
Pure Exploration with Multiple Correct Answers

Rémy Degenne

Centrum Wiskunde & Informatica
Science Park 123, Amsterdam, NL
remy.degenne@cwi.nl

Wouter M. Koolen

Centrum Wiskunde & Informatica
Science Park 123, Amsterdam, NL
wmkoolen@cwi.nl

Abstract

We determine the sample complexity of pure exploration bandit problems with multiple good answers. We derive a lower bound using a new game equilibrium argument. We show how continuity and convexity properties of *single-answer* problems ensure that the existing Track-and-Stop algorithm has asymptotically optimal sample complexity. However, that convexity is lost when going to the *multiple-answer* setting. We present a new algorithm which extends Track-and-Stop to the multiple-answer case and has asymptotic sample complexity matching the lower bound.

1 Introduction

In *pure exploration* aka *active testing* problems the learning system interacts with its environment by sequentially performing experiments to quickly and reliably identify the answer to a particular pre-specified question. Practical applications range from simple queries for cost-constrained physical regimes, i.e. clinical drug testing, to complex queries in structured environments bottlenecked by computation, i.e. simulation-based planning. The theory of pure exploration is studied in the multi-armed bandit framework. The scientific challenge is to develop tools for characterising the sample complexity of new pure exploration problems, and methodologies for developing (matching) algorithms. With the aim of understanding power and limits of existing methodology, we study an extended problem formulation where each instance may have multiple correct answers. We find that multiple-answer problems present a phase transition in complexity, and require a change in our thinking about algorithms.

The existing methodology for developing asymptotically instance-optimal algorithms, Track-and-Stop by Garivier and Kaufmann [2016], exploits the so-called *oracle weights*. These probability distributions on arms naturally arise in sample complexity lower bounds, and dictate the optimal sampling proportions for an “oracle” algorithm that needs to be sample efficient only on exactly the current problem instance. The main idea is to track the oracle weights computed at a converging estimate of the instance. The analysis of Track-and-Stop requires continuity of the oracle weights as a function of the bandit model. For the core Best Arm Identification problem, discontinuity only occurs at degenerate instances where the sample complexity explodes. So this assumption seemed harmless.

Our contribution We show that the oracle weights may be non-unique, even for single-answer problems, and hence need to be regarded as a set-valued mapping. We show this mapping is always (upper hemi-)continuous. We show that each instance maps to a convex set for single-answer problems, and this allows us to extend the Track-and-Stop methodology to all such problems. At instances with non-singleton set-valued oracle weights more care is needed: of the two classical tracking schemes “C” converges to the convex set, while “D” may fail entirely.

We show that for multiple-answer problems convexity is violated. There are instances where two distinct oracle weights are optimal, while no mixture is. Unmodified tracking converges in law

(experimentally) to a distribution on the full convex hull, and suffers as a result. We propose a “sticky” modification to stabilise the approach, and show that now it converges to only the corners. We prove that Sticky Track-and-Stop is asymptotically optimal.

Related work Multi-armed bandits have been the subject of intense study in their role as a model for medical testing and reinforcement learning. For the objective of reward maximisation [Berry and Fristedt, 1985, Lai and Robbins, 1985, Auer et al., 2002, Bubeck and Cesa-Bianchi, 2012] the main challenge is balancing exploration and exploitation. The field of pure exploration (active testing) focuses on generalisation vs sample complexity, in the fixed confidence, fixed budget and simple regret scalarisations. It took off in machine learning with the *multiple-answer* problem of (ϵ, δ) -Best Arm Identification (BAI) [Even-Dar et al., 2002]. Early results focused on worst-case sample complexity guarantees in sub-Gaussian bandits. Successful approaches include *Upper and Lower confidence bounds* [Bubeck et al., 2011, Kalyanakrishnan et al., 2012, Gabillon et al., 2012, Kaufmann and Kalyanakrishnan, 2013, Jamieson et al., 2014], *Racing or Successive Rejects/Eliminations* [Maron and Moore, 1997, Even-Dar et al., 2006, Audibert et al., 2010, Kaufmann and Kalyanakrishnan, 2013, Karnin et al., 2013].

Fundamental information-theoretic barriers [Castro, 2014, Kaufmann et al., 2016, Garivier and Kaufmann, 2016] for each specific problem instance refined the worst-case picture, and sparked the development of instance-optimal algorithms for single-answer problems based on *Track-and-Stop* [Garivier and Kaufmann, 2016] and *Thompson Sampling* [Russo, 2016]. For multiple-answer problems the elegant KL-contraction-based lower bound is not sharp, and new techniques were developed by Garivier and Kaufmann [2019].

Recent years also saw a surge of interest in pure exploration with *complex queries* and *structured environments*. Kalyanakrishnan and Stone [2010] identify the top- M set, Locatelli et al. [2016] the arm closest to a threshold, and Chen et al. [2014], Gabillon et al. [2016] the optimiser over an arbitrary combinatorial class. For arms arranged in a matrix Katariya et al. [2017] study BAI under a rank-one assumption, while Zhou et al. [2017] seek to identify a Nash equilibrium. For arms arranged in a minimax tree there is significant interest in finding the optimal move at the root Teraoka et al. [2014], Garivier et al. [2016], Huang et al. [2017], Kaufmann and Koolen [2017], Kaufmann et al. [2018], as a theoretical model for studying Monte Carlo Tree search (which is a planning sub-module of many advanced reinforcement learning systems).

