Studying the performance of ε -BAI algorithms on BAI tasks is another avenue for future work. ε -BAI algorithms don't commit to the greedy answer unlike BAI methods, and we might have $z_F(\mu_{t-1}, N_{t-1}) = z^*(\mu)$ before $z^*(\mu_{t-1}) = z^*(\mu)$. For this reason we conjecture that using an ε -BAI method with a well-tuned decreasing sequence $(\varepsilon_t)_{t \in \mathbb{N}^*}$ might outperform BAI algorithms on best arm tasks.

Finally, since the existence of a tight finite-time lower bound for multiple-correct answer setting is still an open problem, it remains unclear how to assess the theoretical performance of algorithms in this regime. We believe that, once derived, this lower bound would reveal the existence of strong moderate confidence terms (independent of δ) affecting the sample complexity, which could then be used to design ε -BAI algorithms with theoretical guarantees in both regime.

Acknowledgements

Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>).

References

- Abbasi-yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems*, 2011.
- Al Marjani, A., Kocak, T., and Garivier, A. On the complexity of All ε -Best Arms Identification. 2022.
- Audibert, J., Bubeck, S., and Munos, R. Best arm identification in multi-armed bandits. In *The 23rd Conference on Learning Theory (COLT)*, 2010.
- Bubeck, S., Munos, R., and Stoltz, G. Pure exploration in multi-armed bandits problems. In *Algorithmic Learning Theory (ALT)*, 2009.
- Chen, L., Li, J., and Qiao, M. Towards instance optimal bounds for best arm identification. In *Proceedings of the 2017 Conference on Learning Theory*, 2017.
- Chen, S., Lin, T., King, I., Lyu, M. R., and Chen, W. Combinatorial Pure Exploration of Multi-Armed Bandits. In *Advances in Neural Information Processing Systems 27*, 2014.
- Chernoff, H. Sequential Design of Experiments. *The Annals of Mathematical Statistics*, 30(3):755–770, 1959.
- Cheshire, J., Menard, P., and Carpentier, A. Problem Dependent View on Structured Thresholding Bandit Problems. In *International Conference on Machine Learning*, 2021.
- De Rooij, S., Van Erven, T., Grünwald, P. D., and Koolen, W. M. Follow the leader if you can, hedge if you must. *The Journal of Machine Learning Research*, 15(1):1281–1316, 2014.
- Degenne, R. and Koolen, W. M. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems*, 2019.
- Degenne, R., Koolen, W. M., and Ménard, P. Non-Asymptotic Pure Exploration by Solving Games. In *Advances in Neural Information Processing Systems*, 2019.
- Degenne, R., Ménard, P., Shang, X., and Valko, M. Gamification of pure exploration for linear bandits. In *International Conference on Machine Learning*, 2020a.
- Degenne, R., Shao, H., and Koolen, W. Structure Adaptive Algorithms for Stochastic Bandits. In *International Conference on Machine Learning*, 2020b.
- Even-Dar, E., Mannor, S., and Mansour, Y. PAC bounds for multi-armed bandit and markov decision processes. In *Computational Learning Theory (COLT)*, 2002.
- Even-Dar, E., Mannor, S., and Mansour, Y. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7:1079–1105, 2006.
- Fiez, T., Jain, L., Jamieson, K. G., and Ratliff, L. J. Sequential experimental design for transductive linear bandits.

- In *Advances in Neural Information Processing Systems*, 2019.
- Gabillon, V., Ghavamzadeh, M., and Lazaric, A. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *Advances in Neural Information Processing Systems 25*, 2012.
- Garivier, A. and Kaufmann, E. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory (COLT)*, 2016.
- Garivier, A. and Kaufmann, E. Nonasymptotic sequential tests for overlapping hypotheses applied to near-optimal arm identification in bandit models. *Sequential Analysis*, 40(1):61–96, 2021.
- Hong, L. J., Fan, W., and Luo, J. Review on ranking and selection: A new perspective. *Frontiers of Engineering Management*, 8(3):321–343, 2021.
- Jamieson, K. and Nowak, R. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *Annual Conference on Information Sciences and Systems (CISS)*, 2014.
- Jedra, Y. and Proutière, A. Optimal best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, 2020.
- Jourdan, M., Mutn , M., Kirschner, J., and Krause, A. Efficient pure exploration for combinatorial bandits with semi-bandit feedback. In *Algorithmic Learning Theory (ALT)*, 2021.
- Jun, K.-S., Jain, L., Nassif, H., and Mason, B. Improved Confidence Bounds for the Linear Logistic Model and Applications to Bandits. In *International Conference on Machine Learning*, 2021.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2012.
- Katz-Samuels, J. and Jamieson, K. The true sample complexity of identifying good arms. In *International Conference on Artificial Intelligence and Statistics*, 2020.
- Katz-Samuels, J. and Scott, C. Top Feasible Arm Identification. In *International Conference on Artificial Intelligence and Statistics*, 2019.
- Katz-Samuels, J., Jain, L., Jamieson, K. G., et al. An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. *Advances in Neural Information Processing Systems*, 2020.
- Kaufmann, E. and Koolen, W. M. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22 (246):1–44, 2021.
- Kaufmann, E., Capp , 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.
- Kaufmann, E., Koolen, W. M., and Garivier, A. Sequential test for the lowest mean: From thompson to murphy sampling. In *Advances in Neural Information Processing Systems*, 2018.
- Koc k, T. and Garivier, A. Epsilon best arm identification in spectral bandits. In *International Joint Conference on Artificial Intelligence (IJCAI)*, 2021.
- Lattimore, T. and Szepesv ri, C. *Bandit Algorithms*. Cambridge University Press, 2020.
- Locatelli, A., Gutzeit, M., and Carpentier, A. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning (ICML)*, 2016.
- Mannor, S. and Tsitsiklis, J. N. The Sample Complexity of Exploration in the Multi-Armed Bandit Problem. *The Journal of Machine Learning Research*, 5:623–648, 2004.
- Mason, B., Jain, L., Tripathy, A., and Nowak, R. Finding all ε -good arms in stochastic bandits. *Advances in Neural Information Processing Systems*, 2020.
- R da, C., Kaufmann, E., and Delahaye-Duriez, A. Top-m identification for linear bandits. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2021.
- R da, C., Tirinzoni, A., and Degenne, R. Dealing with misspecification in fixed-confidence linear top-m identification. *Advances in Neural Information Processing Systems*, 2021.
- Robbins, H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58:527–535, 1952.
- Sabato, S. Epsilon-Best-Arm Identification in Pay-Per-Reward Multi-Armed Bandits. In *Advances in Neural Information Processing Systems*, 2019.
- Sentenac, F., Yi, J., Calauzenes, C., Perchet, V., and Vojnovic, M. Pure Exploration and Regret Minimization in Matching Bandits. In *International Conference on Machine Learning*, 2021.

- Soare, M., Lazaric, A., and Munos, R. Best-arm identification in linear bandits. In *Advances in Neural Information Processing Systems*, 2014.
- Tirinzi, A., Pirotta, M., Restelli, M., and Lazaric, A. An asymptotically optimal primal-dual incremental algorithm for contextual linear bandits. *Advances in Neural Information Processing Systems*, 2020.
- Wang, P.-A., Tzeng, R.-C., and Proutiere, A. Fast pure exploration via frank-wolfe. *Advances in Neural Information Processing Systems*, 2021.
- Xu, L., Honda, J., and Sugiyama, M. A fully adaptive algorithm for pure exploration in linear bandits. In *International Conference on Artificial Intelligence and Statistics*, 2018.
- Zaki, M., Mohan, A., and Gopalan, A. Explicit Best Arm Identification in Linear Bandits Using No-Regret Learners. *arXiv:2006.07562 [cs, stat]*, June 2020.

A. Outline

The appendices are organized as follows:

- Notations are summarized in Appendix B.
- In Appendix C, we prove all the results presented in Section 2 on the (greedy) characteristic time.
- The proof of Lemma 3.1 on the stopping-recommendation pairs is given in Appendix D.
- The full proof of Theorem 4.1 is given in Appendix E.
- The implementation details and additional experiments are presented in Appendix F.

Table 2. Notation for the setting

Notation	Type	Description
$\mathcal{M} \subseteq \mathbb{R}^d$		Set of possible parameters subset of \mathbb{R}^d , $\max_{\lambda \in \mathcal{M}} \ \lambda\ _2 \leq M$
μ	\mathcal{M}	Bandit mean parameter, $(\mu_i)_{i \in [d]}$
$\mathcal{K} \subseteq \mathbb{R}^d$		Set of arms of cardinality K , an arm is denoted by a and $L_{\mathcal{K}} \stackrel{\text{def}}{=} \max_{a \in \mathcal{K}} \ a\ _2$
μ^a	\mathbb{R}	Mean reward of arm a , $\mu^a \stackrel{\text{def}}{=} \langle \mu, a \rangle$
$\mathcal{Z} \subseteq \mathbb{R}^d$		Set of answers of cardinality Z , an arm is denoted by z
ε	\mathbb{R}_+^*	Approximation error
$\mathcal{Z}_\varepsilon(\mu) \subseteq \mathcal{Z}$		Set of ε -optimal answers, with $\cdot_{\{\text{add}, \text{mul}\}}$
$\neg_\varepsilon z \subseteq \mathcal{M}$		Set of alternative parameters to answer z , with $\cdot_{\{\text{add}, \text{mul}\}}$
$T_\varepsilon(\mu)$	\mathbb{R}_+^*	Characteristic time for μ , with $\cdot_{\{\text{add}, \text{mul}\}}$
$z^*(\mu) \subseteq \mathcal{Z}$		Greedy answers for μ
$z_F(\mu)$	\mathcal{Z}	Furthest answer for μ , with $\cdot_{\{\text{add}, \text{mul}\}}$
$z_F(\mu, N) \subseteq \mathcal{Z}$		Instantaneous furthest answers for μ and allocation N , with $\cdot_{\{\text{add}, \text{mul}\}}$

Table 3. Notation for algorithms

Notation	Type	Description
δ	$(0, 1)$	Confidence parameter
a_t	\mathcal{K}	Arm sampled at time t
$X_t^{a_t}$	\mathbb{R}	Observation at time t , $X_t^{a_t} \sim \mathcal{N}(\mu^{a_t}, \sigma^2)$ where $\sigma^2 = 1$
\mathcal{F}_t		History up to time t , $\sigma(a_1, X_1^{a_1}, \dots, a_t, X_t^{a_t})$
μ_t	\mathbb{R}^d	Maximum likelihood estimator, $\mu_t = V_{N_t}^{-1} \sum_{s=1}^t X_s^{a_s} a_s$
z_t	\mathcal{Z}	Candidate answer at time t
τ_δ	\mathbb{N}	Stopping time for confidence δ
N_t	$(\mathbb{R}_+)^K$	Empirical count of sampled arms at time t
w_t, W_t	$(\mathbb{R}_+)^K$	Pulling distribution over arms and its cumulative sum at time t
$\beta(t, \delta)$	$\mathbb{N} \times (0, 1) \rightarrow \mathbb{R}_+^*$	Stopping threshold at time t for confidence δ
$f(t)$	$\mathbb{N} \rightarrow \mathbb{R}_+^*$	Exploration bonus at time t
λ_t	\mathcal{M}	Most confusing alternative parameter player by nature at time t
U_t	$(\mathbb{R}_+)^K$	Optimistic gain at time t

B. Notation

We recall some commonly used notations: the set of integers $[n] \stackrel{\text{def}}{=} \{1, \dots, n\}$, the interval of integers $\llbracket a, b \rrbracket$, the euclidean inner-product $\langle x, y \rangle \stackrel{\text{def}}{=} \sum_{i \in [d]} x_i y_i$, the design matrix $V_w \stackrel{\text{def}}{=} \sum_{a \in \mathcal{K}} w^a a a^\top$ for an allocation over arms w ,

the norm $\|x\|_V \stackrel{\text{def}}{=} \sqrt{x^\top V x}$ for a positive definite matrix V , the closure \bar{X} of a set X , the cartesian product $\prod_{i \in [n]} X_i$ between sets $(X_i)_{i \in [n]}$, the set of probability distributions $\mathcal{P}(X)$ over X and the n -dimensional probability simplex $\Delta_n \stackrel{\text{def}}{=} \{x \in \mathbb{R}^n \mid x \geq 0, \langle \mathbf{1}_n, x \rangle = 1\}$. The Landau's notations o , \mathcal{O} and Θ are used. In Table 2, we summarize problem-specific notation. Table 3 gathers notation for the algorithms.

C. Proofs Related to the Characteristic Time

In Appendix C.1, we recall the formula of the generalized log-likelihood ratio for Gaussian bandits. The asymptotic lower bound on the expected sample complexity (Theorem 2.1) is stated in Appendix C.2. In Appendix C.3, when $\bar{\mathcal{M}} = \mathbb{R}^d$, we derive explicit formulas on $T_\varepsilon(\mu)^{-1}$ (Lemma C.1) and exhibit hard BAI instances that can be solved as an ε -BAI problem (Lemma C.3). In Appendix C.4, we prove the results related to the greedy characteristic time (Lemma 2.2-2.3).

C.1. Likelihood Ratio for Gaussian Bandits

For Gaussian bandits such that $\sigma^2 = 1$, the Kullback-Leibler (KL) divergence between two bandits, with mean parameters μ and λ , on the arm $a \in \mathcal{K}$ is:

$$d_{\text{KL}}(\mu^a, \lambda^a) = \frac{1}{2}(\mu^a - \lambda^a)^2 = \frac{1}{2}(a^\top(\mu - \lambda))^2$$

The generalized log-likelihood ratio (GLR) between the whole model space \mathcal{M} and a subset $\Lambda \subseteq \mathcal{M}$ is

$$\text{GLR}_t^{\mathcal{M}}(\Lambda) \stackrel{\text{def}}{=} \ln \frac{\sup_{\tilde{\mu} \in \mathcal{M}} \mathcal{L}_{\tilde{\mu}}(X_1^{a_1}, \dots, X_t^{a_t})}{\sup_{\lambda \in \Lambda} \mathcal{L}_\lambda(X_1^{a_1}, \dots, X_t^{a_t})}$$

where $\mathcal{L}_\lambda(X_1^{a_1}, \dots, X_t^{a_t})$ denotes the likelihood of the observations $X_1^{a_1}, \dots, X_t^{a_t}$ for a bandit with parameter λ .

Given two mean parameters $(\theta, \lambda) \in \mathcal{M}^2$, we have

$$\ln \frac{\mathcal{L}_\theta(X_1^{a_1}, \dots, X_t^{a_t})}{\mathcal{L}_\lambda(X_1^{a_1}, \dots, X_t^{a_t})} = \sum_{a \in \mathcal{K}} N_t^a (d_{\text{KL}}(\mu_t^a, \lambda^a) - d_{\text{KL}}(\mu_t^a, \theta^a))$$

where $\mu_t \stackrel{\text{def}}{=} V_{N_t-1}^{-1} \sum_{s=1}^{t-1} X_s^{a_s} a_s$ is a sufficient statistic of our observations. Note that, when $\mu_t \in \mathcal{M}$, μ_t coincide with the MLE $\tilde{\mu}_t$ defined as:

$$\tilde{\mu}_t = \arg \min_{\lambda \in \mathcal{M}} \sum_{a \in \mathcal{K}} N_t^a d_{\text{KL}}(\mu_t^a, \lambda^a) = \arg \min_{\lambda \in \mathcal{M}} \frac{1}{2} \|\mu_t - \lambda\|_{V_{N_t}}^2$$

The GLR for the set Λ is

$$\begin{aligned} \text{GLR}_t^{\mathcal{M}}(\Lambda) &= \min_{\lambda \in \Lambda} \sum_{a \in \mathcal{K}} N_t^a d_{\text{KL}}(\mu_t^a, \lambda^a) - \sum_{a \in \mathcal{K}} N_t^a d_{\text{KL}}(\mu_t^a, \tilde{\mu}_t^a) \\ &= \min_{\lambda \in \Lambda} \frac{1}{2} \|\mu_t - \lambda\|_{V_{N_t}}^2 - \frac{1}{2} \|\mu_t - \tilde{\mu}_t\|_{V_{N_t}}^2 \end{aligned}$$

C.2. Proof of Theorem 2.1

Theorem (Theorem 2.1, Corollary of Theorem 1 in (Degenne & Koolen, 2019)). For all (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_\varepsilon(\mu)$$

where the inverse of the characteristic time is

$$T_\varepsilon(\mu)^{-1} \stackrel{\text{def}}{=} \max_{z \in \mathcal{Z}_\varepsilon(\mu)} \max_{w \in \Delta_{\mathcal{K}}} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2$$

Proof. Since ε -BAI for Gaussian distributions is a special case of identification in bandits with multiple correct answers and sub-Gaussian distributions, Theorem 1 in [Degenne & Koolen \(2019\)](#) applies to our setting. Therefore, we obtain that: for all (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T^*(\mu)$$

where

$$T^*(\mu)^{-1} \stackrel{\text{def}}{=} \max_{i \in i^*(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg i} \sum_{a \in \mathcal{K}} w^a d_{\text{KL}}(\mu^a, \lambda^a)$$

with $i^*(\mu)$ is the set of correct answers for μ and $\neg i \stackrel{\text{def}}{=} \{\mu \in \overline{\mathcal{M}} : i \notin i^*(\mu)\}$.

To conclude, we need to specify to our setting and some simple manipulations. First, under the Gaussian assumptions with $\sigma^2 = 1$, notice that $\mu^a = \langle a, \mu \rangle$ and $d_{\text{KL}}(\mu^a, \lambda^a) = \frac{1}{2} \langle a, \mu - \lambda \rangle^2 = \frac{1}{2} \|\mu - \lambda\|_{a a^\top}^2$ yield that $\sum_{a \in \mathcal{K}} w^a d_{\text{KL}}(\mu^a, \lambda^a) = \frac{1}{2} \|\mu - \lambda\|_{V_w}^2$.

When considering additive ε -optimality, the set of correct answer $i^*(\mu)$ is $\mathcal{Z}_\varepsilon^{\text{add}}(\mu) = \{z \in \mathcal{Z} : \langle \mu, z \rangle \geq \max_{z \in \mathcal{Z}} \langle \mu, z \rangle - \varepsilon\}$, and $\neg i$ corresponds to $\neg_\varepsilon^{\text{add}} z$. For the multiplicative ε -optimality, the set of correct answer $i^*(\mu)$ is $\mathcal{Z}_\varepsilon^{\text{mul}}(\mu) = \{z \in \mathcal{Z} : \langle \mu, z \rangle \geq (1 - \varepsilon) \max_{z \in \mathcal{Z}} \langle \mu, z \rangle\}$, and $\neg i$ corresponds to $\neg_\varepsilon^{\text{mul}} z$. Note that for the multiplicative setting, we assume $\max_{z \in \mathcal{Z}} \langle \mu, z \rangle > 0$ and $\varepsilon \geq 0$ (or $\max_{z \in \mathcal{Z}} \langle \mu, z \rangle < 0$ and $\varepsilon \leq 0$), while for the additive setting, we can consider $\varepsilon \geq 0$ and $\max_{z \in \mathcal{Z}} \langle \mu, z \rangle \in \mathbb{R}$. Using the correspondences between sets for each notion of ε -optimality, Theorem 1 in [Degenne & Koolen \(2019\)](#) rewrites as: for all (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_\varepsilon(\mu)$$

where

$$T_\varepsilon^{\text{add}}(\mu)^{-1} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{add}}(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon^{\text{add}} z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2$$

$$T_\varepsilon^{\text{mul}}(\mu)^{-1} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{mul}}(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon^{\text{mul}} z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2$$

□

C.3. Explicit Formulas and Hard BAI Instances

In [Appendix C.3.1](#), we derive explicit formulas on $T_\varepsilon(\mu)^{-1}$ ([Lemma C.1](#)), while [Appendix C.3.2](#) proved that there exists hard BAI instances that can be solved as an ε -BAI problem ([Lemma C.3](#)).

C.3.1. EXPLICIT FORMULAS

[Lemma C.1](#) shows more explicit formulas for $T_\varepsilon(\mu)$ when $\overline{\mathcal{M}} = \mathbb{R}^d$.

Lemma C.1. Assume $\overline{\mathcal{M}} = \mathbb{R}^d$,

$$\frac{2}{T_\varepsilon^{\text{add}}(\mu)} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{add}}(\mu)} \max_{w \in \Delta_K} \min_{x \in \mathcal{Z} \setminus \{z\}} \frac{(\varepsilon + \langle \mu, z - x \rangle)^2}{\|z - x\|_{V_w^\dagger}^2}$$

$$\frac{2}{T_\varepsilon^{\text{mul}}(\mu)} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{mul}}(\mu)} \max_{w \in \Delta_K} \min_{x \in \mathcal{Z} \setminus \{z\}} \frac{\langle \mu, z - (1 - \varepsilon)x \rangle^2}{\|z - (1 - \varepsilon)x\|_{V_w^\dagger}^2}$$

where V_w^\dagger is the Moore-Penrose pseudo-inverse of V_w . The associated alternative parameters are

$$\lambda_\varepsilon^{\text{add}}(\mu, z, w, x) = \mu - \frac{\varepsilon + \langle \mu, z - x \rangle}{\|z - x\|_{V_w^\dagger}^2} V_w^\dagger (z - x),$$

$$\lambda_\varepsilon^{\text{mul}}(\mu, z, w, x) = \mu - \frac{\langle \mu, z - (1 - \varepsilon)x \rangle}{\|z - (1 - \varepsilon)x\|_{V_w^\dagger}^2} V_w^\dagger (z - (1 - \varepsilon)x).$$

For the multiplicative ε -optimality, we recognize a term similar to the characteristic time $T_0(\mu, \mathcal{Z})$ of BAI for transductive bandits with answer set \mathcal{Z} . Defining $\mathcal{Z}_\varepsilon^z \stackrel{\text{def}}{=} \{z\} \cup \{(1-\varepsilon)x : x \in \mathcal{Z} \setminus \{z\}\}$, we see that $T_\varepsilon^{\text{mul}}(\mu) = \min_{z \in \mathcal{Z}^{\text{mul}}(\mu)} T_0(\mu, \mathcal{Z}_\varepsilon^z)$. Therefore, the complexity of the multiplicative setting is equals to the easiest BAI for modified transductive bandits, in which \mathcal{K} is unchanged and $\mathcal{Z}_\varepsilon^z$ is the set of answers. Since $T_\varepsilon^{\text{mul}}(\mu) = T_0(\mu, \mathcal{Z}_\varepsilon^{z_F^{\text{mul}}(\mu)})$, the unique correct answer that has to be identified in this modified BAI is the furthest answer $z_F^{\text{mul}}(\mu)$.

The proof of Lemma C.1 uses Lemma C.2, a known result in the literature which we prove for completeness, and the fact that $\neg_\varepsilon z$ is a union over a finite number of half-spaces.

Lemma C.2. For all $\mu, \lambda \in \mathbb{R}^d$, $y \in \mathbb{R}^d$, $x \in \mathbb{R}$, we have

$$\inf_{\lambda \in \mathbb{R}^d : x - \langle \lambda, y \rangle \leq 0} \|\mu - \lambda\|_{V_w}^2 = \begin{cases} 0 & \text{if } x - \langle \mu, y \rangle \leq 0 \text{ or } y \notin \text{Im}(V_w) \\ \left(\frac{x - \langle \mu, y \rangle}{\|y\|_{V_w^\dagger}} \right)^2 & \text{else} \end{cases}$$

where V_w^\dagger is the Moore-Penrose pseudo-inverse of V_w . It is achieved at

$$\lambda = \begin{cases} \mu & \text{if } x - \langle \mu, y \rangle \leq 0 \\ \mu + \alpha u & \text{if } y \notin \text{Im}(V_w), \text{ s.t. } (u, \alpha) \in \text{Ker}(V_w) \times \mathbb{R}, x - \langle \mu + \alpha u, y \rangle \leq 0 \\ \mu + \frac{x - \langle \mu, y \rangle}{\|y\|_{V_w^\dagger}^2} V_w^\dagger y & \text{else} \end{cases}$$

Proof. Since V_w is self-adjoint, $\text{Ker}(V_w)^\perp = \text{Im}(V_w)$. If $y \notin \text{Im}(V_w)$, then there exists $u \in \text{Ker}(V_w)$ with $\langle u, y \rangle \neq 0$. There exists also $\alpha \in \mathbb{R}$ such that $x - \langle \mu + \alpha u, y \rangle \leq 0$ and $\|\mu - (\mu + \alpha u)\|_{V_w} = 0$, hence the value of the objective is 0.

Otherwise, $y \in \text{Im}(V_w)$. We can restrict the problem to an infimum over $\lambda - \mu \in \text{Im}(V_w)$ since for any λ satisfying the constraint, its projection on that space also satisfies the constraint and has lower objective value.

Let V_w^\dagger be the Moore-Penrose pseudo-inverse of V_w . Restricted to $\text{Im}(V_w)$, this is a true inverse.

By using the Lagrangian of the problem, we obtain the following:

$$\begin{aligned} \inf_{\lambda - \mu \in \text{Im}(V_w) : x - \langle \lambda, y \rangle \leq 0} \|\mu - \lambda\|_{V_w}^2 &= \sup_{\alpha \geq 0} \inf_{\lambda - \mu \in \text{Im}(V_w)} \|\mu - \lambda\|_{V_w}^2 + \alpha(x - \langle \lambda, y \rangle) \\ &= \sup_{\alpha \geq 0} \alpha(x - \langle \mu, y \rangle) - \alpha^2 \frac{\|y\|_{V_w^\dagger}^2}{4} \\ &= \begin{cases} 0 & \text{if } x - \langle \mu, y \rangle \leq 0 \\ \left(\frac{x - \langle \mu, y \rangle}{\|y\|_{V_w^\dagger}} \right)^2 & \text{else} \end{cases} \end{aligned}$$

where the infimum in the first equality is reached at $\lambda = \mu + \frac{1}{2}\alpha V_w^\dagger y$ and the supremum in the last equality is reached at $\alpha = 2 \frac{x - \langle \mu, y \rangle}{\|y\|_{V_w^\dagger}^2}$ if $x - \langle \mu, y \rangle \geq 0$ and at $\alpha = 0$ else. \square

We are now ready to prove Lemma C.1.

Proof. Let $w \in \Delta_K$. We will conduct the proof of both results in parallel. When $\overline{\mathcal{M}} = \mathbb{R}^d$, by definitions of $\neg_\varepsilon^{\text{add}} z$ and $\neg_\varepsilon^{\text{mul}} z$ which involve the closure of a set that can be written as a union over a finite number of half-spaces, we can rewrite:

$$\begin{aligned} \neg_\varepsilon^{\text{add}} z &= \bigcup_{x \neq z} \{\lambda \in \mathbb{R}^d : \langle \lambda, z - x \rangle + \varepsilon \leq 0\} & \text{and} & \quad \inf_{\lambda \in \neg_\varepsilon^{\text{add}} z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2 = \min_{x \neq z} \inf_{\lambda \in \mathbb{R}^d : \langle \lambda, z - x \rangle + \varepsilon \leq 0} \|\mu - \lambda\|_{V_w}^2 \\ \neg_\varepsilon^{\text{mul}} z &= \bigcup_{x \neq z} \{\lambda \in \mathbb{R}^d : \langle \lambda, z - (1 - \varepsilon)x \rangle \leq 0\} & \text{and} & \quad \inf_{\lambda \in \neg_\varepsilon^{\text{mul}} z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2 = \min_{x \neq z} \inf_{\lambda \in \mathbb{R}^d : \langle \lambda, z - (1 - \varepsilon)x \rangle \leq 0} \|\mu - \lambda\|_{V_w}^2 \end{aligned}$$

Since $z^*(\mu) = \arg \max_{z \in \mathcal{Z}} \langle \mu, z \rangle$, for all $z \in \mathcal{Z}_\varepsilon^{\text{add}}(\mu)$ (resp. $z \in \mathcal{Z}_\varepsilon^{\text{mul}}(\mu)$) and all $x \neq z$, we have by definition $\langle \mu, z - x \rangle + \varepsilon \geq \langle \mu, z - z^*(\mu) \rangle + \varepsilon \geq 0$ (resp. $\langle \mu, z - (1 - \varepsilon)x \rangle \geq \langle \mu, z - (1 - \varepsilon)x \rangle \geq 0$). Using Lemma C.2 with $\tilde{y} = x - z$ and $\tilde{x} = \varepsilon$ (resp. $\tilde{y} = (1 - \varepsilon)x - z$ and $\tilde{x} = 0$), we obtain directly:

$$2T_\varepsilon^{\text{add}}(\mu)^{-1} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{add}}(\mu)} \max_{w \in \Delta_K} \min_{x \neq z} \left(\frac{\varepsilon + \langle \mu, z - x \rangle}{\|z - x\|_{V_w^\dagger}} \right)^2$$

$$2T_\varepsilon^{\text{mul}}(\mu)^{-1} = \max_{z \in \mathcal{Z}_\varepsilon^{\text{mul}}(\mu)} \max_{w \in \Delta_K} \min_{x \neq z} \left(\frac{\langle \mu, z - (1 - \varepsilon)x \rangle}{\|z - (1 - \varepsilon)x\|_{V_w^\dagger}} \right)^2$$

The associated alternative parameters are:

$$\lambda_\varepsilon^{\text{add}}(\mu, z, w, x) = \mu - \frac{\varepsilon + \langle \mu, z - x \rangle}{\|z - x\|_{V_w^\dagger}^2} V_w^\dagger(z - x)$$

$$\lambda_\varepsilon^{\text{mul}}(\mu, z, w, x) = \mu - \frac{\langle \mu, z - (1 - \varepsilon)x \rangle}{\|z - (1 - \varepsilon)x\|_{V_w^\dagger}^2} V_w^\dagger(z - (1 - \varepsilon)x)$$

□

Link with BAI First, note that $\neg_\varepsilon z \subseteq \neg_0 z$ for all $z \in \mathcal{Z}$, yields $T_\varepsilon(\mu) \leq T_0(\mu)$ for all $\mu \in \mathcal{M}$. When considering the BAI for transductive linear bandits with Gaussian distribution, the characteristic time that lower bounds the expected sample complexity has the following form when $\overline{\mathcal{M}} = \mathbb{R}^d$ (Degenne et al., 2020a),

$$2T_0(\mu)^{-1} = \max_{w \in \Delta_K} \min_{x \neq z^*(\mu)} \left(\frac{\langle \mu, z^*(\mu) - x \rangle}{\|z^*(\mu) - x\|_{V_w^\dagger}} \right)^2 \quad (7)$$

This formula has striking similarities with the one obtained for multiplicative ε -optimality. Defining $\mathcal{Z}_\varepsilon^z \stackrel{\text{def}}{=} \{z\} \cup \{(1 - \varepsilon)x\}_{x \neq z}$, we obtain that

$$T_\varepsilon^{\text{mul}}(\mu) = \min_{z \in \mathcal{Z}_\varepsilon^{\text{mul}}(\mu)} T_0(\mu, \mathcal{Z}_\varepsilon^z)$$

where $T_0(\mu, \mathcal{Z}_\varepsilon^z)$ is the characteristic time for a BAI problem where the set of answers is $\mathcal{Z}_\varepsilon^z$, as defined in (7). Therefore, the complexity of the multiplicative setting is equal to the complexity of a BAI instance in which \mathcal{K} is unchanged, but the set of answers is changed into $\mathcal{Z}_\varepsilon^z$.

C.3.2. HARD BAI INSTANCES

Lemma C.3 shows that there exists arbitrarily hard BAI instances that can be solved if seen as an ε -BAI problem.

Lemma C.3. *Let $\varepsilon > 0$, $\overline{\mathcal{M}} = \mathbb{R}^d$ and $T_\varepsilon(\mu, \mathcal{Z}) = T_\varepsilon(\mu)$. For all $\mu \in \mathcal{M}$, there exists a set of arms \mathcal{K} and a sequence of answers sets $(\mathcal{Z}_t)_{t \in \mathbb{N}}$, such that*

$$\lim_{t \rightarrow +\infty} T_0(\mu, \mathcal{Z}_t) = +\infty \quad \text{and} \quad \lim_{t \rightarrow +\infty} T_\varepsilon(\mu, \mathcal{Z}_t) < +\infty$$

The proof of Lemma C.3 is obtained directly by using Lemma C.4, which provides explicit hard instances.

Lemma C.4. *Let $\varepsilon > 0$, $\overline{\mathcal{M}} = \mathbb{R}^d$ and $\mu \in \mathcal{M}$. Let $\{z_i\}_{i \in [d-1]}$ such that $\{\frac{\mu}{\|\mu\|_2}\} \cup \{z_i\}_{i \in [d-1]}$ is an orthonormal basis of \mathbb{R}^d . Define $\mathcal{K} = \{\frac{\mu}{\|\mu\|_2}\} \cup \{z_i\}_{i \in [d-1]}$. Let $(\theta_t) \in [0, \frac{\pi}{2}]^{\mathbb{N}}$, a decreasing sequence such that $\theta_t \rightarrow_{+\infty} 0$. Define the sequence of answers sets $\mathcal{Z}_t = \{\frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}\} \cup \{z_i\}_{i \in [d-1]}$ for all $t \in \mathbb{N}$ and*

$$T_\varepsilon(\mu, \mathcal{Z})^{-1} \stackrel{\text{def}}{=} \max_{z \in \mathcal{Z}_\varepsilon(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2.$$

Then, we have $T_0(\mu, \mathcal{Z}_t) \rightarrow_{t \rightarrow +\infty} +\infty$ and

$$T_\varepsilon^{\text{add}}(\mu, \mathcal{Z}_t) \rightarrow_{t \rightarrow +\infty} \frac{1}{(1+\varepsilon)^2} \left(\frac{1}{w^*(d-1)} + \frac{d-1}{1-w^*(d-1)} \right) < +\infty$$

$$T_\varepsilon^{\text{mul}}(\mu, \mathcal{Z}_t) \rightarrow_{t \rightarrow +\infty} \frac{1}{w^*((1-\varepsilon)^2(d-1))} + (1-\varepsilon)^2 \frac{d-1}{1-w^*((1-\varepsilon)^2(d-1))} < +\infty$$

where

$$w^*(a) = \begin{cases} \frac{\sqrt{1+(1-a)^2}-1}{a-1} & \text{if } a > 1 \\ \frac{1}{2} & \text{if } a = 1 \\ \frac{\sqrt{1+(1-a)^2}+1}{1-a} & \text{if } a \in (0, 1) \end{cases}$$

Proof. Using the explicit formulas of Lemma C.2, we will perform the computations for both the multiplicative and the additive notions of ε -optimality. Define $\mathcal{K} = \left\{ \frac{\mu}{\|\mu\|_2} \right\} \cup \{z_i\}_{i \in [d-1]}$, where the d -th component is associated with $\frac{\mu}{\|\mu\|}$. Define $\mathcal{Z}_t = \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\} \cup \{z_i\}_{i \in [d-1]}$ for all $t \geq 1$, where $\left\{ \frac{\mu}{\|\mu\|} \right\} \cup \{z_i\}_{i \in [d-1]}$ is an orthonormal basis of \mathbb{R}^d . Note that $\|\cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}\|_2 = 1$ and $\langle \mu, z \rangle = 0$ if $z \in \{z_i\}_{i \in [d-1]}$, else its value is 1 and $\cos(\theta_t)$ for $\frac{\mu}{\|\mu\|_2}$ and $\cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}$ respectively. Therefore, we have directly that $z^*(\mu) = \frac{\mu}{\|\mu\|_2}$ and

$$\forall t \geq t_0, \quad \mathcal{Z}_\varepsilon(\mu) = \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}$$

where $t_0 = \inf\{t \in \mathbb{N}^* \mid \theta_t \leq \arccos(1-\varepsilon)\}$.

In the following, we consider $t \geq t_0$. Let $z \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}$ and $x \in \mathcal{Z}_t \setminus \{z\}$, we have

$$\langle \mu, z - x \rangle = \begin{cases} 1 - \cos(\theta_t) & \text{if } (z, x) = \left(\frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right) \\ \cos(\theta_t) - 1 & \text{if } (z, x) = \left(\cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}, \frac{\mu}{\|\mu\|_2} \right) \\ 1 & \text{if } x \in \{z_i\}_{i \in [d-1]} \end{cases}$$

$$\langle \mu, z - (1-\varepsilon)x \rangle = \begin{cases} 1 - (1-\varepsilon)\cos(\theta_t) & \text{if } (z, x) = \left(\frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right) \\ \cos(\theta_t) - (1-\varepsilon) & \text{if } (z, x) = \left(\cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}, \frac{\mu}{\|\mu\|_2} \right) \\ 1 & \text{if } x \in \{z_i\}_{i \in [d-1]} \end{cases}$$

Using that $\mathcal{K} = \left\{ \frac{\mu}{\|\mu\|_2} \right\} \cup \{z_i\}_{i \in [d-1]}$ is an orthogonal basis of \mathbb{R}^d . Let $w \in \Delta_K^\circ$

$$\|z - x\|_{V_w^\dagger}^2 = \begin{cases} (1 - \cos(\theta_t))^2 \frac{1}{w_d} + \sin(\theta_t)^2 \frac{1}{w_{d-1}} & \text{if } (z, x) \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}^2 \\ \frac{1}{w_d} + \frac{1}{w_i} & \text{if } z = \frac{\mu}{\|\mu\|_2} \text{ and } x \in \{z_i\}_{i \in [d-1]} \\ \cos(\theta_t)^2 \frac{1}{w_d} + \sin(\theta_t)^2 \frac{1}{w_{d-1}} + \frac{1}{w_i} & \text{if } z = \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \text{ and } x \in \{z_i\}_{i \in [d-2]} \\ \cos(\theta_t)^2 \frac{1}{w_d} + (1 - \sin(\theta_t))^2 \frac{1}{w_{d-1}} & \text{if } z = \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \text{ and } x = z_{d-1} \end{cases}$$

$$\|z - (1-\varepsilon)x\|_{V_w^\dagger}^2 = \begin{cases} (1 - (1-\varepsilon)\cos(\theta_t))^2 \frac{1}{w_d} + (1-\varepsilon)^2 \sin(\theta_t)^2 \frac{1}{w_{d-1}} & \text{if } (z, x) = \left(\frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right) \\ (\cos(\theta_t) - (1-\varepsilon))^2 \frac{1}{w_d} + \sin(\theta_t)^2 \frac{1}{w_{d-1}} & \text{if } (z, x) = \left(\cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}, \frac{\mu}{\|\mu\|_2} \right) \\ \frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i} & \text{if } z = \frac{\mu}{\|\mu\|_2} \text{ and } x \in \{z_i\}_{i \in [d-1]} \\ \cos(\theta_t)^2 \frac{1}{w_d} + \sin(\theta_t)^2 \frac{1}{w_{d-1}} + \frac{(1-\varepsilon)^2}{w_i} & \text{if } z = \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}, x \in \{z_i\}_{i \in [d-2]} \\ \cos(\theta_t)^2 \frac{1}{w_d} + (1-\varepsilon - \sin(\theta_t))^2 \frac{1}{w_{d-1}} & \text{if } z = \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1}, x = z_{d-1} \end{cases}$$

Since $\cos(\theta_t) \rightarrow_{t \rightarrow +\infty} 1$ and $\sin(\theta_t) \rightarrow_{t \rightarrow +\infty} 0$, we obtain

$$\begin{aligned} \frac{(\varepsilon + \langle \mu, z - x \rangle)^2}{\|z - x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \begin{cases} +\infty & \text{if } (z, x) \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}^2 \\ \frac{(1+\varepsilon)^2}{\frac{1}{w_d} + \frac{1}{w_i}} & \text{if } x \in \{z_i\}_{i \in [d-1]} \end{cases} \\ \frac{\langle \mu, z - (1-\varepsilon)x \rangle^2}{\|z - (1-\varepsilon)x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \begin{cases} w_d & \text{if } (z, x) \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}^2 \\ \frac{1}{\frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i}} & \text{if } x \in \{z_i\}_{i \in [d-1]} \end{cases} \end{aligned}$$

Since $|\mathcal{Z}_t| = d + 1$, the minimum over a d values $g_z(w, \mathcal{Z}_t) = \min_{x \in \mathcal{Z}_t} g_z(w, x)$ is a continuous function of (w, \mathcal{Z}_t) , where $g_z^{\text{add}}(w, x) = \frac{(\varepsilon + \langle \mu, z - x \rangle)^2}{\|z - x\|_{V_w^\dagger}^2}$ and $g_z^{\text{mul}}(w, x) = \frac{\langle \mu, z - (1-\varepsilon)x \rangle^2}{\|z - (1-\varepsilon)x\|_{V_w^\dagger}^2}$. Therefore, we have for all $z \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}$,

$$\begin{aligned} \min_{x \in \mathcal{Z}_t \setminus \{z\}} \frac{(\varepsilon + \langle \mu, z - x \rangle)^2}{\|z - x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \min_{i \in [d-1]} \frac{(1+\varepsilon)^2}{\frac{1}{w_d} + \frac{1}{w_i}} \\ \min_{x \in \mathcal{Z}_t \setminus \{z\}} \frac{\langle \mu, z - (1-\varepsilon)x \rangle^2}{\|z - (1-\varepsilon)x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \min_{i \in [d-1]} \frac{1}{\frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i}} \end{aligned}$$

where we used that $w_d > \frac{1}{\frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i}}$. Since $g_z(w, \mathcal{Z}_t)$ is a continuous function and Δ_K is compact, Berge's theorem yields that $g_z(\mathcal{Z}_t) = \max_{w \in \Delta_K} g_z(w, \mathcal{Z}_t)$ is continuous function of \mathcal{Z}_t . Therefore, we have for all $z \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}$,

$$\begin{aligned} \max_{w \in \Delta_K} \min_{x \in \mathcal{Z}_t \setminus \{z\}} \frac{(\varepsilon + \langle \mu, z - x \rangle)^2}{\|z - x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} (1+\varepsilon)^2 \max_{w \in \Delta_K} \min_{i \in [d-1]} \frac{1}{\frac{1}{w_d} + \frac{1}{w_i}} \\ \max_{w \in \Delta_K} \min_{x \in \mathcal{Z}_t \setminus \{z\}} \frac{\langle \mu, z - (1-\varepsilon)x \rangle^2}{\|z - (1-\varepsilon)x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \max_{w \in \Delta_K} \min_{i \in [d-1]} \frac{1}{\frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i}} \end{aligned}$$

Since $|\mathcal{Z}_\varepsilon(\mu)| = \left| \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\} \right| = 2$, the minimum over 2 values is continuous. Since the inverse is also a continuous function in \mathbb{R}^* and $\mathcal{Z}_\varepsilon(\mu) \rightarrow_{t \rightarrow +\infty} \left\{ \frac{\mu}{\|\mu\|_2} \right\}$, hence

$$\begin{aligned} T_\varepsilon^{\text{add}}(\mu, \mathcal{Z}_t) &\rightarrow_{t \rightarrow +\infty} \frac{1}{(1+\varepsilon)^2} \min_{w \in \Delta_K} \max_{i \in [d-1]} \frac{1}{w_d} + \frac{1}{w_i} \\ T_\varepsilon^{\text{mul}}(\mu, \mathcal{Z}_t) &\rightarrow_{t \rightarrow +\infty} \min_{w \in \Delta_K} \max_{i \in [d-1]} \frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i} \end{aligned}$$

Rewriting those optimization problems, we obtain

$$\begin{aligned} \min_{w \in \Delta_K} \max_{i \in [d-1]} \frac{1}{w_d} + \frac{1}{w_i} &= \min_{w \in (0,1)} \left(\frac{1}{w} + \frac{1}{1-w} \left(\max_{\tilde{w} \in \Delta_{d-1}} \min_{i \in [d-1]} \tilde{w}_i \right)^{-1} \right) = \min_{w \in (0,1)} \left(\frac{1}{w} + \frac{d-1}{1-w} \right) \\ \min_{w \in \Delta_K} \max_{i \in [d-1]} \frac{1}{w_d} + \frac{(1-\varepsilon)^2}{w_i} &= \min_{w \in (0,1)} \left(\frac{1}{w} + \frac{(1-\varepsilon)^2}{1-w} \left(\max_{\tilde{w} \in \Delta_{d-1}} \min_{i \in [d-1]} \tilde{w}_i \right)^{-1} \right) = \min_{w \in (0,1)} \left(\frac{1}{w} + \frac{(1-\varepsilon)^2(d-1)}{1-w} \right) \end{aligned}$$

where the first equality uses a rewriting of the simplex. The second inequalities are obtained since $\max_{\tilde{w} \in \Delta_{d-1}} \min_{i \in [d-1]} \tilde{w}_i = \frac{1}{d-1}$, achieved for $\tilde{w} = \frac{1}{d-1} \mathbf{1}_{d-1}$ and would lead to smaller values if there exists $i_0 \in [d-1]$ such that $\tilde{w}_{i_0} < \frac{1}{d-1}$, namely \tilde{w}_{i_0} .