2 Notations

We work in a given one-parameter one-dimensional canonical exponential family, with mean parameter in an open interval $\mathcal{O} \subseteq \mathbb{R}$. We denote by $d(\mu, \lambda)$ the KL divergence from the distribution with mean μ to that with mean λ . A K -armed bandit model is identified by its vector $\boldsymbol{\mu} \in \mathcal{O}^K$ of mean parameters. We denote by $\mathcal{M} \subseteq \mathcal{O}^K$ the set of possible mean parameters in a specific bandit problem. Excluding parts of \mathcal{O}^K from \mathcal{M} may be used to encode a known structure of the problem. We assume that there is a finite domain \mathcal{I} of answers, and that the *correct answer* for each bandit model is specified by a set-valued function $i^* : \mathcal{M} \rightarrow 2^{\mathcal{I}}$.

Example 1. As our running example, we will use the *Any Low Arm* multiple-answer problem. Given a threshold $\gamma \in \mathcal{O}$, the goal is return the index k of any arm with $\mu_k < \gamma$ if such an arm exists, or to return “no” otherwise. Formally, we have possible answers $\mathcal{I} = [K] \cup \{\text{no}\}$, and correct answers

$$i^*(\boldsymbol{\mu}) = \begin{cases} \{k \mid \mu_k \leq \gamma\} & \text{if } \min_k \mu_k < \gamma, \\ \{\text{no}\} & \text{if } \min_k \mu_k > \gamma. \end{cases}$$

We exclude the case $\min_k \mu_k = \gamma$ from \mathcal{M} (as such instances require infinite sample complexity).

Additional examples of multiple-answer identification problems are visualised in Table 1 in Appendix B.

Algorithms. A learning strategy is parametrised by a stopping rule $\tau_\delta \in \mathbb{N}$ depending on a parameter $\delta \in [0, 1]$, a sampling rule $A_1, A_2, \dots \in [K]$, and a recommendation rule $\hat{i} \in \mathcal{I}$. When a learning strategy meets a bandit model $\boldsymbol{\mu}$, they interactively generate a history $A_1, X_1, \dots, A_\tau, X_\tau, \hat{i}$, where $X_t \sim \mu_{A_t}$. We allow the possibility of non-termination $\tau_\delta = \infty$, in which case there is no recommendation \hat{i} . At stage $t \in \mathbb{N}$, we denote by $N_t = (N_{t,1}, \dots, N_{t,K})$ the number of samples (or “pulls”) of the arms, and by $\hat{\boldsymbol{\mu}}_t$ the vector of empirical means of the samples of each arm.

Confidence and sample complexity. For confidence parameter $\delta \in (0, 1)$, we say that a strategy is δ -correct (or δ -PAC) for bandit model μ if it recommends a correct answer with high probability, i.e. $\mathbb{P}_\mu(\tau_\delta < \infty \text{ and } \hat{i} \in i^*(\mu)) \geq 1 - \delta$. We call a given strategy δ -correct if it is δ -correct for every $\mu \in \mathcal{M}$. We measure the statistical efficiency of a strategy on a bandit model μ by its *sample complexity* $\mathbb{E}_\mu[\tau_\delta]$. We are interested in δ -correct algorithms minimizing sample complexity.

Divergences. For any answer $i \in \mathcal{I}$, we define the *alternative to i* , denoted $\neg i$, to be the set of bandit models on which answer i is incorrect, i.e.

$$\neg i := \{\mu \in \mathcal{M} | i \notin i^*(\mu)\}.$$

We define the (w -weighted) divergence from $\mu \in \mathcal{M}$ to $\lambda \in \mathcal{M}$ or $\Lambda \subseteq \mathcal{M}$ by

$$\begin{aligned} D(w, \mu, \lambda) &= \sum_k w_k d(\mu_k, \lambda_k) & D(w, \mu, \Lambda) &= \inf_{\lambda \in \Lambda} D(w, \mu, \lambda) \\ D(\mu, \Lambda) &= \sup_{w \in \Delta_K} D(w, \mu, \Lambda) & D(\mu) &= \max_{i \in \mathcal{I}} D(\mu, \neg i) \end{aligned}$$

Note that $D(w, \mu, \Lambda) = 0$ whenever $\mu \in \Lambda$. We denote by $i_F(\mu)$ the argmax (set of maximisers) of $i \mapsto D(\mu, \neg i)$, and we call $i_F(\mu)$ the *oracle answer(s)* at μ . We write $w^*(\mu, \neg i)$ for the maximisers of $w \mapsto D(w, \mu, \neg i)$, and call these the *oracle weights for answer i* at μ . We write $w^*(\mu) = \bigcup_{i \in i_F(\mu)} w^*(\mu, \neg i)$ for the set of *oracle weights* among all oracle answers. We include expressions for the divergence when i^* is generated by half-spaces, minima and spheres in Appendix H.