Considering the function $f_a(w) = \frac{1}{w} + \frac{a}{1-w}$, we have $f'_a(w) = \frac{(a-1)w^2+2w-1}{w^2(1-w)^2}$. Let $w^*(a) = \arg \min_{w \in (0,1)} f_a(w)$. Therefore, by solving the second order equation with discriminant $\Delta = 4 + 4(1-a)^2$ and keeping the solution belonging to $(0, 1)$, we obtain

$$w^*(a) = \begin{cases} \frac{\sqrt{1+(1-a)^2}-1}{a-1} & \text{if } a > 1 \\ \frac{1}{2} & \text{if } a = 1 \\ \frac{\sqrt{1+(1-a)^2}+1}{1-a} & \text{if } a \in (0, 1) \end{cases}$$

Therefore, we have shown

$$\begin{aligned} T_\varepsilon^{\text{add}}(\mu, \mathcal{Z}_t) &\rightarrow_{t \rightarrow +\infty} \frac{1}{(1+\varepsilon)^2} \left(\frac{1}{w^*(d-1)} + \frac{d-1}{1-w^*(d-1)} \right) \\ T_\varepsilon^{\text{mul}}(\mu, \mathcal{Z}_t) &\rightarrow_{t \rightarrow +\infty} \frac{1}{w^*((1-\varepsilon)^2(d-1))} + (1-\varepsilon)^2 \frac{d-1}{1-w^*((1-\varepsilon)^2(d-1))} \end{aligned}$$

Now, let's show that the corresponding BAI problem diverges. Using the same arguments of continuity as before and taking the limit when $t \rightarrow +\infty$, we have

$$\begin{aligned} \frac{\langle \mu, z-x \rangle^2}{\|z-x\|_{V_w^\dagger}^2} &\rightarrow_{t \rightarrow +\infty} \frac{1}{\frac{1}{w_d} + \frac{1}{w_i}} && \text{if } x \in \{z_i\}_{i \in [d-1]} \\ &\sim_{+\infty} \frac{1}{\frac{1}{w_d} + \frac{\sin(\theta_t)^2}{(1-\cos(\theta_t))^2} \frac{1}{w_i}} \rightarrow_{t \rightarrow +\infty} 0 && \text{if } (z, x) \in \left\{ \frac{\mu}{\|\mu\|_2}, \cos(\theta_t) \frac{\mu}{\|\mu\|_2} + \sin(\theta_t) z_{d-1} \right\}^2 \end{aligned}$$

where the last part is obtained since $\frac{\sin(\theta_t)^2}{(1-\cos(\theta_t))^2} \rightarrow +\infty$. Therefore, using the same continuity arguments, we have

$$\min_{x \in \mathcal{Z}_t \setminus \left\{ \frac{\mu}{\|\mu\|} \right\}} \frac{\langle \mu, z-x \rangle^2}{\|z-x\|_{V_w^\dagger}^2} \rightarrow_{t \rightarrow +\infty} 0 \quad \text{and} \quad T_0(\mu, \mathcal{Z}_t)^{-1} \rightarrow_{t \rightarrow +\infty} 0.$$

This yields that $T_0(\mu, \mathcal{Z}_t) \rightarrow_{t \rightarrow +\infty} +\infty$. □

C.4. Greedy Characteristic Time

In Appendix C.4.1, we prove the lower bound for asymptotically greedy algorithm (Lemma 2.2). In Appendix C.4.2, we show that using $z_t \in z^*(\mu_t)$ yields an asymptotically greedy strategy (Lemma 2.3).

C.4.1. PROOF OF LEMMA 2.2

Lemma 3 in Garivier & Kaufmann (2021) is a useful change of measures result in the low-level form (involving probabilities and not expectation). This Lemma is key to derive a lower bound on the sample complexity of ε -BAI, as done in Garivier & Kaufmann (2021) and Degenne & Koolen (2019).

Lemma C.5 (Lemma 3 in Garivier & Kaufmann (2021)). *Consider two distributions \mathbb{P} and \mathbb{Q} . Let us denote the log-likelihood ratio after t rounds by $\mathcal{L}_t = \ln \frac{d\mathbb{P}}{d\mathbb{Q}}$. Then for any measurable event $A \in \mathcal{F}_t$ and threshold $\gamma \in \mathbb{R}$,*

$$\mathbb{P}[A] \leq e^\gamma \mathbb{Q}[A] + \mathbb{P}\{\mathcal{L}_t > \gamma\}$$

We rewrite Lemma 2 in Degenne & Koolen (2019) in the setting of (ε, δ) -PAC BAI for transductive linear bandits with Gaussian distribution.

Lemma C.6 (Lemma 2 in Degenne & Koolen (2019)). *For any answer $z \in \mathcal{Z}$, the divergence from μ to $\neg_\varepsilon z$ equals*

$$T_\varepsilon(\mu, z)^{-1} = \sup_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_w}^2 = \inf_{\mathbb{P}} \max_{a \in \mathcal{K}} \mathbb{E}_{\lambda \sim \mathbb{P}} [\|\mu - \lambda\|_{aa^\top}^2] \quad (8)$$

where the infimum ranges over probability distributions on $\neg_\varepsilon z$ supported on (at most) K points.

We rewrite Lemma 19 in [Degenne & Koolen \(2019\)](#) in the setting of (ε, δ) -PAC BAI for transductive linear bandits with Gaussian distribution. In words, if A is likely under μ , it must also be likely under at least one λ^k for sample sizes $t \ll T_\varepsilon(\mu, z)$. It was proven using Lemma C.5.

Lemma C.7 (Lemma 19 in [Degenne & Koolen \(2019\)](#)). *Let $z \in \mathcal{Z}$. Let w and $\lambda_1, \dots, \lambda_K$ be a minimax witness from Lemma C.6, and let us introduce the abbreviation $\alpha_a = \|\mu - \sum_{k \in [K]} w_k \lambda_k\|_{aa^\top}^2$ for all $a \in \mathcal{K}$. Fix a sample size t , and consider any event $A \in \mathcal{F}_t$. Then, for any $\beta > 0$*

$$\max_{k \in [K]} \mathbb{P}_{\lambda_k} \{A\} \geq e^{-t T_\varepsilon(\mu, z)^{-1} - \beta} \left(\mathbb{P}_\mu \{A\} - \exp\left(\frac{-\beta^2}{2t \max_{a \in \mathcal{K}} \alpha_a}\right) \right).$$

where $T_\varepsilon(\mu, z)^{-1} = \sup_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_w}^2$.

We are now ready to prove Lemma 2.2, which we recall below.

Lemma (Lemma 2.2). For all asymptotically greedy (ε, δ) -PAC strategy, for all $\mu \in \mathcal{M}$,

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_{g,\varepsilon}(\mu)$$

where the inverse of the greedy characteristic time is

$$T_{g,\varepsilon}(\mu)^{-1} \stackrel{\text{def}}{=} \max_{z \in z^*(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z} \frac{1}{2} \|\mu - \lambda\|_{V_w}^2,$$

and $T_{g,\varepsilon}(\mu) > T_\varepsilon(\mu)$ if and only if $z_F(\mu) \notin z^*(\mu)$.

Proof. We will bound the expectation of the stopping time τ_δ through Markov's inequality. For $T > 0$,

$$\mathbb{E}_\mu[\tau_\delta] \geq T(1 - \mathbb{P}_\mu[\tau_\delta \leq T]).$$

The event $\{\tau_\delta \leq T\}$ can be partitioned depending on the answer whether the answer is ε -optimal or not, and then whether it's $z^*(\mu)$ or not. By hypothesis, $\mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] \leq \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] \leq \delta$ and $0 \leq \lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] \leq \lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] = 0$. This yields

$$\mathbb{P}_\mu[\tau_\delta \leq T] = \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] + \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] + \sum_{z \in z^*(\mu)} \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} = z],$$

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta \leq T] \leq \sum_{z \in z^*(\mu)} \lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} = z]$$

Let $z \in z^*(\mu)$, w and $\lambda_1, \dots, \lambda_K$ be a minimax witness from Lemma C.6. Then by Lemma C.7, for any $\beta > 0$

$$\begin{aligned} \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} = z] &\leq \exp\left(\frac{T}{T_\varepsilon(\mu, z)} + \beta\right) \max_{k \in [K]} \mathbb{P}_{\lambda_k}[\tau_\delta \leq T, \hat{z} = z] + \exp\left(\frac{-\beta^2}{2T \max_{a \in \mathcal{K}} \alpha_a}\right) \\ &\leq \delta \exp\left(\frac{T}{T_\varepsilon(\mu, z)} + \beta\right) + \exp\left(\frac{-\beta^2}{2T \max_{a \in \mathcal{K}} \alpha_a}\right) \end{aligned}$$

where the second inequality uses that $\lambda_k \in \neg_\varepsilon z$ for all $k \in [K]$, hence $z \in z^*(\mu) \subseteq \mathcal{Z} \setminus \mathcal{Z}_\varepsilon(\lambda_k)$ and that the strategy satisfies $\mathbb{P}_\lambda[\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\lambda)] \leq \delta$ for all $\lambda \in \mathcal{M}$.

Let $\alpha = \max_{a \in \mathcal{K}} \alpha_a$. For $\eta \in (0, 1)$, $T = (1 - \eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \ln(1/\delta)$, $\beta = \frac{\eta}{2\sqrt{1-\eta}} \sqrt{\frac{T}{\min_{z \in z^*(\mu)} T_\varepsilon(\mu, z)} \log(1/\delta)}$, and all $z \in z^*(\mu)$,

$$\begin{aligned} \mathbb{P}_\mu[\tau_\delta \leq T, \hat{z} = z] &\leq \delta \exp\left(\frac{T}{T_\varepsilon(\mu, z)} + \frac{\eta}{2\sqrt{1-\eta}} \sqrt{\frac{T \log(1/\delta)}{\min_{z \in z^*(\mu)} T_\varepsilon(\mu, z)}}\right) + \exp\left(\frac{-\eta^2 \log(1/\delta)}{8(1-\eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \alpha}\right) \\ &\leq \delta \exp\left((1 - \eta/2) \log \frac{1}{\delta}\right) + \exp\left(\frac{-\eta^2 \log(1/\delta)}{8(1-\eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \alpha}\right) \\ &= \delta^{\eta/2} + \delta^{\eta^2/(8(1-\eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \alpha)} \rightarrow_{\delta \rightarrow 0} 0, \end{aligned}$$

where we used that $\min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \leq T_\varepsilon(\mu, z)$.

Since we have just shown $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta \leq T]$ for $T = (1 - \eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \ln(1/\delta)$, we obtain

$$\begin{aligned} \lim_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} &\geq \lim_{\delta \rightarrow 0} \frac{T}{\ln(1/\delta)} (1 - \mathbb{P}_\mu[\tau_\delta \leq T]) \geq (1 - \eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) \left(1 - \lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta \leq T]\right) \\ &= (1 - \eta) \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z). \end{aligned}$$

Letting η go to zero, we obtain that

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\log(1/\delta)} \geq \min_{z \in z^*(\mu)} T_\varepsilon(\mu, z) = T_{g,\varepsilon}(\mu).$$

The fact that $T_{g,\varepsilon}(\mu) > T_\varepsilon(\mu)$ if and only if $z_F(\mu) \notin z^*(\mu)$ is a direct consequence of the definition of $T_\varepsilon(\mu)$, $T_{g,\varepsilon}(\mu)$ and $z_F(\mu)$. \square

C.4.2. PROOF OF LEMMA 2.3

Lemma (Lemma 2.3). Any (ε, δ) -PAC strategy recommending $\hat{z} \in z^*(\mu_{\tau_\delta})$ is asymptotically greedy if the sampling rule ensures that $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, z^*(\mu_{\tau_\delta}) = z^*(\mu)] = 1$.

Proof. By definition of (ε, δ) -PAC, we have $\mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] \leq \delta$. Since $\hat{z} \in z^*(\mu_{\tau_\delta})$, we have $\{\tau_\delta < +\infty, z^*(\mu_{\tau_\delta}) = z^*(\mu)\} \subseteq \{\tau_\delta < +\infty, \hat{z} \in z^*(\mu)\}$. Therefore, the assumption yields $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in z^*(\mu)] = 1$. Partitioning the event $\{\tau_\delta < +\infty, \hat{z} \in \mathcal{Z}\}$ (which obviously holds)

$$\delta \geq \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] = 1 - \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in z^*(\mu)] - \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] \geq 0.$$

Taking the limit $\delta \rightarrow 0$ yields

$$\lim_{\delta \rightarrow 0} \mathbb{P}_\mu[\tau_\delta < +\infty, \hat{z} \in \mathcal{Z}_\varepsilon(\mu) \setminus z^*(\mu)] = 0,$$

i.e. the strategy is asymptotically greedy. \square

D. Proof of Lemma 3.1

The proof leverages the concentration inequalities in the Corollary 10 of [Kaufmann & Koolen \(2021\)](#), which we restate below.

Lemma D.1. Let ν a Gaussian bandit with mean μ . Let $S \subseteq \mathcal{K}$ and $x > 0$.

$$\mathbb{P}_\nu \left[\exists t \in \mathbb{N} : \sum_{a \in S} N_t^a d_{KL}(\mu_t^a, \mu^a) > \sum_{a \in S} 2 \ln(4 + \ln(N_t^a)) + |S| \mathcal{C}_G \left(\frac{x}{|S|} \right) \right] \leq e^{-x}$$

where \mathcal{C}_G is defined in [Kaufmann & Koolen \(2021\)](#) by $\mathcal{C}_G(x) = \max_{\lambda \in [1/2, 1]} \frac{g_G(\lambda) + x}{\lambda}$ and

$$g_G(\lambda) = 2\lambda - 2\lambda \ln(4\lambda) + \ln \zeta(2\lambda) - \frac{1}{2} \ln(1 - \lambda), \quad (9)$$

where ζ is the Riemann ζ function and $\mathcal{C}_G(x) \approx x + \ln(x)$.

Let $\mathcal{T} \subseteq \mathbb{N}$ be a set of times with $|\mathcal{T}| = \infty$. For both notions of ε -optimality (additive or multiplicative), the crucial element is the stopping criterion, which performs a GLRT at each time $t \in \mathcal{T}$ for an arbitrary candidate ε -optimal answer $z_t \in \mathcal{Z}_\varepsilon(\mu_t)$,

$$\inf_{\lambda \in \neg_\varepsilon z_t} \|\mu_{t-1} - \lambda\|_{V_{N_{t-1}}}^2 > 2\beta(t-1, \delta)$$

Lemma 3.1 holds for both notions of ε -optimality (additive or multiplicative) and is agnostic to the sampling rule. Moreover, it holds when the stopping criterion is evaluated only on some predefined indices of time $\mathcal{T} \subseteq \mathbb{N}$ with $|\mathcal{T}| = \infty$, for example on a geometric grid, and for any recommendation rule verifying $z_t \in \mathcal{Z}_\varepsilon(\mu_t)$ for all $t \in \mathcal{T}$.

Lemma (Lemma 3.1). Let $\mathcal{T} \subseteq \mathbb{N}$ be a set of times with $|\mathcal{T}| = \infty$. Regardless of the sampling rule and for any recommendation rule such that $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$ for all $t \in \mathcal{T}$, then evaluating the stopping criterion (5) at each time $t \in \mathcal{T}$ with the threshold

$$\beta(t, \delta) = 2K \ln \left(4 + \ln \left(\frac{t}{K} \right) \right) + K \mathcal{C}_G \left(\frac{\ln \left(\frac{1}{\delta} \right)}{K} \right)$$

yields an (ε, δ) -PAC strategy. \mathcal{C}_G is defined in (9) and $\mathcal{C}_G(x) \approx x + \ln(x)$.

Proof. For simplicity, let's first prove the result when $\mathcal{T} = \mathbb{N}$. Considering a notion of ε -optimality (additive or multiplicative), an arbitrary sampling rule and a recommendation rule such that for all t , $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$. By evaluating the stopping criterion (5) at each time $t \in \mathbb{N}$, we obtain the following inequalities:

$$\begin{aligned} \mathbb{P}_\mu [\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] &\leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, z_t \notin \mathcal{Z}_\varepsilon(\mu), \inf_{\lambda \in \neg_\varepsilon z_t} \|\mu_{t-1} - \lambda\|_{V_{N_{t-1}}}^2 > 2\beta(t-1, \delta) \right] \\ &\leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, z_t \notin \mathcal{Z}_\varepsilon(\mu), \|\mu_{t-1} - \mu\|_{V_{N_{t-1}}}^2 > 2\beta(t-1, \delta) \right] \\ &\leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, \|\mu_{t-1} - \mu\|_{V_{N_{t-1}}}^2 > 2\beta(t-1, \delta) \right] \\ &= \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, \sum_{a \in \mathcal{K}} N_{t-1}^a d_{\text{KL}}(\mu_{t-1}^a, \mu^a) > \beta(t-1, \delta) \right] \end{aligned}$$

where the first inequality is obtained by definition of the stopping criterion and using that $\mathbb{P}[\exists t > n_0, A_t] \leq \mathbb{P}[\exists t \in \mathbb{N}, A_t]$. Crucially, since $\beta(t-1, \delta) > 0$, it implicitly requires that $z_t \in \mathcal{Z}_\varepsilon(\mu_t)$, otherwise $\inf_{\lambda \in \neg_\varepsilon z_t} \|\mu_{t-1} - \lambda\|_{V_{N_{t-1}}}^2 = 0$. The second inequality is obtained since $\mu \in \neg_\varepsilon z_t$ (as $z_t \notin \mathcal{Z}_\varepsilon(\mu)$) and the third inequality uses $\mathbb{P}[A, B] \leq \mathbb{P}[A]$. The equality is a direct consequence of our notations and the Gaussian assumption with $\sigma^2 = 1$, i.e. $V_N = \sum_{a \in \mathcal{K}} N^a a a^\top$ and $d_{\text{KL}}(\mu^a, \lambda^a) = \frac{1}{2} \|\mu - \lambda\|_{aa^\top}^2$.

Using the concavity of $x \mapsto 2 \ln(4 + \ln(x))$ and the fact that $\sum_{a \in \mathcal{K}} N_t^a = t$, we obtain $\sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_t^a)) \leq 2K \ln(4 + \ln(\frac{t}{K}))$. Since for $u \leq v$, $\mathbb{P}[X > v] \leq \mathbb{P}[X > u]$, by defining the stopping threshold as in (6) and using Lemma D.1 for $S = \mathcal{K}$, we obtain:

$$\begin{aligned} \mathbb{P}_\mu [\tau_\delta < +\infty, \hat{z} \notin \mathcal{Z}_\varepsilon(\mu)] &\leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, \sum_{a \in \mathcal{K}} N_{t-1}^a d_{\text{KL}}(\mu_{t-1}^a, \mu^a) > \beta(t-1, \delta) \right] \\ &\leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N} : \sum_{a \in \mathcal{K}} N_{t-1}^a d_{\text{KL}}(\mu_{t-1}^a, \mu^a) > \sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_{t-1}^a)) + K \mathcal{C}_G \left(\frac{\ln \left(\frac{1}{\delta} \right)}{K} \right) \right] \\ &\leq \delta \end{aligned}$$

This concludes the proof for stopping-recommendation pair evaluating the stopping criterion at each time t . When the stopping criterion is evaluated only on some predefined indices of time $\mathcal{T} \subseteq \mathbb{N}$ with $|\mathcal{T}| = \infty$, for example on a geometric grid, the proof is identical. Since $\mathbb{P}[\exists t \in \mathcal{T}, A_t] \leq \mathbb{P}[\exists t \in \mathbb{N}, A_t]$, we perform the same manipulations with \mathcal{T} instead of \mathbb{N} and conclude

$$\mathbb{P}_\mu \left[\exists t \in \mathcal{T}, \sum_{a \in \mathcal{K}} N_{t-1}^a d_{\text{KL}}(\mu_{t-1}^a, \mu^a) > \beta(t-1, \delta) \right] \leq \mathbb{P}_\mu \left[\exists t \in \mathbb{N}, \sum_{a \in \mathcal{K}} N_{t-1}^a d_{\text{KL}}(\mu_{t-1}^a, \mu^a) > \beta(t-1, \delta) \right]$$

□

Note that Lemma 3.1 covers all the stopping-recommendation pairs considered in Section 3.1: the different candidate answers (greedy, instantaneous furthest and furthest answers) and the computational relaxations (each time t , sticky and lazy updates).