Example 1 (Continued). Consider an *Any Low Arm* instance μ with $\min_k \mu_k < \gamma$. For any arm $i \in i^*(\mu)$, we have $D(w, \mu, \neg i) = w_i d(\mu_i, \gamma)$ and $D(\mu, \neg i) = d(\mu_i, \gamma)$. Hence $D(\mu) = d(\min_i \mu_i, \gamma)$, and $i_F(\mu) = \operatorname{argmin}_i \mu_i$. On the other hand, when $\min_k \mu_k > \gamma$, we have $i^*(\mu) = \{\text{no}\}$, so that $D(w, \mu, \text{no}) = \min_k w_k d(\mu_k, \gamma)$ and $D(\mu, \text{no}) = D(\mu) = 1 / \sum_{k=1}^K \frac{1}{d(\mu_k, \gamma)}$.

The function $i_F(\mu) = \{i \in \mathcal{I} : D(\mu, \neg i) = D(\mu)\}$ is set valued, as is w^* . They are singletons with continuous value over some connected subsets of \mathcal{M} , and are multi-valued on common boundaries of two or more such sets. Both i_F and w^* can be thought of as having single values, unless μ sits on such a boundary, in which case we will prove that they are equal to the union (or convex hull of the union) of their values in the neighbouring regions.

3 Lower Bound

We show a lower bound for any algorithm for multiple-answer problems. Our lower bound extends the single-answer result of Garivier and Kaufmann [2016]. We are further inspired by Garivier and Kaufmann [2019], who analyse the ϵ -BAI problem. They prove lower bounds for algorithms with a sampling rule independent of δ , imposing the further restriction that either $K = 2$ or that the algorithm ensures that $N_{t,k}/t$ converges as $t \rightarrow \infty$. The new ingredient in this section is a game-theoretic equilibrium argument, which allows us to analyse any δ -correct algorithm in any multiple answer problem. Our main lower bound is the following.

Theorem 1. *Any δ -correct algorithm verifies*

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\log(1/\delta)} \geq T^*(\mu) := D(\mu)^{-1} \quad \text{where} \quad D(\mu) = \max_{i \in i^*(\mu)} \max_{w \in \Delta_K} \inf_{\lambda \in \neg i} \sum_{k=1}^K w_k d(\mu_k, \lambda_k)$$

for any multiple answer instance μ with sub-Gaussian arm distributions.

The proof is in Appendix C, where we also discuss how the convenient sub-Gaussian assumption can be relaxed. We would like to point out one salient feature here. To show sample complexity lower bounds at μ , one needs to find problems that are hard to distinguish from it statistically, yet require a different answer. We obtain these problems by means of a minimax result.

Lemma 2. *For any answer $i \in \mathcal{I}$, the divergence from μ to $\neg i$ equals*

$$D(\mu, \neg i) = \inf_{\mathbb{P}} \max_{k \in [K]} \mathbb{E}_{\lambda \sim \mathbb{P}} [d(\mu_k, \lambda_k)].$$

where the infimum ranges over probability distributions on $\neg i$ supported on (at most) K points.

The proof of Theorem 1 then challenges any algorithm for μ by obtaining a witness \mathbb{P} for $D(\mu) = \max_i D(\mu, \neg i)$ from Lemma 2, sampling a model $\lambda \sim \mathbb{P}$, and showing that if the algorithm stops early, it must make a mistake w.h.p. on at least one model from the support. The equilibrium property of \mathbb{P} allows us to control a certain likelihood ratio martingale regardless of the sampling strategy employed by the algorithm.

We discuss the novel aspect of Theorem 1 and its lessons for the design of optimal algorithms. First of all, for single-answer instances $|i^*(\mu)|=1$ we recover the known asymptotic lower bound [Garivier and Kaufmann, 2016, Remark 2]. For multiple-answer instances the bound says the following. The optimal sample complexity hinges on the *oracle answers* $i_F(\mu)$. That is, for $i_f \in i_F(\mu)$, the complexity of problem μ is governed by the difficulty of discriminating μ from the set of models on which answer i_f is incorrect.

Is the bound tight? We argue yes. Consider the following oracle strategy, which is specifically designed to be very good at μ . First, it computes a pair (i, w) witnessing the two outer maxima in $D(\mu)$. The algorithm samples according to w . It stops when it can statistically discriminate $\hat{\mu}_t$ from $\neg i$, and outputs $\hat{i} = i$. This algorithm will indeed have expected sample complexity equal to $D(\mu)^{-1}$ at μ , and it will be correct.

The above oracle viewpoint presents an idea for designing algorithms, following Garivier and Kaufmann [2016] and Chen et al. [2017]. Perform a lower-order amount of forced exploration of all arms to ensure $\hat{\mu}_t \rightarrow \mu$. Then at each time point compute the empirical mean vector $\hat{\mu}_t$ and oracle weights $w_t \in w^*(\hat{\mu}_t)$. Then sample according to w_t . This approach is successful for single-answer bandits with unique and continuous oracle weights. We argue in Section 4.3 below that it extends to points of discontinuity by exploiting upper hemicontinuity and convexity of w^* .

For multiple-answer bandits, we argue that the set of maximisers $w^*(\mu)$ is no longer convex when $i_F(\mu)$ is not a singleton. It can then happen that $\hat{\mu}_t \rightarrow \mu$, while at the same time $w^*(\hat{\mu}_t)$ keeps oscillating. If the algorithm tracks $w^*(\hat{\mu}_t)$, its sampling proportions will end up in the convex hull of $w^*(\mu)$. However, as $w^*(\mu)$ is not convex itself, these proportions will not be optimal. We present empirical evidence for that effect in Appendix D. The lesson here is that the oracle needs to pick an answer and “stick with it”. This will be the basis of our algorithm design in Section 5.

4 Properties of the Optimal Allocation Sets

The Track-and-Stop sampling strategy aims at ensuring that the sampling proportions converge to oracle weights. In the case of a singleton-valued oracle weights set $w^*(\mu)$ for single answer problems, that convergence was proven in [Garivier and Kaufmann, 2016]. We study properties of that set with the double aim of extending Track-and-Stop to points μ where $w^*(\mu)$ is not a singleton and of highlighting what properties hold only for the single-answer case, but not in general.

4.1 Continuity

We first prove continuity properties of D and w^* . We show how the convergence of $\hat{\mu}_t$ to μ translates into properties of the divergences from $\hat{\mu}_t$ to the alternative sets.

For a set B , let $\mathbb{S}(B) = 2^B \setminus \{\emptyset\}$ be the set of all *non-empty* subsets of B .

Definition 3 (Upper hemicontinuity). A set-valued function $\Gamma : A \rightarrow \mathbb{S}(B)$ is upper hemicontinuous at $a \in A$ if for any open neighbourhood V of $\Gamma(a)$ there exists a neighbourhood U of a such that for all $x \in U$, $\Gamma(x)$ is a subset of V .

Theorem 4. For all $i \in \mathcal{I}$,

1. the function $(w, \mu) \mapsto D(w, \mu, \neg i)$ is continuous on $\Delta_K \times \mathcal{M}$,
2. $\mu \mapsto D(\mu, \neg i)$ and $\mu \mapsto D(\mu)$ are continuous on \mathcal{M} ,
3. $\mu \mapsto w^*(\mu, \neg i)$, $\mu \mapsto w^*(\mu)$ and $\mu \mapsto i_F(\mu)$ are upper hemicontinuous on \mathcal{M} with non-empty and compact values,

The proof is in Appendix F. It uses Berge’s maximum theorem and a modification thereof due to [Feinberg et al., 2014]. Related continuity results using this type of arguments, but restricted to single-valued functions, appeared for the regret minimization problem in [Combes et al., 2017].

4.2 Convexity

Next we establish convexity.

Proposition 5. *For each $i \in \mathcal{I}$, for all $\boldsymbol{\mu} \in \mathcal{M}$ the set $\boldsymbol{w}^*(\boldsymbol{\mu}, \neg i)$ is convex.*

This is a consequence of the concavity of $\boldsymbol{w} \mapsto D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ (which is an infimum of linear functions). In single-answer problems, we obtain that the oracle weights set $\boldsymbol{w}^*(\boldsymbol{\mu})$ is convex everywhere. This is however not the case in general for multiple-answer problems, as illustrated by the next example.

Example 1 (Continued). Consider a $K = 2$ -arm *Any Low Arm* instance $\boldsymbol{\mu}$ with $\mu_1 < \gamma$ and $\mu_2 < \gamma$, so that both answers 1 and 2 are correct. Recall that $D(\boldsymbol{\mu}) = \max_{k=1,2} d(\mu_k, \gamma)$. Now for $\mu_1 < \mu_2 < \gamma$, $\boldsymbol{w}^*(\boldsymbol{\mu}) = \{(1, 0)\}$ and symmetrically for $\mu_2 < \mu_1 < \gamma$, $\boldsymbol{w}^*(\boldsymbol{\mu}) = \{(0, 1)\}$. However, for $\mu_1 = \mu_2 < \gamma$, $\boldsymbol{w}^*(\boldsymbol{\mu}) = \{(1, 0), (0, 1)\}$, which is not convex. Playing intermediate weights $\boldsymbol{w} = (\alpha, 1 - \alpha)$ results in strictly sub-optimal $D(\boldsymbol{\mu}, \boldsymbol{w}) = \max\{\alpha, 1 - \alpha\} d(\boldsymbol{\mu}, \gamma) < d(\boldsymbol{\mu}, \gamma) = D(\boldsymbol{\mu})$.

This example also illustrates the upper hemicontinuity of $\boldsymbol{w}^*(\boldsymbol{\mu})$: since $\boldsymbol{\mu}$ of the form (μ, μ) is the limit of a sequence $(\boldsymbol{\mu}_t)_{t \in \mathbb{N}}$ with $\mu_{t,1} < \mu_{t,2}$, we obtain that $\{(1, 0)\} \subseteq \boldsymbol{w}^*(\boldsymbol{\mu})$. Similarly, using a sequence with $\mu_{t,1} > \mu_{t,2}$, $\{(0, 1)\} \subseteq \boldsymbol{w}^*(\boldsymbol{\mu})$.

The example scales up to K arms, and shows that the sample complexity guarantee for vanilla TaS (Theorem 9) may exceed by a factor K the optimal complexity, which is matched by our new method (Theorem 11).