E. Proof of Theorem 4.1

In this appendix, we provide the full proof of Theorem 4.1, which is recalled below. It is organized as follows:

- In Appendix E.1, the proof scheme is first sketched, then detailed and proved.
- In Appendix E.2, we show that combining an instantaneous furthest answer $z_t \in z_F(\mu_{t-1}, N_{t-1})$ as recommendation rule and the stopping criterion (5) yield the desired property (17) to obtain asymptotic optimality. Moreover, we introduce the *lazy* and *sticky* computational relaxations for the stopping-recommendation pair, which both satisfy the property (17).
- In Appendix E.3, the sampling rule used in $L\varepsilon$ BAI is shown to satisfy the desired property (18) on the sampling rule.
- We assemble the different blocks of the proof in Appendix E.4, yielding the Theorem 4.1.
- Technical results are stated and proved in Appendix E.5.

We shall assume that $\|\mu_t\|_2 \leq M$ holds for all t . This can be guaranteed by projecting the estimates onto the set of realizable models. We will abuse the Landau's notations o , \mathcal{O} and Θ to cope for t and $\ln\left(\frac{1}{\delta}\right)$ simultaneously, while discarding crossed terms. The notations \tilde{o} , $\tilde{\mathcal{O}}$ and $\tilde{\Theta}$ are used similarly with the simplification of dropping the poly-logarithmic multiplicative factors.

For the sake of generality, the proofs will be conducted with parameters $b > 0$ and $\alpha > 1$ in the definitions of the exploration bonus $f(t) \stackrel{\text{def}}{=} 2\beta(t, t^{1/\alpha})$ and the slacks $(c_s^a)_{(s,a) \in \llbracket n_0, t-1 \rrbracket \times \mathcal{K}}$ defined below in (20). In the main content, we directly used the values chosen thanks to the analysis $(b, \alpha) = (1, 3)$ as discussed in Appendix E.4.

E.1. Proof Scheme

The proof scheme sketched below is inspired by recent works using a game approach (Degenne et al., 2020a). We present below a sketch of the proof, see Appendix E.1.1 for a more detail proof scheme.

In order to obtain an asymptotic upper bound on $\mathbb{E}_\mu[\tau_\delta]$, we derive a non-asymptotic one and take the limit $\delta \rightarrow 0$. Having multiple ε -optimal answers is a key difficulty in several arguments. Our main contribution with respect to this proof lies in overcoming this hurdle.

Using Lemma E.1 (Lemma 13 in Degenne et al. (2020a)), the proof boils down to proving the existence and an upper bound on $T_1(\delta) \in \mathbb{N}$, such that for all $t \geq T_1(\delta)$ if a concentration event \mathcal{E}_t holds then the algorithm has already stopped, $\mathcal{E}_t \subseteq \{\tau_\delta \leq t\}$. To obtain the asymptotic optimality of the identification strategy, the upper bound should satisfy $T_1(\delta) \leq T_\varepsilon(\mu) \ln\left(\frac{1}{\delta}\right) + o\left(\ln\left(\frac{1}{\delta}\right)\right)$, where the dependency in t vanishes when $\delta \rightarrow 0$.

Lemma E.1 (Lemma 13 in Degenne et al. (2020a)). *Let $(\mathcal{E}_t)_{t \geq 1}$ be a sequence of concentration events, such that for all $t \geq 1$, $\mathbb{P}_\mu[\mathcal{E}_t^c] \leq \frac{1}{t^\alpha}$ and an identification strategy such that for all $\delta \in (0, 1)$, there exists $T_1(\delta) \in \mathbb{N}$ such that for $t \geq T_1(\delta)$, $\mathcal{E}_t \subseteq \{\tau_\delta \leq t\}$. Then, $\mathbb{E}_\mu[\tau_\delta] \leq T_1(\delta) + \frac{1}{\alpha-1}$.*

Proof. Since $\alpha > 1$, using an integral-sum comparison we obtain that: $\sum_{t=1}^{+\infty} \frac{1}{t^\alpha} \leq \int_{x=1}^{+\infty} \frac{1}{x^\alpha} dx = \frac{1}{\alpha-1}$. Since τ_δ is a positive random variable, we have:

$$\mathbb{E}_\mu[\tau_\delta] = \sum_{t \in \mathbb{N}} \mathbb{P}_\mu[\tau_\delta > t] \leq T_1(\delta) + \sum_{t \geq T_1(\delta)} \mathbb{P}_\mu[\mathcal{E}_t^c] \leq T_1(\delta) + \sum_{t \geq T_1(\delta)} \frac{1}{t^\alpha} \leq T_1(\delta) + \frac{1}{\alpha-1}$$

where we split the sum in two terms and used that for $t \geq T_1(\delta)$, $\{\tau_\delta > t\} \subseteq \mathcal{E}_t^c$. □

Considering the sequence of concentration events $(\mathcal{E}_t)_{t \geq 1}$, where for all $t \geq 1$

$$\mathcal{E}_t \stackrel{\text{def}}{=} \left\{ \forall s \leq t : \|\mu_s - \mu\|_{V_{N_s}}^2 \leq f(t) \right\} \quad (10)$$

ensures that $\mathbb{P}_\mu \left[\mathcal{E}_t^{\mathcal{G}} \right] \leq \frac{1}{t^\alpha}$ (Lemma E.2), where $f(t) \stackrel{\text{def}}{=} 2\beta(t, t^{1/\alpha})$. To show the existence of $T_1(\delta)$ and an upper bound leading to asymptotic optimality, it is sufficient to show that under \mathcal{E}_t , if the algorithm does not stop at time $t + 1$, then

$$tT_\varepsilon(\mu)^{-1} \leq \ln\left(\frac{1}{\delta}\right) + o\left(t + \ln\left(\frac{1}{\delta}\right)\right)$$

To derive the above inequality, the analysis distinguishes between two independent components. Under \mathcal{E}_t , if the algorithm does not stop at time $t + 1$, the stopping-recommendation pair should satisfy

$$2\beta(t, \delta) \geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 - o\left(t + \ln\left(\frac{1}{\delta}\right)\right)$$

while the sampling rule has to verify

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq 2tT_\varepsilon(\mu)^{-1} - o\left(t + \ln\left(\frac{1}{\delta}\right)\right)$$

The expression $\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2$ only feature the empirical counts. Therefore, our proof scheme allows to combine any stopping-recommendation pair and any sampling rule, provided they satisfy the corresponding inequality. The fact that our algorithms are (ε, δ) -PAC is a direct consequence of Lemma 3.1.

E.1.1. DETAILS AND PROOFS

As stated in Subsection 4.1, the proof scheme is inspired by recent works using a game approach (Degenne et al., 2019; 2020a). To obtain an asymptotic upper bound on the expected sample complexity $\mathbb{E}_\mu[\tau_\delta]$, we first derive a non-asymptotic upper bound, which might be loose as a function of t . Then, we consider the limit $\delta \rightarrow 0$.

The first step towards the non-asymptotic upper-bound on $\mathbb{E}_\mu[\tau_\delta]$ is Lemma E.1, whose proof is inspired by Lemma 13 in Degenne et al. (2020a). This approach is due to Garivier & Kaufmann (2016).

In Lemma E.2, it is shown that the sequence of concentration events $(\mathcal{E}_t)_{t \geq 1}$ defined in (10), for all $t \geq 1$

$$\mathcal{E}_t \stackrel{\text{def}}{=} \left\{ \forall s \leq t : \|\mu_s - \mu\|_{V_{N_s}}^2 \leq f(t) \right\}$$

satisfies the first condition of Lemma E.1. Recall that $\beta(t, \delta) = \Theta\left(\ln\left(\frac{1}{\delta}\right) + \ln(\ln(t))\right)$. Since $f(t) \stackrel{\text{def}}{=} 2\beta(t, t^{1/\alpha})$, f is a logarithmic function of t , i.e. $f(t) = \Theta(\ln(t))$.

Lemma E.2. *Considering $(\mathcal{E}_t)_{t \geq 1}$ in (10), we have: for all $t \geq 1$, $\mathbb{P}_\mu \left[\mathcal{E}_t^{\mathcal{G}} \right] \leq \frac{1}{t^\alpha}$*

Proof. Similarly to the proof of Lemma 3.1, we will apply Corollary 10 in Kaufmann & Koolen (2021):

$$\mathbb{P}_\mu \left[\exists s \in \mathbb{N} : \sum_{a \in \mathcal{K}} N_s^a d_{\text{KL}}(\mu_s^a, \mu^a) > \sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_s^a)) + KC_G \left(\frac{\ln\left(\frac{1}{\delta}\right)}{K} \right) \right] \leq \delta$$

By concavity of $x \mapsto \ln(4 + \ln(x))$ (which is also an increasing function) and $\sum_{a \in \mathcal{K}} N_s^a = s$, we have: for all $s \in [t]$,

$$\sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_s^a)) \leq 2K \ln\left(4 + \ln\left(\frac{s}{K}\right)\right) \leq 2K \ln\left(4 + \ln\left(\frac{t}{K}\right)\right)$$

Recall that for $u \leq v$, $\mathbb{P}[X > v] \leq \mathbb{P}[X > u]$ and $\mathbb{P}_\mu \left[\bigcup_{s \in [t]} A_s \right] \leq \mathbb{P}_\mu \left[\bigcup_{s \in \mathbb{N}} A_s \right]$.

$$\begin{aligned} \mathbb{P}_\mu \left[\mathcal{E}_t^{\mathcal{G}} \right] &= \mathbb{P}_\mu \left[\exists s \in [t], \|\mu_s - \mu\|_{V_{N_s}}^2 > 2\beta(t, t^{-\alpha}) \right] \\ &\leq \mathbb{P}_\mu \left[\exists s \in [t], \sum_{a \in \mathcal{K}} N_s^a d_{\text{KL}}(\mu_s^a, \mu^a) > \sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_s^a)) + KC_G \left(\frac{\ln(t^\alpha)}{K} \right) \right] \\ &\leq \mathbb{P}_\mu \left[\exists s \in \mathbb{N} : \sum_{a \in \mathcal{K}} N_s^a d_{\text{KL}}(\mu_s^a, \mu^a) > \sum_{a \in \mathcal{K}} 2 \ln(4 + \ln(N_s^a)) + KC_G \left(\frac{\ln(t^\alpha)}{K} \right) \right] \leq \frac{1}{t^\alpha} \end{aligned}$$

□

Lemma E.3 gives a sufficient condition on the identification strategy to ensure there exists $T_1(\delta) \in \mathbb{N}$ satisfying the second condition of Lemma E.1.

Lemma E.3. *Let $(\mathcal{E}_t)_{t \geq 1}$ be a sequence of concentration events. Assume that the identification strategy verifies the following property: there exists $\beta_1, \beta_2 \in (0, 1]$ such that if the algorithm does not stop at time $t + 1$ and \mathcal{E}_t holds then*

$$H_\mu(t, \delta) \geq tT_\varepsilon(\mu)^{-1} \quad \text{where } H_\mu(t, \delta) = \ln\left(\frac{1}{\delta}\right) + \tilde{\mathcal{O}}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right) \quad (11)$$

Then, for all $\delta \in (0, 1)$, there exists $T_1(\delta) \in \mathbb{N}$ such that for $t \geq T_1(\delta)$, $\mathcal{E}_t \subseteq \{\tau_\delta \leq t\}$.

Proof. Assume there exists an identification strategy satisfying the above condition. Let $T(\delta)$ be the maximum of the $t \in \mathbb{N}$ such that

$$H_\mu(t, \delta) \geq tT_\varepsilon(\mu)^{-1} \quad (12)$$

where $T(\delta)$ always exists since $H_\mu(t, \delta) = \mathcal{O}(t^{1-\beta_1}) = o(t)$. Let δ_{\min} be the largest $\delta \in (0, 1)$ such that

$$H_\mu\left(\ln(1/\delta)^{(1-\beta_1)^{-1}}, \delta\right) < \ln(1/\delta)^{(1-\beta_1)^{-1}} T_\varepsilon(\mu)^{-1} \quad (13)$$

where δ_{\min} always exists since $H_\mu\left(\ln(1/\delta)^{(1-\beta_1)^{-1}}, \delta\right) = \mathcal{O}(\ln(1/\delta)) = o\left(\ln(1/\delta)^{(1-\beta_1)^{-1}}\right)$. It depends only on the parameters of the problem: $T_\varepsilon(\mu)$, L , M , d and α .

For $\delta \leq \delta_{\min}$, we have $H_\mu\left(\ln(1/\delta)^{(1-\beta_1)^{-1}}, \delta\right) < \ln(1/\delta)^{(1-\beta_1)^{-1}} T_\varepsilon(\mu)^{-1}$, hence $T(\delta) < \ln(1/\delta)^{(1-\beta_1)^{-1}}$. Since $t \mapsto H_\mu(t, \delta)$ is increasing, plugging $\ln(1/\delta)^{(1-\beta_1)^{-1}}$ in the left-hand side of (12), we obtain $T(\delta) < T_0(\delta)$ where

$$T_0(\delta) \stackrel{\text{def}}{=} T_\varepsilon(\mu) H_\mu\left(\ln(1/\delta)^{(1-\beta_1)^{-1}}, \delta\right) \quad (14)$$

Moreover, under \mathcal{E}_t and if we do not stop, it implies that $t \leq T(\delta)$. To sum up, we have shown: (1) for $\delta \leq \delta_{\min}$, for $t \geq T_0(\delta)$, we know that $\tau_\delta \leq t$ and (2) for $\delta > \delta_{\min}$, for $t \geq T(\delta) + 1$, we know that $\tau_\delta \leq t$. In both cases, this means that $\mathcal{E}_t \subseteq \{\tau_\delta \leq t\}$. Therefore, there exists $T_1(\delta)$ defined as

$$T_1(\delta) \stackrel{\text{def}}{=} \begin{cases} T_0(\delta) & \text{if } \delta \leq \delta_{\min} \\ T(\delta) + 1 & \text{else.} \end{cases} \quad (15)$$

such that for $t \geq T_1(\delta)$, $\mathcal{E}_t \subseteq \{\tau_\delta \leq t\}$. □

By definition of $T_1(\delta)$ in (15), we have $T_1(\delta) = T_\varepsilon(\mu) \ln\left(\frac{1}{\delta}\right) + o(\ln(1/\delta))$. This fact is crucial to obtain the upper bound with the right constant when considering $\delta \rightarrow 0$, hence it is necessary to prove asymptotic optimality.

The last and hardest component of the proof is to show that: under $(\mathcal{E}_t)_{t \geq 1}$ (10) and if the algorithm does not stop at time $t + 1$, then the identification strategy satisfies the condition (11) of Lemma E.3. When this holds, we say that the strategy is *sample-efficient*.

Plugging together all the previous results, we obtain directly Lemma E.4.

Lemma E.4. *Let $(\mathcal{E}_t)_{t \geq 1}$ as in (10). Using a sample-efficient (ε, δ) -PAC strategy, i.e. (11) holds true, yields an (ε, δ) -PAC strategy and, for all $\mu \in \mathcal{M}$ such that $|z_F(\mu)| = 1$,*

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln\left(\frac{1}{\delta}\right)} \leq T_\varepsilon(\mu)$$

Proof. Under these hypotheses, the two conditions of Lemma E.1 are fulfilled. Using the definition of $(\mathcal{E}_t)_{t \geq 1}$ in (10), the first condition holds thanks to Lemma E.2. By definition of a sample-efficient identification strategy we can apply Lemma E.3, hence the second condition is also true. Therefore, applying Lemma E.1, we have: for all $\delta \in (0, 1)$

$$\mathbb{E}_\mu[\tau_\delta] \leq T_1(\delta) + \frac{1}{\alpha - 1} \quad (16)$$

with $T_1(\delta) = T_\varepsilon(\mu) \ln\left(\frac{1}{\delta}\right) + o(\ln(1/\delta))$. Dividing by $\ln\left(\frac{1}{\delta}\right)$ on both side and taking $\limsup_{\delta \rightarrow 0}$, we obtain the asymptotic optimality of this algorithm.

By assumption, the algorithm was (ε, δ) -PAC. By the inequality (16), we have that τ_δ is finite almost surely, i.e. $\mathbb{P}_\mu[\tau_\delta < +\infty] = 1$. \square

Sample-Efficient Identification Strategy When proving that a strategy is sample-efficient, we distinguish two independent conditions. Under \mathcal{E}_t if the algorithms does not stop at time $t + 1$, the stopping-recommendation pair should satisfy

$$2\beta(t, \delta) \geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu - \lambda\|_{V_{N_t}}^2 - \tilde{\mathcal{O}}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right) \quad (17)$$

while the sampling rule has to verify

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq 2tT_\varepsilon(\mu)^{-1} - \tilde{\mathcal{O}}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right) \quad (18)$$

with $(\beta_1, \beta_2) \in (0, 1]^2$. By grouping together both inequalities (17-18), we obtain directly the condition of Lemma E.3. The above lemmas are assembled in a global proof in Appendix E.4.

E.2. Stopping-Recommendation Pairs

Lemma E.5 shows that the desired condition on the stopping-recommendation pair (17) holds when considering an instantaneous furthest answer as recommendation rule and evaluating the stopping rule at each time t . In particular, it holds for $(\beta_1, \beta_2) = (1, \frac{1}{2})$ (the poly-logarithmic dependence in t is hidden in the notation $\tilde{\mathcal{O}}$). The extension of the proof for other update schemes is detailed in Appendix E.2.1.

Note that while the condition on the recommendation rule to have a (ε, δ) -PAC algorithm is very mild, i.e. $z_t \in \mathcal{Z}_\varepsilon(\mu_{t-1})$, the choice of z_t is crucial to obtain an asymptotically optimal algorithm. Intuitively it should converge asymptotically towards $z_F(\mu)$. For this reason, using a greedy answer $z_t \in z^*(\mu_t)$ as recommendation rule is doomed to fail whenever $z_F(\mu) \notin z^*(\mu)$ (which is often the case). It remains unclear whether using a furthest answer for the current estimator $z_t \in z_F(\mu_t)$ as recommendation rule would result in an asymptotically optimal algorithm. While converging to $z_F(\mu)$, it might be inefficient when associated with the stopping criterion (5). This is an interesting open question to investigate in future work.

Lemma E.5. *Regardless of the sampling rule, an identification strategy evaluating at each time t the stopping rule (5) with stopping threshold $\beta(t, \delta)$ (6) for an instantaneous furthest answer $z_t \in z_F(\mu_{t-1}, N_{t-1})$ satisfies that, under \mathcal{E}_t as in (10), if the algorithm does not stop at time $t + 1$, then*

$$2\beta(t, \delta) \geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu - \lambda\|_{V_{N_t}}^2 - h_\delta(t)$$

where $h_\delta(t) = \sqrt{8f(t)\beta(t, \delta)} + 4f(t) = \Theta\left(\ln(t) + \sqrt{\ln\left(\frac{1}{\delta}\right)}\right)$.

Proof. Fix any time $t \geq 1$. Suppose that \mathcal{E}_t holds and the algorithm does not stop at time $t + 1$. From the stopping rule and the definition of $z_{t+1} \in z_F(\mu_t, N_t)$, we obtain

$$2\beta(t, \delta) \geq \inf_{\lambda \in \neg_{\varepsilon} z_{t+1}} \|\mu_t - \lambda\|_{V_{N_t}}^2 = \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu_t - \lambda\|_{V_{N_t}}^2 \geq \max_{z \in z_F(\mu, N_t)} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu_t - \lambda\|_{V_{N_t}}^2 \quad (19)$$

where the last inequality uses Lemma E.15. Combining the triangular inequality and concentration event \mathcal{E}_t ,

$$\begin{aligned} \|\mu_t - \lambda\|_{V_{N_t}}^2 &\geq (\|\mu - \lambda\|_{V_{N_t}} - \|\mu_t - \mu\|_{V_{N_t}})^2 \\ &\geq \|\mu - \lambda\|_{V_{N_t}}^2 - 2\|\mu - \lambda\|_{V_{N_t}} \|\mu_t - \mu\|_{V_{N_t}} \\ &\geq \|\mu - \lambda\|_{V_{N_t}}^2 - 2\|\mu - \lambda\|_{V_{N_t}} \sqrt{f(t)}. \end{aligned}$$

Let $\tilde{z} \in \arg \max_{z \in z_F(\mu, N_t)} \inf_{\lambda \in \neg_{\varepsilon} z} \|\mu_t - \lambda\|_{V_{N_t}}^2$ and $\tilde{\lambda} \in \arg \min_{\lambda \in \neg_{\varepsilon} \tilde{z}} \|\mu_t - \lambda\|_{V_{N_t}}^2$. Then,

$$\|\mu_t - \tilde{\lambda}\|_{V_{N_t}}^2 \geq \|\mu - \tilde{\lambda}\|_{V_{N_t}}^2 - 2\sqrt{f(t)} \sqrt{\|\mu - \tilde{\lambda}\|_{V_{N_t}}^2}$$

Using Lemma E.10 for $y = \|\mu_t - \tilde{\lambda}\|_{V_{N_t}}^2$, $\alpha = 2\sqrt{f(t)}$ and $x = \|\mu - \tilde{\lambda}\|_{V_{N_t}}^2$, we obtain

$$\begin{aligned} \|\mu_t - \tilde{\lambda}\|_{V_{N_t}}^2 &\geq \|\mu - \tilde{\lambda}\|_{V_{N_t}}^2 - 2\sqrt{f(t)}\sqrt{\|\mu_t - \tilde{\lambda}\|_{V_{N_t}}^2} - 4f(t) \\ &\geq \|\mu - \tilde{\lambda}\|_{V_{N_t}}^2 - 2\sqrt{f(t)}\sqrt{2\beta(t, \delta)} - 4f(t) \\ &\geq \inf_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_{N_t}}^2 - \sqrt{8f(t)\beta(t, \delta)} - 4f(t) \\ &= \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 - \sqrt{8f(t)\beta(t, \delta)} - 4f(t). \end{aligned}$$

The second inequality is obtained by using (19) and the definition of $\tilde{\lambda}$. The third is obtained by taking the $\inf_{\lambda \in \neg_\varepsilon \tilde{z}}$, which is possible since $\tilde{\lambda} \in \neg_\varepsilon \tilde{z}$. The equality uses that $\tilde{z} \in z_F(\mu, N_t)$. Therefore, we have shown that

$$2\beta(t, \delta) \geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 - h_\delta(t)$$

where $h_\delta(t) = \sqrt{8f(t)\beta(t, \delta)} + 4f(t) = \Theta\left(\ln(t) + \sqrt{\ln\left(\frac{1}{\delta}\right)}\right)$, hence sub-linear function of both t and $\ln\left(\frac{1}{\delta}\right)$. \square

E.2.1. COMPUTATIONAL RELAXATIONS

Depending on \mathcal{Z} , computing an instantaneous furthest answer at each time t might be too costly. To reduce the computational cost, we propose two different relaxations for any stopping-recommendation pair: the *lazy* version and the *sticky* one. The main idea behind both of them is to avoid the computation of a new recommendation rule at each time. Defining a grid of time $\mathcal{T} \subseteq \mathbb{N}$, we stick to the previous candidate answer when $t \notin \mathcal{T}$, i.e. $z_t = z_{t-1}$, else we compute a new one. To satisfy the sole requirement on the recommendation rule (Lemma 3.1), a new value has to be computed when $t \notin \mathcal{T}$ and $z_{t-1} \notin \mathcal{Z}_\varepsilon(\mu_{t-1})$.