4.3 Consequences for Track-and-Stop

The original analysis of Track-and-Stop excludes the mean vectors $\boldsymbol{\mu} \in \mathcal{M}$ for which $\boldsymbol{w}^*(\boldsymbol{\mu})$ is not a singleton. We show that the upper hemicontinuity and convexity properties of $\boldsymbol{w}^*(\boldsymbol{\mu})$ allow us to extend that analysis to all $\boldsymbol{\mu}$ with a single oracle answer (in particular all single-answer bandit problems), at least for one of the two Track-and-Stop variants. Indeed, that algorithm was introduced with two possible subroutines, dubbed C-tracking and D-tracking [Garivier and Kaufmann, 2016]. Both variants compute oracle weights \boldsymbol{w}_t at the point $\hat{\boldsymbol{\mu}}_t$, but the arm pulled differs.

C-tracking: compute the projection $\boldsymbol{w}_t^{\varepsilon_t}$ of \boldsymbol{w}_t on $\Delta_K^{\varepsilon_t} = \{\boldsymbol{w} \in \Delta_K : \forall k \in [K], w_k \geq \varepsilon_t\}$, where $\varepsilon_t > 0$. Pull the arm with index $k_t = \arg \min_{k \in [K]} N_{t,k} - \sum_{s=1}^t w_{s,k}^{\varepsilon_s}$.

D-tracking: if there is an arm j with $N_{t,j} \leq \sqrt{t} - K/2$, then pull $k_t = j$. Otherwise, pull the arm $k_t = \arg \min_{k \in [K]} N_{t,k} - t w_{t,k}$.

The proof of the optimal sample complexity of Track-and-Stop for C-tracking remains essentially unchanged but we replace Proposition 9 of [Garivier and Kaufmann, 2016] by the following lemma, proved in Appendix G.3.

Lemma 6. *Let a sequence $(\hat{\boldsymbol{\mu}}_t)_{t \in \mathbb{N}}$ verify $\lim_{t \rightarrow +\infty} \hat{\boldsymbol{\mu}}_t = \boldsymbol{\mu}$. For all $t \geq 0$, let $\boldsymbol{w}_t \in \boldsymbol{w}^*(\hat{\boldsymbol{\mu}}_t)$ be arbitrary oracle weights for $\hat{\boldsymbol{\mu}}_t$. If $\boldsymbol{w}^*(\boldsymbol{\mu})$ is convex, then*

$$\lim_{t \rightarrow +\infty} \inf_{\boldsymbol{w} \in \boldsymbol{w}^*(\boldsymbol{\mu})} \left\| \frac{1}{t} \sum_{s=1}^t \boldsymbol{w}_s - \boldsymbol{w} \right\|_{\infty} = 0.$$

The average of oracle weights for $\hat{\boldsymbol{\mu}}_t$ converges to the set of oracle weights for $\boldsymbol{\mu}$. C-tracking then ensures that the proportion of pulls N_t/t is close to that average by Lemma 7 of [Garivier and Kaufmann, 2016], hence N_t/t gets close to oracle weights.

Theorem 7. *For all $\boldsymbol{\mu} \in \mathcal{M}$ such that $i_F(\boldsymbol{\mu})$ is a singleton (in particular all single-answer problems), Track-and-Stop with C-tracking is δ -correct with asymptotically optimal sample complexity.*

Proof in Appendix G.6. We encourage the reader to first proceed to Section 5, since the proof considers the result as a special case of the multiple-answers setting.

Remark 8. *If $\boldsymbol{w}^*(\boldsymbol{\mu})$ is not a singleton, Track-and-Stop using D-tracking may fail to converge to $\boldsymbol{w}^*(\boldsymbol{\mu})$, even when it is convex.*

While we do not prove that D-tracking fails to converge to $\boldsymbol{w}^*(\boldsymbol{\mu})$ on a specific example of a bandit, we provide empirical evidence in Appendix E. The reason for the failure of D-tracking

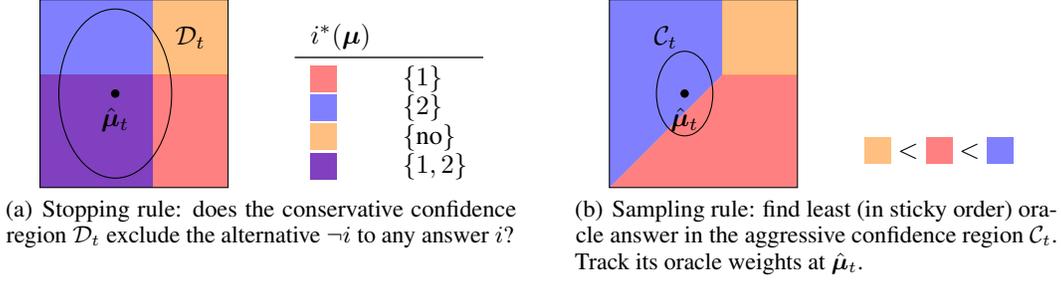


Figure 1: Sticky Track-and-Stop: The two main ideas, illustrated on the *Any Low Arm* problem.

is that it does not in general converge to the convex hull of the points it tracks. Suppose that $w_t = w^{(1)} = (1/2, 1/2, 0)$ for t odd and $w_t = w^{(2)} = (1/2, 0, 1/2)$ for t even. Then D-tracking verifies $\lim_{t \rightarrow +\infty} N_t/t = (1/3, 1/3, 1/3)$. This limit is outside of the convex hull of $\{w^{(1)}, w^{(2)}\}$.

5 Algorithms for the Multiple-Answers Setting

We can prove for Track-and-Stop the following suboptimal upper bound on the sample complexity, based on the fact that it ensures convergence of N_t/t to the convex hull of the oracle weight set.

Theorem 9. *Let $\text{conv}(A)$ be the convex hull of a set A . For all $\mu \in \mathcal{M}$ in a multi-answer problem, Track-and-Stop with C -tracking is δ -correct and verifies*

$$\lim_{\delta \rightarrow 0} \frac{\mathbb{E}_{\mu}[\tau_{\delta}]}{\log(1/\delta)} \leq \max_{w \in \text{conv}(w^*(\mu))} \frac{1}{D(w, \mu)}.$$

5.1 Sticky Track-and-Stop

The cases of multiple-answers problems for which Track-and-Stop is inadequate are $\mu \in \mathcal{M}$ with $i_F(\mu)$ of cardinality greater than 1. When convexity does not hold, $w^*(\mu)$ is the union of the convex sets $(w^*(\mu, \neg i))_{i \in i_F(\mu)}$. If an algorithm can a priori select $i_f \in i_F(\mu)$ and track allocations w_t in $w^*(\hat{\mu}_t, \neg i_f)$, then using Track-and-Stop on that restricted problem will ensure that N_t/t converges to the oracle weights. Our proposed algorithm, Sticky Track-and-Stop, which we display in Algorithm 1, uses a confidence region around the current estimate $\hat{\mu}_t$ to determine what $i \in \mathcal{I}$ can be the oracle answer for μ . It selects one of these answers according to an arbitrary total order on \mathcal{I} and does not change it (sticks to it) until no point in the confidence region has the chosen answer in its set of oracle answers.

Algorithm 1 Sticky Track-and-Stop.

Input: $\delta > 0$, strict total order on \mathcal{I} . Set $t = 1$, $\hat{\mu}_0 = 0$, $N_0 = 0$.

while not stopped do

 Let $\mathcal{C}_t = \{\mu' \in \mathcal{M} : D(N_{t-1}, \hat{\mu}_{t-1}, \mu') \leq \log(f(t-1))\}$. // small conf. reg.

 Compute $I_t = \bigcup_{\mu' \in \mathcal{C}_t} i_F(\mu')$.

 Pick the first alternative $i_t \in I_t$ in the order on \mathcal{I} .

 Compute $w_t \in w^*(\hat{\mu}_{t-1}, \neg i_t)$.

 Pull an arm a_t according to the C -tracking rule and receive $X_t \sim \nu_{a_t}$.

 Set $N_t = N_{t-1} + e_{a_t}$ and $\hat{\mu}_t = \hat{\mu}_{t-1} + \frac{1}{N_{t,a_t}}(X_t - \hat{\mu}_{t-1,a_t})e_{a_t}$.

 Let $\mathcal{D}_t = \{\mu' \in \mathcal{M} : D(N_t, \hat{\mu}_t, \mu') \leq \beta(t, \delta)\}$. // large conf. reg.

if there exists $i \in \mathcal{I}$ such that $\mathcal{D}_t \cap \neg i = \emptyset$ **then**

 | stop and return i .

end

$t \leftarrow t + 1$.

end

Theorem 10. *For $\beta(t, \delta) = \log(Ct^2/\delta)$, with C such that $C \geq e \sum_{t=1}^{+\infty} (\frac{e}{K})^K \frac{(\log^2(Ct^2) \log(t))^K}{t^2}$, Sticky Track-and-Stop is δ -correct.*

That result is a consequence of Proposition 12 of [Garivier and Kaufmann, 2016].

5.2 Sample Complexity

Theorem 11. *Sticky Track-and-Stop is asymptotically optimal, i.e. it verifies for all $\mu \in \mathcal{M}$,*

$$\lim_{\delta \rightarrow 0} \frac{\mathbb{E}_{\mu}[\tau_{\delta}]}{\log(1/\delta)} \rightarrow \frac{1}{D(\mu)}.$$

Let $i_{\mu} = \min i_F(\mu)$ in the arbitrary order on answers. For $\varepsilon, \xi > 0$, we define $C_{\varepsilon, \xi}^*(\mu)$, the minimal value of $D(\mathbf{w}', \mu', \neg i_{\mu})$ for \mathbf{w}' and μ' in ε and ξ -neighbourhoods of $\mathbf{w}^*(\mu)$ and μ .

$$C_{\varepsilon, \xi}^*(\mu) = \inf_{\substack{\mu': \|\mu' - \mu\|_{\infty} \leq \xi \\ \mathbf{w}': \inf_{\mathbf{w} \in \mathbf{w}^*(\mu, \neg i_{\mu})} \|\mathbf{w}' - \mathbf{w}\|_{\infty} \leq 3\varepsilon}} D(\mathbf{w}', \mu', \neg i_{\mu}).$$

Our proof strategy is to show that under a concentration event defined below, for t big enough, $(\hat{\mu}_t, N_t/t)$ belongs to that (ξ, ε) neighbourhood of $(\mu, \mathbf{w}^*(\mu, \neg i_{\mu}))$. From that fact, we obtain $D(N_t, \hat{\mu}_t, \neg i_{\mu}) \geq tC_{\varepsilon, \xi}^*(\mu)$. Furthermore, if the algorithm does not stop at stage t , we also get an upper bound on $D(N_t, \hat{\mu}_t, \neg i_{\mu})$ from the stopping condition. We obtain an upper bound on the stopping time, function of δ and $C_{\varepsilon, \xi}^*(\mu)$. By continuity of $(\mathbf{w}, \mu) \mapsto D(\mathbf{w}, \mu, \neg i_{\mu})$ (from Theorem 4), we have $\lim_{\varepsilon \rightarrow 0, \xi \rightarrow 0} C_{\varepsilon, \xi}^*(\mu) = D(\mu, \neg i_{\mu}) = D(\mu)$.