Both computational relaxations use the stopping criterion defined in (5). While the sticky version evaluates it at each time $t > n_0$, the lazy version only does it when the recommendation rule is updated ($t \in \mathcal{T}$). It is important to notice that thanks to Lemma 3.1, the lazy and sticky computational relaxations yield (ε, δ) -PAC strategy regardless of the sampling rule.

In the following, we consider an instantaneous furthest answer as recommendation rule and a geometric-like grid of time \mathcal{T} . To ensure asymptotic optimality of the algorithm, we need to have a strictly positive and decreasing expansion parameter $(\gamma_i)_{i \in \mathbb{N}^*}$ such that $\gamma_i \rightarrow 0$. More formally, we use $\mathcal{T} \stackrel{\text{def}}{=} \{n_0 + 1\} \cup \{t_i\}_{i \in \mathbb{N}}$ where $t_0 > n_0$ denotes the end of the first phase and for all $i \in \mathbb{N}^*$, $t_i = \lceil (1 + \gamma_i)t_{i-1} \rceil$. Since γ_i is strictly positive, we have $t_i > t_{i-1}$ and \mathcal{T} is an infinite set of times, hence Lemma 3.1 applies. Since the main components of the proof are unchanged, we only state the noteworthy modifications.

Lazy Update When considering the lazy update, both the candidate answer and the stopping criterion are computed only when $t \in \mathcal{T}$. Under \mathcal{E}_t , if the algorithm doesn't stop at time $t + 1$, there exists $i \in \mathbb{N}$ such that $t \in \llbracket t_i, t_{i+1} - 1 \rrbracket$ such that the stopping criterion is not met at time $t_i + 1$:

$$2\beta(t_i, \delta) \geq \inf_{\lambda \in \neg_\varepsilon z_{t_i+1}} \|\mu_{t_i} - \lambda\|_{V_{N_{t_i}}}^2 \geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_{t_i}}}^2 - h_\delta(t_i)$$

where h_δ is defined in Lemma E.5 and the second inequality is obtained as above since $\mathcal{E}_{t_i} \subseteq \mathcal{E}_t$ and $z_{t_i+1} = z_F(\mu_{t_i}, N_{t_i})$. Therefore, using the lazy update with an instantaneous furthest answer as recommendation rule allow to satisfy the condition (17) at time t_i .

Given a sampling rule satisfying the condition (18) at time t_i , the rest of the proof is exactly the same as sketched in Appendix E.1 and proved in Appendix E.3 with t_i instead of t . Therefore, we will obtain

$$H_\mu(t_i, \delta) \geq t_i T_\varepsilon(\mu)^{-1}$$

Since $t \in \llbracket t_i, t_{i+1} - 1 \rrbracket$, we have $H_\mu(t_i, \delta) \leq H_\mu(t, \delta)$ and $t_i \geq \frac{t_{i+1}-1}{1+\gamma_{i+1}} \geq \frac{t}{1+\gamma_{i+1}}$. This yields

$$\tilde{H}_\mu(t, \delta) \stackrel{\text{def}}{=} (1 + \gamma_{i+1})H_\mu(t_i, \delta) \geq t T_\varepsilon(\mu)^{-1}$$

We remark that $t \rightarrow +\infty$ implies $i \rightarrow +\infty$, hence $\gamma_{i+1} \rightarrow_{t \rightarrow +\infty} 0$. Choosing γ_i independently of δ and noting that $\tilde{H}_\mu(t, \delta) = \ln\left(\frac{1}{\delta}\right) + \mathcal{O}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right)$ since $H_\mu(t, \delta) = \ln\left(\frac{1}{\delta}\right) + \mathcal{O}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right)$, we can conclude the proof. Combining the lazy update with a good sampling rule and an instantaneous furthest answer as recommendation rule leads to asymptotically optimal algorithms.

For an expansion parameter which is constant, the manipulation above shows that the upper bound won't match the lower bound by a constant multiplicative factor $1 + \gamma_0$. Therefore, the shrinking expansion parameter is crucial to obtain asymptotic optimality.

Sticky Update When considering the sticky update, the candidate answer is computed only when $t \in \mathcal{T}$, while the stopping criterion is evaluated at each time t . Since the GLRTs conducted with the lazy update are strictly included in the ones conducted with the sticky update (with matching candidate answer), the sticky approach is strictly better in terms of sample complexity than considering the lazy one. Therefore, we have $\tau_\delta^{\text{sticky}} \leq \tau_\delta^{\text{lazy}}$, hence $\lim_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta^{\text{sticky}}]}{\ln(1/\delta)} \leq \lim_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta^{\text{lazy}}]}{\ln(1/\delta)}$. Using the above result showing that lazy update leads to asymptotic optimal algorithms, we conclude that sticky update also leads to asymptotic optimal algorithms.

E.3. Sampling Rule

In this appendix, we show that the desired condition (18) on the sampling rule holds when considering the sampling rule in $\mathbf{L}\varepsilon\mathbf{BAI}$ (Appendix E.3.2). Appendix E.3.1 details the two algorithmic requirements enforced in $\mathbf{L}\varepsilon\mathbf{BAI}$.

E.3.1. REQUIREMENTS

To conduct the analysis of the sampling rule used in $\mathbf{L}\varepsilon\mathbf{BAI}$, two algorithmic requirements were needed to be enforced. However, as experiments show in Appendix F.2, competitive empirical performance are obtained without them.

Forced Exploration To provide an upper bound on the number of times a good event doesn't occur (Lemma E.20), forced exploration has to be introduced by using $w_t = \frac{1}{tK} \mathbf{1}_K + \left(1 - \frac{1}{t}\right) w_t^{\mathcal{K}}$. Good events will be formally introduced below in (22). Since we are using tracking and $\sum_{s=n_0+1}^t \frac{1}{s} = \Theta(\ln t)$, the forced exploration only concerns a logarithmic number of steps.

Exact \mathcal{Z} -Oracle Likewise, to provide an upper bound on the number of times a good event doesn't occur (Lemma E.20), the \mathcal{Z} -oracle has to satisfy $\tilde{z}_t \in z_F(\mu_{t-1})$. Since this requirement is computationally intractable, it doesn't lead to a practical algorithm. In the experiments, we use an instantaneous furthest answer $\tilde{z}_t = z_t \in z_F(\mu_{t-1}, N_{t-1})$ which has good empirical behavior. Since it was already computed for the recommendation rule, no additional computation are needed.

While the proof still eludes us, we believe that using an instantaneous furthest answer would be also a theoretically valid choice for the \mathcal{Z} -oracle. Intuitively an instantaneous furthest answer is a good proxy for $z_F(\mu)$ because the empirical proportions will eventually converge to $w_F(\mu)$. Using our current proof techniques, it is difficult to show this result since the convergence properties and upper bound on the expected sample complexity are highly intertwined. Borrowing the proofs strategy of (Garivier & Kaufmann, 2016), it might be possible to show it with purely asymptotic arguments. We leave this interesting question to future work.

E.3.2. PROOF OF DESIRED PROPERTY

To highlight the structure of the proof, we prove separately two inequalities satisfied by the sampling rule used in $\mathbf{L}\varepsilon\mathbf{BAI}$ (Lemma E.6 and Lemma E.7). This intermediate step involves the cumulative gain of the learner used by the \mathcal{K} -player.

Lemma E.6. *Let $\mathcal{L}^{\mathcal{Z}}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. Let $b > 0$. For all $s \in \llbracket n_0, t-1 \rrbracket$ and $a \in \mathcal{K}$, let $(c_s^a)_{s \geq n_0, a \in \mathcal{K}}$ be defined by*

$$c_s^a \stackrel{\text{def}}{=} \min \left\{ f\left(s^{1+b}\right) \|a\|_{V_{N_s}^{-1}}^2, 4M^2 L_{\mathcal{K}}^2 \right\} \quad (20)$$

The optimistic gain $g_s(w) \stackrel{\text{def}}{=} \left(1 - \frac{1}{s}\right) \langle w, U_s \rangle$ are defined such that: for all $a \in \mathcal{K}$ and $s \in \llbracket n_0 + 1, t \rrbracket$,

$$U_s^a \stackrel{\text{def}}{=} \left(\|\mu_{s-1} - \lambda_s\|_{aa^\top} + \sqrt{c_{s-1}^a} \right)^2 \quad (21)$$

with $\lambda_s \stackrel{\text{def}}{=} \arg \min_{\lambda \in \neg_\varepsilon \tilde{z}_s} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2$. Then, under \mathcal{E}_t if $L\varepsilon\text{BAI}$ does not stop at time $t + 1$,

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq \sum_{s=n_0+1}^t g_s \left(w_s^{L\mathcal{K}} \right) - r_1(t),$$

where

$$r_1(t) = C_1 + \left(N_F(t) + t^{\frac{1}{1+b}} \right) 4KM^2 L_{\mathcal{K}}^2 + h'(t) + h''(t) = \Theta \left(\ln(t)\sqrt{t} + \ln(t)^2 t^{\frac{1}{1+b}} \right)$$

$$\text{with } h''(t) = f(t^{1+b}) \left(K \ln(K) + 2K \ln(t) \right) + 4ML_{\mathcal{K}} \sqrt{f(t^{1+b})} \left(K \ln(K) + \sqrt{8Kt} \right)$$

$$\text{and } h'(t) = 4ML_{\mathcal{K}} \sqrt{\left(K \ln(K) + 2K \ln(t) \right) f(t^{1+b})} t + 8M^2 L_{\mathcal{K}}^2 t^{\frac{1}{1+b}}.$$

Proof. Step 1. From N_t to W_t . Let $W_t \stackrel{\text{def}}{=} \sum_{s \in [t]} w_s$. Using Lemma E.15,

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq \max_{z \in z_F(\mu, W_t)} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2.$$

Recall that for all $N \in \mathbb{R}_+^K$, $\|\mu - \lambda\|_{V_N}^2 = \sum_{a \in \mathcal{K}} N^a \|\mu - \lambda\|_{a, \top}^2$. Using Cauchy-Schwartz and the bounded assumption, we have: $\|\mu - \lambda\|_{a, \top}^2 = \langle \mu - \lambda, a \rangle^2 \leq \|\mu - \lambda\|_2^2 \|a\|_2^2 \leq 4M^2 L_{\mathcal{K}}^2$. For any $\lambda \in \mathcal{M}$, using Lemma E.9,

$$\begin{aligned} \|\mu - \lambda\|_{V_{N_t}}^2 &\geq \|\mu - \lambda\|_{V_{W_t}}^2 - \ln(K) \sum_{a \in \mathcal{K}} \|\mu - \lambda\|_{a, \top}^2 \\ &\geq \|\mu - \lambda\|_{V_{W_t}}^2 - C_1, \end{aligned}$$

where $C_1 \stackrel{\text{def}}{=} 4 \ln(K) K M^2 L_{\mathcal{K}}^2 = \Theta(1)$.

Let $\tilde{z} \in \arg \max_{z \in z_F(\mu, W_t)} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2$ and $\tilde{\lambda} \in \arg \min_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_{N_t}}^2$. Using the above yields

$$\begin{aligned} \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 &\geq \|\mu - \tilde{\lambda}\|_{V_{W_t}}^2 - C_1 \\ &\geq \inf_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_{W_t}}^2 - C_1 \\ &= \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{W_t}}^2 - C_1, \end{aligned}$$

where we used that $\tilde{\lambda} \in \neg_\varepsilon \tilde{z}$ and $\tilde{z} \in z_F(\mu, W_t)$.

Step 2. From $(\mu, z_F(\mu, W_t))$ to (μ_{s-1}, \tilde{z}_s) . Using the concavity of \inf and $\|\mu - \lambda\|_{V_{W_t}}^2 = \sum_{s=n_0+1}^t \|\mu - \lambda\|_{V_{w_s}}^2$, we obtain

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{W_t}}^2 \geq \max_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{w_s}}^2.$$

Using the triangular inequality, for all $z \in \mathcal{Z}$ and $s \in [n_0 + 1, t]$

$$\|\mu - \lambda\|_{V_{w_s}}^2 \geq \left(\|\mu_{s-1} - \lambda\|_{V_{w_s}} - \|\mu_{s-1} - \mu\|_{V_{w_s}} \right)^2 \geq \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - 2\|\mu_{s-1} - \mu\|_{V_{w_s}} \|\mu_{s-1} - \lambda\|_{V_{w_s}}$$

For $\tilde{\lambda}_s^z \in \arg \min_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{w_s}}^2$, we obtain

$$\begin{aligned} \|\mu - \tilde{\lambda}_s^z\|_{V_{w_s}}^2 &\geq \|\mu_{s-1} - \tilde{\lambda}_s^z\|_{V_{w_s}}^2 - 2\|\mu_{s-1} - \mu\|_{V_{w_s}} \|\mu_{s-1} - \tilde{\lambda}_s^z\|_{V_{w_s}} \\ &\geq \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - 2\|\mu_{s-1} - \mu\|_{V_{w_s}} \|\mu_{s-1} - \tilde{\lambda}_s^z\|_{V_{w_s}} \\ &\geq \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - 4ML_{\mathcal{K}} \|\mu_{s-1} - \mu\|_{V_{w_s}} \end{aligned}$$

The second to last inequality is obtained by taking the $\inf_{\lambda \in \neg_\varepsilon z}$, which is possible since $\tilde{\lambda}_s^z \in \neg_\varepsilon z$, and the last one by upper bounding $\|\mu_{s-1} - \tilde{\lambda}_s\|_{V_{w_s}} \leq \|\mu_{s-1} - \lambda_s\|_2 \leq 2ML_{\mathcal{K}}$ (since $w_s \in \Delta_K$). Summing those inequalities together, using Cauchy-Schwartz when $s \geq \lceil t^{\frac{1}{1+b}} \rceil$ and the boundedness assumption when $s < \lceil t^{\frac{1}{1+b}} \rceil$, we obtain: for all $z \in \mathcal{Z}$

$$\begin{aligned} \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{w_s}}^2 &\geq \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - 4ML_{\mathcal{K}} \sum_{s=n_0+1}^t \|\mu_{s-1} - \mu\|_{V_{w_s}} \\ &\geq \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - 4ML_{\mathcal{K}} \sqrt{t - \lceil t^{\frac{1}{1+b}} \rceil} + 1 \sqrt{\sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2} \\ &\quad - 8M^2 L_{\mathcal{K}}^2 \left(\lceil t^{\frac{1}{1+b}} \rceil - n_0 - 1 \right) \\ &\geq \sum_{s=n_0+1}^t \|\mu_{s-1} - \lambda_s\|_{V_{w_s}}^2 - 4ML_{\mathcal{K}} \sqrt{t} \sqrt{\sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2} - 8M^2 L_{\mathcal{K}}^2 t^{\frac{1}{1+b}} \end{aligned}$$

The second inequality is obtained by concavity of $x \mapsto \sqrt{x}$, and the last one uses the definition of the best-response oracle in the sampling rule, $\lambda_s \in \arg \min_{\lambda \in \neg_\varepsilon z_s} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2$. Using Lemma E.25, we know that $\sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2 \leq f(t^{1+b}) (K \ln(K) + 2K \ln(t)) = \mathcal{O}(\ln(t)^2)$. Therefore, we obtain: for all $z \in \mathcal{Z}$

$$\sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{w_s}}^2 \geq \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - h'(t)$$

where $h'(t) \stackrel{\text{def}}{=} 4ML_{\mathcal{K}} \sqrt{(K \ln(K) + 2K \ln(t)) f(t^{1+b}) t} + 8M^2 L_{\mathcal{K}}^2 t^{\frac{1}{1+b}} = \Theta\left(\ln(t) \sqrt{t} + t^{\frac{1}{1+b}}\right)$. Taking the maximum on both sides for the above result and using Lemma E.21, we obtain by chaining the above inequalities

$$\begin{aligned} \max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 &\geq \max_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - C_1 - h'(t) \\ &\geq \sum_{s=n_0+1}^t \inf_{\lambda \in \neg_\varepsilon z_s} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2 - \left(N_F(t) + t^{\frac{1}{1+b}}\right) 4KM^2 L_{\mathcal{K}}^2 - C_1 - h'(t) \\ &= \sum_{s=n_0+1}^t \|\mu_{s-1} - \lambda_s\|_{V_{w_s}}^2 - \left(N_F(t) + t^{\frac{1}{1+b}}\right) 4KM^2 L_{\mathcal{K}}^2 - C_1 - h'(t), \end{aligned}$$

where $N_F(t)$ is defined in (24).

Step 3. From $\|\mu_{s-1} - \lambda_s\|_{V_{w_s}}^2$ to $g_s(w_s^{\mathcal{K}})$. Using Lemma E.17, we obtain: for all $s \in \llbracket n_0 + 1, t \rrbracket$ and $a \in \mathcal{K}$,

$$\|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 \geq U_s^a - c_{s-1}^a - 4ML_{\mathcal{K}} \sqrt{c_{s-1}^a}$$

Summing those inequalities together, we obtain:

$$\begin{aligned}
 & \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 \\
 & \geq \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a U_s^a - \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a c_{s-1}^a - 4ML_{\mathcal{K}} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \sqrt{c_{s-1}^a} \\
 & \geq \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a U_s^a - f(t^{1+b}) (K \ln(K) + 2K \ln(t)) - 4ML_{\mathcal{K}} \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt}) \\
 & \geq \sum_{s=n_0+1}^t \left(1 - \frac{1}{s}\right) \langle w_s^{\mathcal{L}^{\mathcal{K}}}, U_s \rangle - f(t^{1+b}) (K \ln(K) + 2K \ln(t)) \\
 & \quad - 4ML_{\mathcal{K}} \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt}) \\
 & = \sum_{s=n_0+1}^t g_s(w_s^{\mathcal{L}^{\mathcal{K}}}) - f(t^{1+b}) (K \ln(K) + 2K \ln(t)) - 4ML_{\mathcal{K}} \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt})
 \end{aligned}$$

where the second inequality is obtained by using Lemma E.24 and the third since the optimistic gains are positive and the weights are defined as $w_s = \frac{1}{sK} \mathbf{1}_K + (1 - \frac{1}{s}) w_s^{\mathcal{L}^{\mathcal{K}}}$ for all $s \in \llbracket n_0 + 1, t \rrbracket$. The last equality uses the definition of the gain $g_s(w) = (1 - \frac{1}{s}) \langle w, U_s \rangle$. Therefore, since $\|\mu_{s-1} - \lambda_s\|_{V_{w_s}}^2 = \sum_{a \in \mathcal{K}} w_s^a \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2$, we obtain that:

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq \sum_{s=n_0+1}^t g_s(w_s^{\mathcal{L}^{\mathcal{K}}}) - C_1 - (N_F(t) + t^{\frac{1}{1+b}}) 4KM^2 L_{\mathcal{K}}^2 - h'(t) - h''(t)$$

where $h''(t) = f(t^{1+b}) (K \ln(K) + 2K \ln(t)) + 4ML_{\mathcal{K}} \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt}) = \Theta(\sqrt{t \ln(t)})$. \square

Lemma E.7. Let $\mathcal{L}^{\mathcal{Z}}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. Let $\mathcal{L}^{\mathcal{K}}$ be a learner with regret $R_{\mathcal{L}^{\mathcal{K}}}(t) = \mathcal{O}(t^{\alpha_1})$ with $\alpha_1 \in (0, 1)$ and fed with optimistic gain $g_s(w) = (1 - \frac{1}{s}) \langle w, U_s \rangle$ where $(U_s)_{s \in \llbracket n_0+1, t \rrbracket}$ defined in (21). Then, under \mathcal{E}_t if $\mathcal{L}^{\mathcal{BAI}}$ does not stop at time $t + 1$,

$$\sum_{s=n_0+1}^t g_s(w_s^{\mathcal{L}^{\mathcal{K}}}) \geq 2tT_\varepsilon(\mu)^{-1} - r_2(t)$$

where

$$\begin{aligned}
 r_2(t) &= R_{\mathcal{L}^{\mathcal{K}}}(t) + h'''(t) = \mathcal{O}\left(\ln(t)^2 t^{\frac{1}{1+b}} + t^{\max\{\alpha_1, \frac{1}{1+b}, 1 - \frac{1}{1+b}\}}\right) \\
 \text{with } h'''(t) &= 2\left(t^{\frac{1}{1+b}} + N_F(t) + t^{1 - \frac{1}{1+b}}\right) T_\varepsilon(\mu)^{-1}
 \end{aligned}$$