Two concentration events. Let $\mathcal{E}_T = \bigcap_{t=h(T)}^T \{\mu \in \mathcal{C}_t\}$ be the event that the small confidence region contains the true parameter vector μ for $t \geq h(T)$. The function $h : \mathbb{N} \rightarrow \mathbb{R}$, positive, increasing and going to $+\infty$, makes sure that each event $\{\mu \in \mathcal{C}_t\}$ appears in finitely many \mathcal{E}_T , which will be essential in the concentration results. We will eventually use $h(T) = \sqrt{T}$.

In order to define the second event, we first highlight a consequence of Theorem 4.

Corollary 12. *For all $\varepsilon > 0$, for all $\mu \in \mathcal{M}$, for all $i \in \mathcal{I}$, there exists $\xi > 0$ such that*

$$\|\mu' - \mu\|_{\infty} \leq \xi \Rightarrow \forall \mathbf{w}' \in \mathbf{w}^*(\mu', \neg i) \exists \mathbf{w} \in \mathbf{w}^*(\mu, \neg i), \|\mathbf{w}' - \mathbf{w}\|_{\infty} \leq \varepsilon.$$

Let $\mathcal{E}'_T = \bigcap_{t=h(T)}^T \{\|\hat{\mu}_t - \mu\|_{\infty} \leq \xi\}$ be the event that the empirical parameter vector is close to μ , where ξ is chosen as in the previous corollary for $i = i_{\mu}$. The analysis of Sticky Track-and-Stop consists of two parts: first show that \mathcal{E}_T^c and $\mathcal{E}'_T{}^c$ happen rarely enough to lead only to a finite term in $\mathbb{E}_{\mu}[\tau_{\delta}]$; then show that under $\mathcal{E}_T \cap \mathcal{E}'_T$ there is an upper bound on τ_{δ} .

Lemma 13. *Suppose that there exists T_0 such that for $T \geq T_0$, $\mathcal{E}_T \cap \mathcal{E}'_T \subset \{\tau_{\delta} \leq T\}$. Then*

$$\mathbb{E}_{\mu}[\tau_{\delta}] \leq T_0 + \sum_{T=T_0}^{+\infty} \mathbb{P}_{\mu}(\mathcal{E}_T^c) + \sum_{T=T_0}^{+\infty} \mathbb{P}_{\mu}(\mathcal{E}'_T{}^c). \quad (1)$$

Proof. Since τ_{δ} is a non-negative integer-valued random variable, $\mathbb{E}_{\mu}[\tau_{\delta}] = \sum_{T=0}^{+\infty} \mathbb{P}_{\mu}\{\tau_{\delta} > T\}$. For $T \geq T_0$, $\mathbb{P}_{\mu}\{\tau_{\delta} > T\} \leq \mathbb{P}_{\mu}(\mathcal{E}_T^c \cup \mathcal{E}'_T{}^c) \leq \mathbb{P}_{\mu}(\mathcal{E}_T^c) + \mathbb{P}_{\mu}(\mathcal{E}'_T{}^c)$. \square

The sums depending on the events \mathcal{E}_T and \mathcal{E}'_T in (1) are finite for well chosen $h(T)$ and $\mathcal{C}(t)$.

Lemma 14. *For $h(T) = \sqrt{T}$ and $f(t) = \exp(\beta(t, 1/t^5)) = Ct^{10}$ in the definition of the confidence region \mathcal{C}_t , the sum $\sum_{T=T_0}^{+\infty} \mathbb{P}_{\mu}(\mathcal{E}_T^c) + \sum_{T=T_0}^{+\infty} \mathbb{P}_{\mu}(\mathcal{E}'_T{}^c)$ is finite.*

The proof of the Lemma can be found in Appendix G.1. The remainder of the proof is concerned with finding a suitable T_0 . First, we show that if $\hat{\mu}_t$ and N_t/t are in an (ξ, ε) neighbourhood of μ and $\mathbf{w}^*(\mu, \neg i_{\mu})$, then such an upper bound T_0 on τ_{δ} can be obtained.