Proof. **Step 4. No-regret property.** Using the no-regret property of the chosen online learner $\mathcal{L}^{\mathcal{K}}$ fed with gains $(g_s)_{s \geq n_0+1}$ and playing $(w_s^{\mathcal{L}^{\mathcal{K}}})_{s \geq n_0+1}$, whose regret is denoted by $R_{\mathcal{L}^{\mathcal{K}}}(t) = \mathcal{O}(t^{\alpha_1})$, we obtain directly

$$\sum_{s=n_0+1}^t g_s(w_s^{\mathcal{L}^{\mathcal{K}}}) \geq \max_{w \in \Delta_K} \sum_{s=n_0+1}^t g_s(w) - R_{\mathcal{L}^{\mathcal{K}}}(t)$$

Step 5. From the optimal gain to $T_\varepsilon(\mu)^{-1}$. Dropping the first positive terms and the ones for which the good event $A_s \stackrel{\text{def}}{=} \{\lambda_s \in \neg_\varepsilon z_F(\mu) \vee \mu_{s-1} \in \neg_\varepsilon z_F(\mu) \vee \tilde{z}_s = z_F(\mu)\}$ doesn't hold, i.e. $\llbracket n_0 + 1, \lceil t^{\frac{1}{1+b}} \rceil - 1 \rrbracket \cup \left\{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket : \neg A_s\right\}$,

and using Lemma E.18, we obtain: for all $w \in \Delta_K$,

$$\begin{aligned} \sum_{s=n_0+1}^t g_s(w) &\geq \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \mathbf{1}_{(A_s)} g_s(w) \geq \left(1 - t^{-\frac{1}{1+b}}\right) \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \mathbf{1}_{(A_s)} \langle w, U_s \rangle \\ &\geq \left(1 - t^{-\frac{1}{1+b}}\right) \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \mathbf{1}_{(A_s)} \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2 \\ &\geq \left(1 - t^{-\frac{1}{1+b}}\right) \left(t - t^{\frac{1}{1+b}} - N_F(t)\right) \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2 \end{aligned}$$

where $\left| \left\{ s \in \left[\lceil t^{\frac{1}{1+b}} \rceil, t \right] : \neg A_s \right\} \right| \leq N_F(t)$ since $\{\tilde{z}_s = z_F(\mu)\} \subseteq A_s$, where $N_F(t)$ defined in (24) satisfies $N_F(t) = \mathcal{O}\left(\ln(t)^2 t^{\frac{1}{1+b}}\right)$ (Lemma E.20). For the last inequality, we also used that $-\lceil t^{\frac{1}{1+b}} \rceil + 1 \geq -t^{\frac{1}{1+b}}$. Taking the maximum over Δ_K , we obtain:

$$\begin{aligned} \max_{w \in \Delta_K} \sum_{s=n_0+1}^t g_s(w) &\geq \left(1 - t^{-\frac{1}{1+b}}\right) \left(t - t^{\frac{1}{1+b}} - N_F(t)\right) \max_{w \in \Delta_K} \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2 \\ &= 2 \left(1 - t^{-\frac{1}{1+b}}\right) \left(t - t^{\frac{1}{1+b}} - N_F(t)\right) T_\varepsilon(\mu)^{-1} \\ &\geq 2 \left(t - t^{\frac{1}{1+b}} - N_F(t) - t^{1-\frac{1}{1+b}}\right) T_\varepsilon(\mu)^{-1} \end{aligned}$$

where the equality is obtained by definition of $z_F(\mu)$. The last inequality is obtained by dropping some positive terms. Putting everything together, we obtain:

$$\sum_{s=n_0+1}^t g_s(w_s^{\mathcal{L}^\kappa}) \geq 2tT_\varepsilon(\mu)^{-1} - R_{\mathcal{L}^\kappa}(t) - h'''(t)$$

where $h'''(t) = 2 \left(t^{\frac{1}{1+b}} + N_F(t) + t^{1-\frac{1}{1+b}}\right) T_\varepsilon(\mu)^{-1} = \mathcal{O}\left(\ln(t)^2 t^{\frac{1}{1+b}} + t^{\max\{\frac{1}{1+b}, 1-\frac{1}{1+b}\}}\right)$. \square

Chaining the two inequalities in Lemma E.6 and Lemma E.7, we obtain directly Lemma E.8. Therefore, the sampling rule in $\mathbf{L}_\varepsilon\mathbf{BAI}$ satisfies the desired condition (18) for $\beta_2 = 1$ and $\beta_1 = 1 - \max\left\{\frac{1}{2}, \frac{1}{1+b}, 1 - \frac{1}{1+b}, \alpha_1\right\}$.

Lemma E.8. *Let $\mathcal{L}^\mathcal{Z}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. Let \mathcal{L}^κ be a learner with regret $R_{\mathcal{L}^\kappa}(t) = \mathcal{O}(t^{\alpha_1})$ with $\alpha_1 \in (0, 1)$ fed with optimistic gain $g_s(w) = \left(1 - \frac{1}{s}\right) \langle w, U_s \rangle$ where $(U_s)_{s \in \llbracket n_0+1, t \rrbracket}$ defined in (21). Then, under \mathcal{E}_t if the algorithms does not stop at time $t + 1$,*

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 \geq 2tT_\varepsilon(\mu)^{-1} - r(t)$$

where $r(t) = r_1(t) + r_2(t) = \tilde{\mathcal{O}}\left(t^{\max\{\frac{1}{2}, \frac{1}{1+b}, 1-\frac{1}{1+b}, \alpha_1\}}\right)$.

E.4. Summarized Proof

Combining the lemmas obtained in Appendix E.1, E.2 and E.3, we obtain directly Theorem 4.1. Therefore, our algorithm $\mathbf{L}_\varepsilon\mathbf{BAI}$ provides asymptotically optimal algorithms if instantiated properly.

Theorem (Theorem 4.1). Let \mathcal{L}^κ be a learner with sub-linear regret and $\mathcal{L}^\mathcal{Z}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. Let the recommendation rule be an instantaneous furthest answer $z_t \in z_F(\mu_{t-1}, N_{t-1})$ and the stopping rule given by (5), both being either evaluated at each time t or with the sticky or lazy update schedule as described in Appendix E.2.1. Let the stopping threshold $\beta(t, \delta)$ as in (6) and the exploration bonus $f(t) \stackrel{\text{def}}{=} 2\beta(t, t^{1/\alpha})$. Then, $\mathbf{L}_\varepsilon\mathbf{BAI}$ yields an (ε, δ) -PAC algorithm and, for all $\mu \in \mathcal{M}$ such that $|z_F(\mu)| = 1$,

$$\limsup_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln\left(\frac{1}{\delta}\right)} \leq T_\varepsilon(\mu).$$

Proof. First, let's prove the result for the stopping-recommendation pair with an update at each time t . Combining Lemma E.5 and Lemma E.8, we obtain:

$$\begin{aligned} 2\beta(t, \delta) &\geq \max_{z \in \mathcal{Z}} \inf_{\lambda \in \tau_\varepsilon z} \|\mu - \lambda\|_{V_{N_t}}^2 - h_\delta(t) \\ &\geq 2tT_\varepsilon(\mu)^{-1} - r(t) - h_\delta(t) \end{aligned}$$

where $h_\delta(t) = \sqrt{8f(t)\beta(t, \delta)} + 4f(t) = \Theta\left(\ln(t) + \sqrt{\ln\left(\frac{1}{\delta}\right)}\right)$ and $r(t) = r_1(t) + r_2(t) = \tilde{\mathcal{O}}\left(t^{\max\{\frac{1}{2}, \frac{1}{1+b}, 1 - \frac{1}{1+b}, \alpha_1\}}\right)$. Therefore, our identification strategy is sample-efficient, i.e. it verifies (11) with

$$H_\mu(t, \delta) \geq tT_\varepsilon(\mu)^{-1} \quad \text{where } H_\mu(t, \delta) = \beta(t, \delta) + \frac{1}{2}(r(t) + h_\delta(t)) = \mathcal{O}\left(t^{1-\beta_1} + \ln(1/\delta)^{1-\beta_2}\right)$$

where $\beta_1 = 1 - \max\left\{\frac{1}{2}, \frac{1}{1+b}, 1 - \frac{1}{1+b}, \alpha_1\right\} \in (0, 1]$ and $\beta_2 = \frac{1}{2} \in (0, 1]$.

Combining Lemma 3.1 and Lemma E.4, we obtain directly that our $\mathcal{L}\varepsilon$ BAI yields a (ε, δ) -PAC and asymptotically optimal algorithm.

When considering the computational relaxations of the update schemes called sticky or lazy, the proof is identical as explained in Appendix E.2.1 and the fact that Lemma 3.1 can also be used. \square

Values of Parameters In the proof above, the condition (11) holds for $\beta_1 = 1 - \max\left\{\frac{1}{2}, \frac{1}{1+b}, 1 - \frac{1}{1+b}, \alpha_1\right\}$ and $\beta_2 = \frac{1}{2}$. For AdaHedge (De Rooij et al., 2014), Lemma E.14 shows that $\alpha_1 = \frac{1}{2}$. An optimal choice of parameter suggests that $\frac{1}{1+b} = 1 - \frac{1}{1+b}$. Therefore, we choose $b = 1$. Higher b would imply more conservative optimistic gains. Therefore, we have shown that the condition (11) holds for $\beta_1 = \beta_2 = \frac{1}{2}$. As done in Degenne et al. (2020a), we can choose $\alpha = 3$ in the definition of the exploration bonus $f(t) = 2\beta(t, t^{1/\alpha})$. The higher α , the higher the exploration bonus is.

E.5. Technical Arguments

In this appendix, we list and prove the technical arguments used in the core of the proof of Theorem 4.1. For clarity, we distinguish between the lemmas extracted from the literature (Appendix E.5.1), the key new lemmas (Appendix E.5.2) and the technical ones allowing to upper bound the considered cumulative sums (Appendix E.5.3).

E.5.1. LEMMAS FROM THE LITERATURE

We recall the lemmas extracted from the literature, while omitting the proofs for the sake of space. Lemma E.9 is a powerful result allowing to upper and lower bound the difference between the empirical allocation over arms and the cumulative sum of the played proportions when using tracking. Lemma E.10 and Lemma E.11 are technical results on manipulation of inequalities and cumulative sums. Lemma E.13 shows that $\mu \mapsto z_F(\mu)$ is upper hemicontinuous on \mathcal{M} . Upper hemicontinuity is defined in Definition E.12. Lemma E.14 gives an upper bound on the regret incurred by AdaHedge.

Lemma E.9 (Theorem 6 in Degenne et al. (2020b)). *The tracking procedure, which draws $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - W_t^a$ where $W_t = W_{t-1} + w_t$, ensures that for all $t \in \mathbb{N}$, for all $a \in \mathcal{K}$,*

$$-\sum_{j=2}^K \frac{1}{j} \leq N_t^a - W_t^a \leq 1$$

In practice, we will use the following slightly coarser lower bound $-\ln(K) \leq N_t^a - W_t^a$.

Lemma E.10 (Lemma 28 in Degenne et al. (2020a)). *For all $\alpha, y > 0$, if for some $x \geq 0$, it holds that $y \geq x - \alpha\sqrt{x}$, then $x \leq y + \alpha\sqrt{y} + \alpha^2$.*

Lemma E.11 (Lemma 8 in Degenne et al. (2019)). *For $t \geq t_0 \geq 1$ and $(x_s)_{s \in [t]}$ non-negative real numbers such that*

$$\sum_{s=1}^{t_0-1} x_s > 0,$$

$$\begin{aligned} \sum_{s=t_0}^t \frac{x_s}{\sum_{r=1}^s x_r} &\leq \ln \left(\sum_{s=1}^t x_s \right) - \ln \left(\sum_{s=1}^{t_0-1} x_s \right) \\ \sum_{s=t_0}^t \frac{x_s}{\sqrt{\sum_{r=1}^s x_r}} &\leq 2\sqrt{\sum_{s=1}^t x_s} - 2\sqrt{\sum_{s=1}^{t_0-1} x_s} \end{aligned}$$

Definition E.12 (Upper hemicontinuity). For a set B , let $\mathbb{S}(B) = 2^B \setminus \{\emptyset\}$ be the set of all non-empty subsets of B . A set-valued function $\Gamma : A \mapsto \mathbb{S}(B)$ is upper hemicontinuous at $a \in A$ if for any open neighborhood V of $\Gamma(a)$ there exists a neighborhood U of a such that for all $x \in U$, $\Gamma(x)$ is a subset of V .

Lemma E.13 (Theorem 4 in Degenne & Koolen (2019)). *The function $\mu \mapsto z_F(\mu)$ is upper hemicontinuous on \mathbb{R}^d with non-empty and compact values.*

Proof. Theorem 4 in Degenne & Koolen (2019) proves that the function $\mu \mapsto z_F(\mu)$ is upper hemicontinuous on \mathcal{M} with non-empty and compact values. In their paper the function $\mu \mapsto z_F(\mu)$ was only defined on \mathcal{M} . However, a careful analysis of their Appendix D reveals that their result and proof apply without any change when $\mu \mapsto z_F(\mu)$ is defined on \mathbb{R}^d . Since in our setting $z_F(\mu)$ is defined on \mathbb{R}^d , this concludes the proof. Note that while $\mu \mapsto z_F(\mu)$ is defined on \mathcal{M} , the alternatives $(\neg_\varepsilon z)_{z \in \mathcal{Z}}$ are all subsets of $\overline{\mathcal{M}}$. \square

Lemma E.14. *On the online learning problem with K arms and gains $g_s(w) = (1 - \frac{1}{s})\langle w, U_s \rangle$ for $s \in \llbracket n_0 + 1, t \rrbracket$, AdaHedge predicting $(w_s^{\mathcal{L}^\mathcal{K}})_{s \in \llbracket n_0 + 1, t \rrbracket}$ has regret*

$$R_{\mathcal{L}^\mathcal{K}}(t) = \max_{w \in \Delta_K} \sum_{s=n_0+1}^t g_s(w) - g_s(w_s^{\mathcal{L}^\mathcal{K}}) \leq \left(\sqrt{t \ln(K)} + \frac{4}{3} \ln(K) + 2 \right) \sigma = \mathcal{O}(\sqrt{t})$$

where $\sigma = \max_{s \in \llbracket n_0 + 1, t \rrbracket} (1 - \frac{1}{s}) (\max_{a \in \mathcal{K}} U_s^a - \min_{a \in \mathcal{K}} U_s^a)$.

Proof. For scaled losses in $[0, 1]$, Theorem 6 in De Rooij et al. (2014) yields that AdaHedge's cumulative regret R_t satisfies: $R_t \leq 2\sqrt{V_t \ln(K)} + \frac{4}{3} \ln(K) + 2$ where $V_t = \sum_{s \in [t]} v_s$ with $v_s = \sum_{a \in \mathcal{K}} w_{s,a}^{\mathcal{L}^\mathcal{K}} (l_s^a - \langle w_s^{\mathcal{L}^\mathcal{K}}, l_s \rangle)^2$.

Rewriting our gains as losses, we have $l_s^a = (1 - \frac{1}{s}) \frac{\max_{b \in \mathcal{K}} U_s^b - U_s^a}{b_s}$ where $b_s = (1 - \frac{1}{s}) (\max_{a \in \mathcal{K}} U_s^a - \min_{a \in \mathcal{K}} U_s^a)$ is the scale of the loss at time s . We have $v_s \leq \|l_s - \langle w_s^{\mathcal{L}^\mathcal{K}}, l_s \rangle\|_\infty^2 = (1 - \frac{1}{s})^2 \frac{\| \langle w_s^{\mathcal{L}^\mathcal{K}}, U_s \rangle - U_s \|_\infty^2}{b_s^2} \leq \frac{b_s^2}{\sigma^2}$ where $\sigma = \max_{s \in \llbracket n_0 + 1, t \rrbracket} b_s$. The upper bound on R_t rewrites as: $R_t \leq \frac{1}{\sigma} \sqrt{\sum_{s \leq t} b_s^2 \ln(K)} + \frac{4}{3} \ln(K) + 2$. Theorem 16 in De Rooij et al. (2014) yields that $R_{\mathcal{L}^\mathcal{K}}(t) = \sigma R_t$. Therefore, we conclude that:

$$R_{\mathcal{L}^\mathcal{K}}(t) \leq \sqrt{\sum_{s \in \llbracket n_0 + 1, t \rrbracket} b_s^2 \ln(K)} + \left(\frac{4}{3} \ln(K) + 2 \right) \max_{s \in \llbracket n_0 + 1, t \rrbracket} b_s \leq \left(\sqrt{t \ln(K)} + \left(\frac{4}{3} \ln(K) + 2 \right) \right) \max_{s \in \llbracket n_0 + 1, t \rrbracket} b_s$$

Using the definition of U_s^a and Lemma E.17, we obtain:

$$b_s \leq \max_{a \in \mathcal{K}} U_s^a \leq \max_{a \in \mathcal{K}} (\|\mu_{s-1} - \lambda_s\|_{aa}^2 + c_{s-1}^a + 4ML_{\mathcal{K}} \sqrt{c_{s-1}^a}) \leq 16M^2 L_{\mathcal{K}}^2$$

where the last inequality is obtained by the boundedness assumption and the definition of c_s^a . \square

We also recall a trivial result obtained by comparison of sum and integrals: for all $\alpha \in (0, 1)$ and all $t \in \mathbb{N}^*$, $\sum_{s \in [t]} \frac{1}{s^\alpha} = \Theta(t^{1-\alpha})$.

E.5.2. KEY LEMMAS

The following lemmas are key arguments used throughout Appendix E. Lemma E.15 is a simple, yet powerful, result allowing to change the considered alternatives. This subtlety arises only because we are considering multiple correct answers.

Lemma E.15. For all $w \in \mathbb{R}_+^K$, $\mu \in \mathcal{M}$ and $\tilde{z} \in \mathcal{Z}$,

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_w}^2 \geq \inf_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_w}^2.$$

In particular, for all $(w, \tilde{w}) \in (\mathbb{R}_+^K)^2$ and $(\mu, \tilde{\mu}) \in \mathcal{M}^2$,

$$\max_{z \in \mathcal{Z}} \inf_{\lambda \in \neg_\varepsilon z} \|\mu - \lambda\|_{V_w}^2 \geq \max_{\tilde{z} \in \mathcal{Z}_F(\tilde{\mu}, \tilde{w})} \inf_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_w}^2.$$

Proof. **Case 1:** $\tilde{z} \notin \mathcal{Z}_\varepsilon(\mu)$, hence $\mu \in \neg_\varepsilon \tilde{z}$. Therefore $\inf_{\lambda \in \neg_\varepsilon \tilde{z}} \|\mu - \lambda\|_{V_w}^2 = 0$.

Case 2: $\tilde{z} \in \mathcal{Z}_\varepsilon(\mu)$. The result is direct by taking the maximum. \square

Lemma E.16 shows that the slacks $(c_s^a)_{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t-1 \rrbracket, a \in \mathcal{K}}$ are an upper bound of the unknown distance $\|\mu - \mu_s\|_{aa^\top}^2$.

Lemma E.16. Let $(c_s^a)_{s \in \llbracket n_0, t-1 \rrbracket, a \in \mathcal{K}}$ in (20). Under \mathcal{E}_t , for all $(s, a) \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t-1 \rrbracket \times \mathcal{K}$, $\|\mu - \mu_s\|_{aa^\top}^2 \leq c_s^a$.

Proof. Since $(\mu, \mu_s) \in \mathcal{M}^2$, the boundedness yields: $\|\mu - \mu_s\|_{aa^\top}^2 = \langle \mu - \mu_s, a \rangle^2 \leq \|\mu - \mu_s\|_2^2 \|a\|_2^2 \leq 4M^2 L_{\mathcal{K}}^2$.

By definition, under \mathcal{E}_t , for all $s \in [t]$, we have $\|\mu_s - \mu\|_{V_{N_s}}^2 \leq f(t)$. Since $x \mapsto f(x)$ is increasing, for all $s \geq \lceil t^{\frac{1}{1+b}} \rceil$, we have $f(t) \leq f(s^{1+b})$. By initialization V_{N_s} is positive definite, using Cauchy-Schwartz yields:

$$\langle \mu - \mu_s, a \rangle^2 \leq \|\mu_s - \mu\|_{V_{N_s}}^2 \|a\|_{V_{N_s}^{-1}}^2 \leq f(s^{1+b}) \|a\|_{V_{N_s}^{-1}}^2$$

Combining both upper bound and using the definition of $(c_s^a)_{s \in \llbracket n_0, t-1 \rrbracket, a \in \mathcal{K}}$ in (20) conclude the proof. \square

Lemma E.17 shows useful upper and lower bounds on the optimistic gains $(U_s^a)_{s \in \llbracket n_0+1, t \rrbracket, a \in \mathcal{K}}$.

Lemma E.17. Let $b > 0$, $(c_{s-1}^a)_{s \in \llbracket n_0+1, t \rrbracket, a \in \mathcal{K}}$ and $(U_s^a)_{s \in \llbracket n_0+1, t \rrbracket, a \in \mathcal{K}}$ defined in (20-21). For all $s \in \llbracket n_0 + 1, t \rrbracket$, for all $a \in \mathcal{K}$,

$$\|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 \leq U_s^a \leq \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 + c_{s-1}^a + 4ML_{\mathcal{K}} \sqrt{c_{s-1}^a}.$$

For all $s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket$, for all $a \in \mathcal{K}$, $U_s^a \geq \|\mu_{s-1} - \mu\|_{aa^\top}^2$ and $U_s^a \geq \|\mu - \lambda_s\|_{aa^\top}^2$.

Proof. By definition of U_s^a in (21) and using the boundedness to show $\|\mu_{s-1} - \lambda_s\|_{aa^\top} \leq 2ML_{\mathcal{K}}$, we have: for all $s \in \llbracket n_0 + 1, t \rrbracket$, for all $a \in \mathcal{K}$,

$$\begin{aligned} U_s^a &= \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 + c_{s-1}^a + 2\|\mu_{s-1} - \lambda_s\|_{aa^\top} \sqrt{c_{s-1}^a} \geq \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 \\ U_s^a &\leq \|\mu_{s-1} - \lambda_s\|_{aa^\top}^2 + c_{s-1}^a + 4ML_{\mathcal{K}} \sqrt{c_{s-1}^a} \end{aligned}$$

By definition of U_s^a in (21), Lemma E.16 and triangular inequality, we have: for all $s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket$, for all $a \in \mathcal{K}$,

$$\begin{aligned} U_s^a &= (\|\mu_{s-1} - \lambda_s\|_{aa^\top} + \sqrt{c_{s-1}^a})^2 \geq (\|\mu_{s-1} - \lambda_s\|_{aa^\top} + \|\mu_{s-1} - \mu\|_{aa^\top})^2 \geq \|\mu - \lambda_s\|_{aa^\top}^2 \\ U_s^a &\geq c_{s-1}^a \geq \|\mu_{s-1} - \mu\|_{aa^\top}^2 \end{aligned}$$

\square

Lemma E.18 shows that under a good event A_s the gain $\langle w, U_s \rangle$ is an upper bound on the unknown $\inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2$ for all $w \in \Delta_K$. This good event essentially boils down to having as candidate answer the unknown furthest answer $\tilde{z}_s = z_F(\mu)$.

Lemma E.18. *Let $b > 0$ as in Lemma E.17. Let $(A_s)_{s \in \left[\left[t^{\frac{1}{1+b}} \right], t \right]}$ defined as: for all $s \in \left[\left[t^{\frac{1}{1+b}} \right], t \right]$,*

$$A_s \stackrel{\text{def}}{=} \{ \lambda_s \in \neg_\varepsilon z_F(\mu) \vee \mu_{s-1} \in \neg_\varepsilon z_F(\mu) \vee \tilde{z}_s = z_F(\mu) \} \quad (22)$$

Under \mathcal{E}_t , we have for all $w \in \Delta_K$ and $s \in \left[\left[t^{\frac{1}{1+b}} \right], t \right]$, under A_s it holds that

$$\langle w, U_s \rangle \geq \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2$$

Proof. Summing the lower bound of Lemma E.17 yields: for all $s \in \left[\left[t^{\frac{1}{1+b}} \right], t \right]$ and all $w \in \Delta_K$, $\langle w, U_s \rangle \geq \|\mu - \lambda_s\|_{V_w}^2$ and $\langle w, U_s \rangle \geq \|\mu - \mu_{s-1}\|_{V_w}^2$.

Case 1: $\lambda_s \in \neg_\varepsilon z_F(\mu)$. Taking the infimum on the first inequality, we have: for all $s \in \left[\left[t^{\frac{1}{1+b}} \right], t \right]$ and all $w \in \Delta_K$,

$$\langle w, U_s \rangle \geq \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2$$

Case 2: $\mu_{s-1} \in \neg_\varepsilon z_F(\mu)$. Taking the infimum on the second inequality, we obtain directly: for all $w \in \Delta_K$ and $s \geq t^{\frac{1}{1+b}}$,

$$\langle w, U_s \rangle \geq \inf_{\lambda \in \neg_\varepsilon z_F(\mu)} \|\mu - \lambda\|_{V_w}^2$$

Case 3: $\tilde{z}_s = z_F(\mu)$. By definition $\lambda_s \in \neg_\varepsilon \tilde{z}_s$, hence $\lambda_s \in \neg_\varepsilon z_F(\mu)$. Therefore, we can use the first case. \square

We introduce a strictly positive geometrical quantity Δ_{\min}^2 (23) in Lemma E.19. This result crucially uses that \mathcal{Z} is finite and $|z_F(\mu)| = 1$ for the unknown parameter μ . It means that we can't be arbitrarily close to μ while having a unique furthest answer which is different from $z_F(\mu)$.

Δ_{\min}^2 represents the distance of μ to the *furthest alternative* to $z_F(\mu)$, denoted by $\neg_F z_F(\mu)$. The furthest alternative is fundamentally different from the alternative $\neg_\varepsilon z_F(\mu)$ (which is considered in the rest of the paper). While the furthest alternative corresponds to an identification problem with a unique correct answer, the alternative corresponds to an identification problem with multiple correct answers. Therefore, we see that to achieve asymptotic optimality in this multiple correct answer setting, we need to solve the harder problem of exactly identifying the unique $z_F(\mu)$.

Lemma E.19. *Defining*

$$\Delta_{\min}^2 \stackrel{\text{def}}{=} \inf_{\lambda \in \neg_F z_F(\mu)} \max_{a \in \mathcal{K}} \|\mu - \lambda\|_{aa^\top}^2 > 0 \quad (23)$$

where $\neg_F z_F(\mu) \stackrel{\text{def}}{=} \{ \lambda \in \mathbb{R}^d : z_F(\mu) \neq z_F(\lambda) \}$. Then, we have: $\Delta_{\min}^2 > 0$.

Proof. By Lemma E.13, we know that $\lambda \mapsto z_F(\lambda)$ is upper hemicontinuous on \mathbb{R}^d with non-empty and compact values. Since \mathcal{Z} is finite and $z_F(\mu)$ is a singleton by assumption on μ , the upper hemicontinuity implies that there exists an open neighborhood U of μ such that for all $\tilde{\mu} \in U$, $z_F(\tilde{\mu}) = z_F(\mu)$. This is obtained directly by using Definition E.12 and taking an open neighborhood V of $z_F(\mu)$ such that $V \cap \mathcal{Z} = z_F(\mu)$, which exists since $|z_F(\mu)| = 1$ and \mathcal{Z} finite. Noticing that $U \subseteq (\neg_F z_F(\mu))^c$ and using that U is an open neighborhood of μ , we can conclude that $\Delta_{\min}^2 > 0$. \square

Lemma E.20 shows that the good event $\{ \tilde{z}_s = z_F(\mu) \} \subseteq A_s$ doesn't happen only for a sub-linear number of times. A similar argument was shown in Lemma 25 of [Degegne et al. \(2020a\)](#). This crucially uses the fact that $\Delta_{\min}^2 > 0$.

Lemma E.20. Let $\mathcal{L}^{\mathcal{Z}}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. For all $s \in \left[\left\lceil t^{\frac{1}{1+b}} \right\rceil, t \right]$, let

$$B_s \stackrel{\text{def}}{=} \{\tilde{z}_s = z_F(\mu)\} \quad \text{and} \quad N_F(t) \stackrel{\text{def}}{=} \left| \left\{ s \in \left[\left\lceil t^{\frac{1}{1+b}} \right\rceil, t \right] : \neg B_s \right\} \right|. \quad (24)$$

Then, under \mathcal{E}_t ,

$$N_F(t) \leq \frac{K}{\Delta_{\min}^2} t^{\frac{1}{1+b}} f(t^{1+b}) (K \ln(K) + 2K \ln(t)) = \mathcal{O}\left(\ln(t)^2 t^{\frac{1}{1+b}}\right)$$

Proof. Using Lemma E.25, under \mathcal{E}_t ,

$$\begin{aligned} f(t^{1+b}) (K \ln(K) + 2K \ln(t)) &\geq \sum_{s=\left\lceil t^{\frac{1}{1+b}} \right\rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2 \geq \sum_{s=\left\lceil t^{\frac{1}{1+b}} \right\rceil}^t \mathbf{1}_{(\neg B_s)} \|\mu_{s-1} - \mu\|_{V_{w_s}}^2 \\ &\geq \sum_{s=\left\lceil t^{\frac{1}{1+b}} \right\rceil}^t \mathbf{1}_{(\neg B_s)} \frac{1}{sK} \sum_{a \in \mathcal{K}} \|\mu_{s-1} - \mu\|_{aa^\top}^2 \end{aligned}$$

where the last inequality is obtained since $w_s = \frac{1}{sK} \mathbf{1}_K + \left(1 - \frac{1}{s}\right) w_s^{\mathcal{K}}$ for all $s \in \llbracket n_0 + 1, t \rrbracket$.

Since $B_s = \{\tilde{z}_s = z_F(\mu)\}$ and $\tilde{z}_s \in z_F(\mu_{s-1})$, we obtain:

$$\neg B_s = \{\tilde{z}_s \notin z_F(\mu), \tilde{z}_s \in z_F(\mu_{s-1})\} \subseteq \{z_F(\mu_{s-1}) \neq z_F(\mu)\}$$

where in the above $z_F(\mu)$ is viewed as the set containing one element. Therefore, since $\mu_{s-1} \in \mathbb{R}^d$ and $z_F(\mu_{s-1}) \neq z_F(\mu)$, we have shown that $\mu_{s-1} \in \neg_F z_F(\mu)$. By the definition of Δ_{\min}^2 in (23), there exists $a_s \in \mathcal{K}$ such that $\|\mu_{s-1} - \mu\|_{a_s a_s^\top}^2 \geq \Delta_{\min}^2$. This yields $\sum_{a \in \mathcal{K}} \|\mu_{s-1} - \mu\|_{aa^\top}^2 \geq \Delta_{\min}^2$. Therefore, we obtain: $\frac{1}{s} \sum_{a \in \mathcal{K}} \|\mu_{s-1} - \mu\|_{aa^\top}^2 \geq \frac{\Delta_{\min}^2}{s} \geq \frac{\Delta_{\min}^2}{t^{\frac{1}{1+b}}}$ for all $s \in \left[\left\lceil t^{\frac{1}{1+b}} \right\rceil, t \right]$. Putting everything together, we have shown that:

$$N_F(t) \leq \frac{K}{\Delta_{\min}^2} t^{\frac{1}{1+b}} f(t^{1+b}) (K \ln(K) + 2K \ln(t))$$

□

Lemma E.21. Let $\mathcal{L}^{\mathcal{Z}}$ be the \mathcal{Z} -oracle such that $\tilde{z}_s \in z_F(\mu_{s-1})$. Then, under \mathcal{E}_t ,

$$\min_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t [f_s(\tilde{z}_s) - f_s(z)] \leq \left(N_F(t) + t^{\frac{1}{1+b}} \right) 4KM^2 L_{\mathcal{K}}^2 = \mathcal{O}\left(\ln(t)^2 t^{\frac{1}{1+b}}\right),$$

where $f_s(z) = \inf_{\lambda \in \neg_\varepsilon z} \|\mu_{s-1} - \lambda\|_{V_{w_s}}^2$.

Proof. Let B_s as in (24) for all $s \in \left[\left\lceil t^{\frac{1}{1+b}} \right\rceil, t \right]$. Since $f_s(z) \geq 0$, we can drop the the first positive terms and the ones for

which the good event B_s doesn't hold, i.e. $\llbracket n_0 + 1, \lceil t^{\frac{1}{1+b}} \rceil - 1 \rrbracket \cup \left\{ s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket : \neg B_s \right\}$, hence

$$\begin{aligned}
 & \min_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t [f_s(\tilde{z}_s) - f_s(z)] \\
 & \leq \min_{z \in \mathcal{Z}} \sum_{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket} \mathbf{1}_{(B_s)} [f_s(\tilde{z}_s) - f_s(z)] + \sum_{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket} \mathbf{1}_{(\neg B_s)} f_s(\tilde{z}_s) + \sum_{s \in \llbracket n_0+1, \lceil t^{\frac{1}{1+b}} \rceil - 1 \rrbracket} f_s(\tilde{z}_s) \\
 & = \min_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t \mathbf{1}_{(B_s)} [f_s(z_F(\mu)) - f_s(z)] + \sum_{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket} \mathbf{1}_{(\neg B_s)} f_s(\tilde{z}_s) + \sum_{s \in \llbracket n_0+1, \lceil t^{\frac{1}{1+b}} \rceil - 1 \rrbracket} f_s(\tilde{z}_s) \\
 & \leq \sum_{s \in \llbracket \lceil t^{\frac{1}{1+b}} \rceil, t \rrbracket} \mathbf{1}_{(\neg B_s)} f_s(\tilde{z}_s) + \sum_{s \in \llbracket n_0+1, \lceil t^{\frac{1}{1+b}} \rceil - 1 \rrbracket} f_s(\tilde{z}_s)
 \end{aligned}$$

where we used that the best constant policy for (f_s) is by definition better than the constant policy playing $z_F(\mu)$. By boundedness assumption, we have $f_s(\tilde{z}_s) \leq 4KM^2L_K^2$ for all s . Therefore, under \mathcal{E}_t ,

$$\min_{z \in \mathcal{Z}} \sum_{s=n_0+1}^t [f_s(\tilde{z}_s) - f_s(z)] \leq \left(N_F(t) + t^{\frac{1}{1+b}} \right) 4KM^2L_K^2 = \mathcal{O} \left(\ln(t)^2 t^{\frac{1}{1+b}} \right).$$

□

E.5.3. UPPER BOUNDING CUMULATIVE SUMS

Lemma E.22 is very similar to Lemma 9 in [Degenne et al. \(2019\)](#) and Lemma 15 in [Jourdan et al. \(2021\)](#), the novelty lies in using Lemma E.9 to obtain tighter upper bounds. It gives an upper bound on the cumulative sum of the reweighted inverse of the empirical allocations.

Lemma E.22. *The tracking procedure, which draws $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - W_t^a$ where $W_t = W_{t-1} + w_t$, ensures that for all $t \in \mathbb{N}$,*

$$\begin{aligned}
 \sum_{a \in \mathcal{K}} \sum_{s=n_0}^t \frac{w_s^a}{N_s^a} &\leq K \ln(K) + K \ln(t) & \text{and} & \quad \sum_{a \in \mathcal{K}} \sum_{s=n_0+1}^t \frac{w_s^a}{N_{s-1}^a} \leq K \ln(K) + 2K \ln(t) \\
 \sum_{a \in \mathcal{K}} \sum_{s=n_0}^t \frac{w_s^a}{\sqrt{N_s^a}} &\leq K \ln(K) + 2\sqrt{Kt} & \text{and} & \quad \sum_{a \in \mathcal{K}} \sum_{s=n_0+1}^t \frac{w_s^a}{N_{s-1}^a} \leq K \ln(K) + 2\sqrt{2Kt}
 \end{aligned}$$

Proof. Let's prove the first and the third inequality. Let $a \in \mathcal{K}$ and t_0^a be the first time such that: $\sum_{s=1}^{t_0^a-1} w_s^a > \ln(K) - 1$. Since $w_{t_0^a-1}^a \leq 1$, this yields $\sum_{s=1}^{t_0^a-1} w_s^a \leq \ln(K)$. Since $N_s^a \geq 1$ for $s \geq n_0$, we obtain:

$$\begin{aligned}
 \sum_{s=n_0}^t \frac{w_s^a}{N_s^a} &= \sum_{s=n_0}^{t_0^a-1} \frac{w_s^a}{N_s^a} + \sum_{s=t_0^a}^t \frac{w_s^a}{N_s^a} \leq \sum_{s=n_0}^{t_0^a-1} w_s^a + \sum_{s=t_0^a}^t \frac{w_s^a}{N_s^a} \leq \ln(K) + \sum_{s=t_0^a}^t \frac{w_s^a}{N_s^a} \\
 \sum_{s=n_0}^t \frac{w_s^a}{\sqrt{N_s^a}} &= \sum_{s=n_0}^{t_0^a-1} \frac{w_s^a}{\sqrt{N_s^a}} + \sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{N_s^a}} \leq \sum_{s=n_0}^{t_0^a-1} w_s^a + \sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{N_s^a}} \leq \ln(K) + \sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{N_s^a}}
 \end{aligned}$$

Combining the Lemma E.9 and Lemma E.11 for $x_s = w_s^a$, we obtain:

$$\begin{aligned} \sum_{s=t_0^a}^t \frac{w_s^a}{N_s^a} &\leq \sum_{s=t_0^a}^t \frac{w_s^a}{\sum_{r=1}^s w_r^a - \ln(K)} \leq \ln \left(\sum_{s=1}^t w_s^a - \ln(K) \right) - \ln \left(\sum_{s=1}^{t_0^a-1} w_s^a - \ln(K) \right) \\ \sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{N_s^a}} &\leq \sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{\sum_{r=1}^s w_r^a - \ln(K)}} \leq 2\sqrt{\sum_{s=1}^t w_s^a - \ln(K)} - 2\sqrt{\sum_{s=1}^{t_0^a-1} w_s^a - \ln(K)} \end{aligned}$$

Since $\sum_{s=1}^{t_0^a-1} w_s^a > \ln(K) - 1$, we have $\ln \left(\sum_{s=1}^{t_0^a-1} w_s^a - \ln(K) \right) \geq 0$. Therefore, we have shown: $\sum_{s=t_0^a}^t \frac{w_s^a}{N_s^a} \leq \ln \left(\sum_{s=1}^t w_s^a \right)$ and $\sum_{s=t_0^a}^t \frac{w_s^a}{\sqrt{N_s^a}} \leq 2\sqrt{\sum_{s=1}^t w_s^a}$. By concavity of $x \mapsto \ln(x)$ (resp. $x \mapsto \sqrt{x}$) and the fact that $\sum_{s=1}^t \sum_{a \in \mathcal{K}} w_s^a = t$, we obtain:

$$\sum_{a \in \mathcal{K}} \sum_{s=n_0}^t \frac{w_s^a}{N_s^a} \leq K \ln(K) + K \ln(t) \quad \text{and} \quad \sum_{a \in \mathcal{K}} \sum_{s=n_0}^t \frac{w_s^a}{\sqrt{N_s^a}} \leq K \ln(K) + 2\sqrt{Kt}$$

For all $s \geq n_0$, we have $N_{s-1}^a \geq 1$, hence $N_{s-1}^a \geq \frac{1}{2}N_{s,a}$. Plugging this inequality in the sum starting from t_0^a yields: $\sum_{a \in \mathcal{K}} \sum_{s=n_0+1}^t \frac{w_s^a}{N_{s-1}^a} \leq K \ln(K) + 2K \ln(t)$ and $\sum_{a \in \mathcal{K}} \sum_{s=n_0}^t \frac{w_s^a}{\sqrt{N_{s-1}^a}} \leq K \ln(K) + 2\sqrt{2Kt}$. \square

Lemma E.23 gives an upper bound on the norm of each arms once reweighted by the inverse of the empirical allocation.

Lemma E.23. For all $s \in \llbracket n_0 + 1, t \rrbracket$ and all $a \in \mathcal{K}$, we have $\|a\|_{V_{N_{s-1}}^{-1}}^2 \leq \frac{1}{N_{s-1}^a}$.

Proof. Recall that \mathcal{K} is arbitrary and by initialization $N_{s-1}^a \geq 1$ for all $(s, a) \in \llbracket n_0 + 1, t \rrbracket \times \mathcal{K}$. Let $a \in \mathcal{K}$. Let's rewrite $V_{N_{s-1}} = C_{s-1}^a + N_{s-1}^a a a^\top$ where $C_{s-1}^a = \sum_{b \in \mathcal{K}: b \neq a} N_{s-1}^b b b^\top$. If $\mathcal{K} = \{e_a\}_{a \in [d]}$, we have directly that $\|a\|_{V_N^{-1}}^2 = \frac{1}{N^a}$. In the following, we consider the more general case where $\mathcal{K} \neq \{e_a\}_{a \in [d]}$.

Case 1: C_{s-1}^a is invertible. Using the Sherman-Morrison formula, we obtain: for all $s \in \llbracket n_0 + 1, t \rrbracket$ and all $a \in \mathcal{K}$

$$\begin{aligned} V_{N_{s-1}}^{-1} &= (C_{s-1}^a)^{-1} - N_{s-1}^a \frac{(C_{s-1}^a)^{-1} a a^\top (C_{s-1}^a)^{-1}}{1 + N_{s-1}^a a^\top (C_{s-1}^a)^{-1} a} \\ \|a\|_{V_{N_{s-1}}^{-1}}^2 &= a^\top V_{N_{s-1}}^{-1} a = \frac{a^\top (C_{s-1}^a)^{-1} a}{1 + N_{s-1}^a a^\top (C_{s-1}^a)^{-1} a} \leq \frac{1}{N_{s-1}^a} \end{aligned}$$

where for the last inequality, we used that $\frac{x}{1+xy} \leq \frac{1}{y}$ (since it is equivalent with $1 + xy \geq xy$ which is true) for $x = a^\top (C_{s-1}^a)^{-1} a$ and $y = N_{s-1}^a$.

Case 2: C_{s-1}^a is not invertible. By initialization, we know that $V_{N_{s-1}}$ is invertible for all $s \in \llbracket n_0 + 1, t \rrbracket$. Let $u \in \text{Ker}(C_{s-1}^a) \setminus \{0_d\}$. We have $N_{s-1}^a \langle a, u \rangle a = V_{N_{s-1}} u \neq 0_d$ since $V_{N_{s-1}}$ is invertible and $u \neq 0_d$. Given that $N_{s-1}^a > 0$ (otherwise $V_{N_{s-1}} = C_{s-1}^a$, hence contradiction with invertible), we obtain that $\langle a, u \rangle > 0$, hence $u \notin \text{Span}(a)^\perp$. By dimension consideration, we obtain $\text{Span}(a) = \text{Ker}(C_{s-1}^a)$. This yields directly that a is an eigenvector of $V_{N_{s-1}}$ with eigenvalue $N_{s-1}^a \|a\|_2^2$. Therefore, a is also an eigenvector of $V_{N_{s-1}}^{-1}$ with eigenvalue $\frac{1}{N_{s-1}^a \|a\|_2^2}$ and we can conclude:

$$a^\top V_{N_{s-1}}^{-1} a = a^\top \left(\frac{1}{N_{s-1}^a \|a\|_2^2} a \right) = \frac{1}{N_{s-1}^a}. \quad \square$$

Lemma E.24 gives an upper bound on the cumulative sum of the reweighted slacks involved in the optimistic reward.

Lemma E.24.

$$\begin{aligned} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a c_{s-1}^a &\leq f(t^{1+b}) (K \ln(K) + 2K \ln(t)) \\ \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \sqrt{c_{s-1}^a} &\leq \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt}) \end{aligned}$$

Proof. Using the definition of c_s^a in (20), $f((s-1)^{1+b}) \leq f(t^{1+b})$ for all $s \leq t$, we obtain:

$$\begin{aligned} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a c_{s-1}^a &\leq \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a f((s-1)^{1+b}) \|a\|_{V_{N_{s-1}}^{-1}}^2 \leq f(t^{1+b}) \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \|a\|_{V_{N_{s-1}}^{-1}}^2 \\ &\leq f(t^{1+b}) \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} \frac{w_s^a}{N_{s-1}^a} \\ &\leq f(t^{1+b}) (K \ln(K) + 2K \ln(t)) \end{aligned}$$

where the second to last inequality is obtained by Lemma E.23 and the last one by Lemma E.22. Using the same arguments, we obtain:

$$\begin{aligned} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \sqrt{c_{s-1}^a} &\leq \sqrt{f(t^{1+b})} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a \|a\|_{V_{N_{s-1}}^{-1}} \leq \sqrt{f(t^{1+b})} \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} \frac{w_s^a}{\sqrt{N_{s-1}^a}} \\ &\leq \sqrt{f(t^{1+b})} (K \ln(K) + \sqrt{8Kt}) \end{aligned}$$

□

Lemma E.25 gives an upper bound on the cumulative sum of the reweighted KL divergence between the true parameter and its MLE.

Lemma E.25. *Let $b > 0$ as in (20). Under \mathcal{E}_t ,*

$$\sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2 \leq f(t^{1+b}) (K \ln(K) + 2K \ln(t))$$

Proof. Using Lemma E.16, we obtain:

$$\begin{aligned} \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \|\mu_{s-1} - \mu\|_{V_{w_s}}^2 &= \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \sum_{a \in \mathcal{K}} w_s^a \|\mu_{s-1} - \mu\|_{aa^\top}^2 \leq \sum_{s=\lceil t^{\frac{1}{1+b}} \rceil}^t \sum_{a \in \mathcal{K}} w_s^a c_{s-1}^a \leq \sum_{s=n_0+1}^t \sum_{a \in \mathcal{K}} w_s^a c_{s-1}^a \\ &\leq f(t^{1+b}) (K \ln(K) + 2K \ln(t)) \end{aligned}$$

where the second inequality uses $w_s^a c_{s-1}^a \geq 0$ and the last inequality is obtained by Lemma E.24. □

F. Implementation Details and Additional Experiments

After presenting the implementations details in Appendix F.1, we display supplementary experiments in Appendix F.2.

F.1. Implementation Details

Computational Cost of $L\varepsilon$ BAI When $\overline{\mathcal{M}} = \mathbb{R}^d$, the non-convex set of alternatives can be rewritten as a union over $Z - 1$ half-spaces, on which there is a closed-form formula for the closest alternative of Lemma C.1. Therefore, the stopping-recommendation pair has a computational cost in $\mathcal{O}(Zd^2|Z_\varepsilon(\mu_{t-1})|)$. For the heuristic $L\varepsilon$ BAI, i.e. $\tilde{z}_t = z_t$, the sampling rule has a computational cost in $\mathcal{O}((K + Z)d^2)$. Computing $z_F(\mu_{t-1})$ requires solving $|Z_\varepsilon(\mu_{t-1})|$ separate optimization problems. Each one can be rewritten as an easier one-dimensional optimization problem (Garivier & Kaufmann, 2016) which can be solved numerically by using two nested binary searches. Our experiments testify of the feasibility of that procedure.

Algorithm Implementations We list below more clarifications on the exact implementation of each individual algorithm.

- The discretization of the simplex Δ_2 and Δ_4 (with 500 and 10000 vectors) is obtained by drawing uniformly vectors in the simplex, in practice we used a Dirichlet distribution with parameters $\frac{1}{2}\mathbf{1}_2$ and $\frac{1}{4}\mathbf{1}_4$.
- The algorithms DKM, LinGame and $L\varepsilon$ BAI are implemented without the boundedness assumption. In practice, given $(\mu, z, w, x) \in \mathcal{M} \times \mathcal{Z}_\varepsilon(\mu) \times \Delta_K \times \mathcal{Z} \setminus \{z\}$, we used the closed-form formulas for the closest alternative $\lambda_0(\mu, z, w, x)$ and $\lambda_\varepsilon(\mu, z, w, x)$ as detailed in Appendix C.3.1.
- We consider the greedy version of LinGapE, which does not have a theoretical guarantee in the general case. We pull an arm with $z_t \in z^*(\mu_{t-1})$ and $x_t = \arg \max_{x \neq z_t} \langle \mu_{t-1}, x - z_t \rangle + \|x - z_t\|_{V_{N_{t-1}}^{-1}} \sqrt{2\beta(t-1, \delta)}$.
- Likewise, we use the greedy version of \mathcal{XY} -Static in order to avoid computing the optimal allocation at each step.
- We set the hyper-parameter of \mathcal{XY} -Adaptive to 0.1 as done in Soare et al. (2014); Degenne et al. (2020a). It controls the length of each phase.

Stopping Threshold Instead of the stopping threshold (6) supported by the theory, we use as heuristic $\beta(t, \delta) = 4 \ln \left(\frac{4 + \ln(t/2)}{\delta} \right)$, where we take the main term and plug in d instead of K . This is similar to what we could obtain with a threshold tailored to linear bandits and featuring only d . With this choice, the empirical error (number of runs such that $\hat{z} \notin \mathcal{Z}_\varepsilon(\mu)$) is lower than the confidence parameter δ . Previous algorithms on BAI settings (Garivier & Kaufmann, 2016; Degenne et al., 2019; 2020a; Jourdan et al., 2021) were using $\beta(t, \delta) = \ln \left(\frac{1 + \ln(t)}{\delta} \right)$ as heuristic for the stopping threshold. Interestingly, using this coarser stopping threshold for ε -BAI leads to an empirical error which is higher than δ , violating the (ε, δ) -PAC property. Therefore, we need to be closer to the theoretically validated threshold.

An additional argument in favor of considering the same stopping rule for all BAI algorithm can be seen in the choice of the stopping threshold itself. DKM (Degenne et al., 2019) uses the concentration results in Garivier & Kaufmann (2016), LinGame (Degenne et al., 2020a) the ones in (Lattimore & Szepesvári, 2020) and $L\varepsilon$ BAI the ones in (Kaufmann & Koolen, 2021). Therefore, a fair comparison of the sampling rule requires using the same $\beta(t, \delta)$, as defined above.

Reproducibility To assess our code and reproduce the experiments presented in this paper, you need to unzip the provided code by running `unzip code.zip`. All the algorithms and experiments are implemented in Julia 1.6.3 (but also run on Julia 1.1.1). Plots were generated with the `StatsPlots.jl` package. Other dependencies are listed in the `Readme.md`. The `Readme.md` file provides detailed julia instructions to reproduce our experiments, as well as a `script.sh` to run them all at once. The general structure of the code (and some functions) is taken from the `tidnabbil` library.²

F.2. Supplementary Experiments

Supplementary experiments for the multiplicative ε -optimality are shown in Appendix F.2.1. The equivalent experiments for the additive ε -optimality are displayed in Appendix F.2.2.

In the following, the hard/random instances that are considered are the same as in Section 6. We consider the same choice of parameters: $(\varepsilon, \delta) = (0.05, 0.01)$, discretization of the simplex with 500 (resp. 10000) vectors when $K = 2$ (resp.

²This library was created by (Degenne et al., 2019), see <https://bitbucket.org/wmkoolen/tidnabbil>. No license were available on the repository, but we obtained the authorization from the authors.

$K = 4$), average (resp. standard deviation) on 5000 runs (resp. sub-samples of size 100). Unless specified otherwise, the recommendation rule and the \mathcal{Z} -oracle return an instantaneous furthest answer, i.e. $z_t = \tilde{z}_t \in z_F(\mu_{t-1}, N_{t-1})$, and the stopping-recommendation pair is updated/evaluated at each time t .

F.2.1. MULTIPLICATIVE OPTIMALITY

Below, we present experiments on the multiplicative ε -optimality that were conducted to highlight algorithmic choices.




Figure 5. Empirical stopping time of the modified BAI algorithms with $z_t \in z_F(\mu_{t-1}, N_{t-1})$ on the hard instance (star equals mean). “-G” denotes when $z_t \in z^*(\mu_{t-1})$.

Modified BAI Figure 5 conveys the same message as Figure 2. For all modified BAI using (5), considering an instantaneous furthest answer instead of a greedy answer leads to lower empirical stopping time. Their ratio is 0.928 on average (coherent with Figure 1(b)). Note that the extension of LinGapE to the multiplicative ε -optimality is not obtained directly as was done for the additive setting, hence we didn’t displayed it.

Table 4. Empirical stopping time ($\pm \sigma$) with their original stopping rule or with ours (5) on the hard instance ($\mathcal{K} = \mathcal{Z}$). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

	LinGame	LinGapE	\mathcal{XY} -Adaptive
Original	151473 (± 18082)	121716 (± 14771)	263806 (± 36896)
Modified	423 (± 51)	414 (± 51)	399 (± 52)

Original Stopping Rule In Table 4, we displayed the difference between the empirical stopping time of BAI algorithms using their original stopping rule and ours (5). Their stopping time is 438 times higher when using their original stopping rule which was designed for an harder problem, i.e. identifying the unique best-answer which is also an ε -optimal answer.

Empirical Allocation Table 5 details the empirical allocations of pulls. Uniform and DKM sample all arms similarly. Therefore, they stop with almost twice as many samples as the other algorithms whose empirical allocations are close to the oracle allocation used in the fixed algorithm. As expected, \mathcal{XY} -Adaptive outperforms \mathcal{XY} -Static and G-Static. We see that \mathcal{LEBAI} slightly outperforms LinGapE and \mathcal{XY} -Adaptive, performs better than LinGame and \mathcal{XY} -Static and is on par with the “oracle” *fixed* algorithm.

Computational Relaxations In Appendix E.2.1, we introduced the sticky and the lazy schemes. Instead of updating the candidate answer and evaluating the stopping criterion at each time t , the lazy scheme only does it on all time $t \in \mathcal{T} \subseteq \mathbb{N}$. The sticky scheme evaluates the stopping criterion at each time $t \geq n_0 + 1$, but update the candidate answer when $t \in \mathcal{T} \subseteq \mathbb{N}$. With a good choice of \mathcal{T} , using the sticky/lazy schemes with \mathcal{LEBAI} yield an asymptotically optimal algorithm

Table 5. Average number of pulls per arm and empirical stopping time ($\pm \sigma$) on the hard instance ($\mathcal{K} = \mathcal{Z}$). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

	a_1	a_2	a_3	a_4	Total
$L\varepsilon$ BAI	115	247	19	3	384 (± 17)
LinGame	120	239	52	12	423 (± 16)
DKM	174	223	173	175	745 (± 29)
LinGapE	63	349	1	1	414 (± 16)
G-Static	204	227	7	19	458 (± 18)
\mathcal{XY} -Static	223	224	1	1	449 (± 17)
\mathcal{XY} -Adaptive	82	315	1	1	399 (± 18)
Fixed	96	271	1	1	370 (± 16)
Uniform	214	214	214	213	856 (± 33)




Figure 6. Empirical stopping time on the hard instance ($\mathcal{K} = \mathcal{Z}$) for (a) the lazy and sticky update, and different implementations of (b) the lazy scheme and (c) the sticky scheme. “-S” denotes the sticky scheme and “-L” the lazy one. The notations for implementations are: “-C” for the constant one with $T_0 = 10$, “-G” for the geometric one with $(T_0, \gamma) = (10, 0.2)$, “-D” for geometrically decreasing one with $(T_0, \gamma) = (10, 0.2)$ and “-B” for the Bernoulli one with parameter $p = 0.1$.

(Appendix E.2.1). To ensure that the strategy is (ε, δ) -PAC, it is sufficient to consider $|\mathcal{T}| = +\infty$ (Lemma 3.1), property satisfied by all the implementations below.

In Figure 6, we test several implementations of lazy/sticky schemes to assess their impact on the empirical stopping time. They differ by how the infinite grid of time \mathcal{T} is defined. The constant implementation uses an arithmetic grid of time with parameter $T_0 > n_0$, i.e. $\mathcal{T} = \{n_0 + 1\} \cup \{iT_0\}_{i \in \mathbb{N}^*}$. The geometric implementation uses a geometric grid of time with parameter $T_0 > n_0$ and $\gamma > 0$, i.e. $\mathcal{T} \stackrel{\text{def}}{=} \{n_0 + 1\} \cup \{T_i\}_{i \in \mathbb{N}^*}$ where $T_i = \lceil (1 + \gamma)T_{i-1} \rceil$ for $i \in \mathbb{N}^*$. The geometrically decreasing implementation uses a grid of time with parameter $T_0 > n_0$ and $\gamma > 0$, i.e. $\mathcal{T} \stackrel{\text{def}}{=} \{n_0 + 1\} \cup \{T_i\}_{i \in \mathbb{N}^*}$ where $T_i = \lceil (1 + \frac{\gamma}{i})T_{i-1} \rceil$ for $i \in \mathbb{N}^*$. The Bernoulli implementation with parameter p is slightly different as it adds internal randomness, hence it is not a deterministic strategy anymore. The idea is simple: draw $X \sim \mathcal{B}(p)$, if $X = 1$ we update the candidate answer, else we stick to it. The geometrically decreasing implementation is the only implementation ensuring that $L\varepsilon$ BAI is an asymptotically optimal algorithm (Appendix E.2.1).

In Figure 6(a), we see that the sticky relaxation allows to perform on par with the algorithm updating the recommendation rule at each time t . However, when considering the lazy relaxation, we pay the price of not evaluating the stopping rule at each time t by incurring a slightly higher empirical stopping time. Therefore, depending on the constraints of the practitioner, the sticky/lazy schemes allow to reduce the computational cost per time step while keeping similar sample complexity.

In Figure 6(b), we observe that the empirical stopping time might be higher depending on the implementation of the lazy scheme. Overall, algorithms using a lazy scheme suffer from slightly worse sample complexity. This is the price to pay to drastically reduce the computational cost.

In Figure 6(c), we see that the exact implementation of the sticky scheme has few consequences as regards the empirical stopping time since they all perform on par with the algorithm updating the recommendation rule at each time t . Therefore, the sticky scheme is the computational relaxation to adopt when one wishes to reduce the computational cost in a significant

manner without damaging the empirical performance.




Figure 7. Empirical stopping time on the hard instance ($\mathcal{K} = \mathcal{Z}$). “-D” denotes when the D-Tracking is used instead of C-Tracking and “-NF” denotes the removal of forced exploration.

Tracking And Forced Exploration The tracking procedure that we use throughout the paper is referred as C-Tracking in the literature, since it tracks the cumulative sum of pulling proportions over arms played by the agent, i.e. $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - W_t^a$ where $W_t = \sum_{s=n_0+1}^t w_s$. Another commonly considered tracking procedure is D-Tracking (Garivier & Kaufmann, 2016), in which the pulling proportion over arms played by the agent at time t is directly tracked, i.e. $a_t \in \arg \min_{a \in \mathcal{K}} N_{t-1}^a - (t - n_0)w_t^a$. In Degenne & Koolen (2019), it is shown that D-Tracking might fail when several correct answer exist. This explains why we only considered C-Tracking in this paper. In Figure 7, we observe that considering D-Tracking instead of C-Tracking leads to similar empirical stopping time.

As discussed in Appendix E.3.1, we need a logarithmic forced exploration to conclude the proof, which was not needed by previous game-based approaches Degenne et al. (2020a). In Figure 7, we see that removing the forced exploration has almost no impact on the empirical stopping time.

F.2.2. ADDITIVE OPTIMALITY

When comparing the following plots and their equivalent in the main content of the paper, we observe that the characteristic time and the empirical stopping time in the additive ε -optimality is always smaller than for the multiplicative one (all other parameters being identical). If there isn’t a natural choice of ε -optimality, this fact might influence the choice of the practitioner. Overall the messages conveyed by the following plots are the same as the ones highlighted in the multiplicative ε -optimality case, we report them for the sake of completeness.

Furthest Answer The results observed in Figure 8 are similar to the ones in Figure 1, while the discrepancy is less noteworthy. The proportion of draws where $z_F(\mu) \notin z^*(\mu)$ is not negligible, reaching on average 10%. The impact on the characteristic time is visible with a ratio on average at 0.95, but it is very slim when considering $\varepsilon \leq 0.01$. Up to $\varepsilon \approx 0.25$ the proportion of draws where $z_F(\mu) \notin z^*(\mu)$ has a logarithmic increase with ε . Therefore, when they are different, the furthest answer still outperforms a greedy answer achieving $T_{g,\varepsilon}^{\text{add}}(\mu)$. We note that this difference is smaller in the additive ε -optimality than for the multiplicative ε -optimality.

Table 6. Empirical stopping time ($\pm \sigma$) with their original stopping rule or with ours (5) on the hard instance ($\mathcal{K} = \mathcal{Z}$). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

	LinGame	LinGapE	\mathcal{XY} -Adaptive
Original	151473 (± 18082)	121716 (± 14771)	263806 (± 36896)
Modified	393 (± 49)	297 (± 39)	314 (± 38)




Figure 8. Influence of ε on (a) the proportion of draws where $z_F(\mu) \notin z^*(\mu)$, (b) the median (and first/third quartile) of $\frac{T_{g,\varepsilon}^{\text{add}}(\mu)}{T_g^{\text{add}}(\mu)}$, when $z_F(\mu) \notin z^*(\mu)$.

Original Stopping Rule In Table 6, the *Original* line is unchanged compared to Table 4 since BAI algorithms consider $\varepsilon = 0$, hence no multiplicative/additive notions of ε -optimality. Similarly, their stopping time is 545 times higher when using their original stopping rule which was designed for an harder problem, i.e. identifying the unique best-answer which is also an ε -optimal answer.

Table 7. Empirical stopping time ($\pm \sigma$) for different combinations of sampling rule and recommendation rule on the hard instance with $\mathcal{K} = \{e_1, e_2\}$.

	$z^*(\mu_{t-1})$	$z_F(\mu_{t-1})$	$z_F(\mu_{t-1}, N_{t-1})$
$L\varepsilon$ BAI	350 (± 14)	318 (± 12)	318 (± 13)
ε -TaS	323 (± 12)	298 (± 13)	298 (± 13)
Fixed	324 (± 10)	300 (± 11)	300 (± 11)
Uniform	477 (± 15)	434 (± 16)	434 (± 16)

Choosing Answers Table 7 conveys the same messages as Table 1. It shows that using greedy is consistently worse than an (instantaneous) furthest answer, the ratio of their stopping time is 0.917 on average (coherent with Figure 8(b)) and that instantaneous furthest answer and furthest answer have the same empirical performance. While $L\varepsilon$ BAI consistently outperform uniform sampling (73%), it performs slightly worse than ε -TaS and the “oracle” *fixed* algorithm (tracking the optimal allocation $w_F(\mu)$).

Modified BAI Algorithms and Empirical Allocation The detailed empirical allocations presented in Table 8 bare similarity with the results shown in Table 5. Likewise, Figure 9 shows similar results as Figure 3. Uniform and DKM sample all arms equally. Their stopping time is twice as high as the other algorithms whose empirical allocations are close to the oracle allocation (fixed algorithm). $\mathcal{X}\mathcal{Y}$ -Adaptive outperforms $\mathcal{X}\mathcal{Y}$ -Static and G-Static. We see that $L\varepsilon$ BAI slightly outperforms LinGapE and $\mathcal{X}\mathcal{Y}$ -Adaptive, performs better than LinGame and $\mathcal{X}\mathcal{Y}$ -Static and is on par with the “oracle” *fixed* algorithm.

Random Instances In Higher Dimensions For random instances and increasing dimension, Figure 10 reveals the same trends as Figure 4. $L\varepsilon$ BAI shows competitive empirical performance with modified BAI algorithms. Even though it is outperformed by LinGapE, $L\varepsilon$ BAI is almost twice as fast as $\mathcal{X}\mathcal{Y}$ -Adaptive and appears to be slightly more robust than LinGame to increasing dimension.

Table 8. Average number of pulls per arm and empirical stopping time ($\pm \sigma$) on the hard instance ($\mathcal{K} = \mathcal{Z}$). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.

	a_1	a_2	a_3	a_4	Total
$L\varepsilon$ BAI	77	228	13	3	321 (± 13)
LinGame	112	221	49	11	393 (± 15)
DKM	169	217	168	169	723 (± 27)
LinGapE	47	248	1	1	297 (± 13)
G-Static	196	217	7	18	438 (± 18)
\mathcal{XY} -Static	215	217	1	1	434 (± 17)
\mathcal{XY} -Adaptive	76	236	1	1	314 (± 12)
Fixed	48	251	1	1	300 (± 11)
Uniform	211	211	211	210	844 (± 33)




Figure 9. Empirical stopping time on the hard instance ($\mathcal{K} = \mathcal{Z}$), the star represents the mean. The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.




Figure 10. Empirical stopping time on random instances ($\mathcal{K} = \mathcal{Z}$) for $d \in \{6, 8, 10, 12\}$ (from top left to bottom right). The modified BAI algorithms use (5) with $z_t \in z_F(\mu_{t-1}, N_{t-1})$.