Lemma 15. *Let t_1 be an integer and suppose that for all $t \geq t_1$, $D(N_t, \hat{\mu}_t, \neg i_{\mu}) \geq tC_{\varepsilon, \xi}^*(\mu)$. Let $T_{\beta} = \inf\{t : t > \beta(t, \delta)/C_{\varepsilon, \xi}^*(\mu)\}$. Then $\tau_{\delta} \leq \max(t_1, T_{\beta})$.*

Proof. Take $t \geq t_1$. If $\tau_{\delta} > t$ then by hypothesis and the stopping condition, $t \leq D(N_t, \hat{\mu}_t, \neg i_{\mu})/C_{\varepsilon, \xi}^*(\mu) \leq \beta(t, \delta)/C_{\varepsilon, \xi}^*(\mu)$. Conversely, for $t \geq t_1$, if $t > \beta(t, \delta)/C_{\varepsilon, \xi}^*(\mu)$ then $\tau_{\delta} \leq t$. We obtain that $\tau_{\delta} \leq \max(t_1, \inf\{t : t > \beta(t, \delta)/C_{\varepsilon, \xi}^*(\mu)\})$. \square

The oracle answer i_t becomes constant. Due to the forced exploration present in the C-tracking procedure, the confidence region \mathcal{C}_t shrinks. After some time, when concentration holds, the set of possible oracle answers I_t becomes constant over t and equal to $i_F(\boldsymbol{\mu})$.

Lemma 16. *If an algorithm guaranties that for all $k \in [K]$ and all $t \geq 1$, $N_{t,k} \geq n(t) > 0$ with $\lim_{t \rightarrow +\infty} n(t)/\log(f(t)) = +\infty$, then there exists T_Δ such that under the event \mathcal{E}_T , for $t \geq \max(h(T), T_\Delta)$, $I_t = i_F(\boldsymbol{\mu})$ and $\min I_t = i_\mu = \min i_F(\boldsymbol{\mu})$.*

Proof in Appendix G.4. Note that Lemma 16 depends only on the amount of forced exploration and not on other details of the algorithm. Any algorithm using C-tracking verifies the hypothesis with $n(t) = \sqrt{t + K^2} - 2K$ by Lemma 34 [Garivier and Kaufmann, 2016, Lemma 7].

Convergence to the neighbourhood of $(\boldsymbol{\mu}, \mathbf{w}^*(\boldsymbol{\mu}, -i_\mu))$. Once $i_t = i_\mu$, we fall back to tracking points from a convex set of oracle weights. The estimate $\hat{\boldsymbol{\mu}}_t$ and N_t/t both converge, to $\boldsymbol{\mu}$ and to the set $\mathbf{w}^*(\boldsymbol{\mu}, -i_\mu)$. The Lemma below is proved in Appendix G.5.

Lemma 17. *Let T_Δ be defined as in Lemma 16. For T such that $h(T) \geq T_\Delta$, it holds that on $\mathcal{E}_T \cap \mathcal{E}'_T$ Sticky Track-and-Stop with C-Tracking verifies*

$$\forall t \geq h(T), \|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi, \quad \text{and} \quad \forall t \geq 4\frac{K^2}{\varepsilon^2} + 3\frac{h(T)}{\varepsilon}, \quad \inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, -i_\mu)} \left\| \frac{N_t}{t} - \mathbf{w} \right\|_\infty \leq 3\varepsilon.$$

Remainder of the proof. Suppose that the event $\mathcal{E}_T \cap \mathcal{E}'_T$ holds. Let T_Δ be defined as in Lemma 16 and T be such that $h(T) \geq T_\Delta$. Let $\eta(T) = 4K^2/\varepsilon^2 + 3h(T)/\varepsilon$. For all $t \geq \eta(T)$ we have $D(N_t, \hat{\boldsymbol{\mu}}_t, -i_\mu) \geq tC_{\varepsilon, \xi}^*(\boldsymbol{\mu})$ by Lemma 17. For $h(T)$ bigger than some T_η we have $\eta(T) \leq T$. We suppose $h(T) \geq \max(T_\Delta, T_\eta)$. We apply Lemma 15 with $t_1 = \eta(T)$. We obtain that $\tau_\delta \leq \max(\eta(T), T_\beta) \leq \max(T, T_\beta)$. Conclusion: for $T \geq T_0 = \max(h^{-1}(T_\Delta), h^{-1}(T_\eta), T_\beta)$, under the concentration event, $\tau_\delta \leq T$ and we can apply Lemma 13.

Note that $\lim_{\delta \rightarrow 0} \frac{T_0}{\log(1/\delta)} = \frac{1}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}$. Taking $\varepsilon \rightarrow 0$ (hence $\xi \rightarrow 0$ as well), we obtain $\lim_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\log(1/\delta)} = \frac{1}{\lim_{\varepsilon \rightarrow 0} C_{\varepsilon, \xi}^*(\boldsymbol{\mu})} = \frac{1}{D(\boldsymbol{\mu})}$. We proved Theorem 11.

6 Conclusion

We characterized the complexity of multiple-answers pure exploration bandit problems, showing a lower bound and exhibiting an algorithm with asymptotically matching sample complexity on all such problems. That study could be extended in several interesting directions and we now list a few.

- The computational complexity of Track-and-Stop is an important issue: it would be desirable to design a pure exploration algorithm with optimal sample complexity which does not need to solve a min-max problem at each step. Furthermore, the same would need to be done for the sticky selection of an answer for the multiple-answers setting.
- Both lower bounds and upper bounds in this paper are asymptotic. In the upper bound case, only the forced exploration rounds are considered when evaluating the convergence of $\hat{\boldsymbol{\mu}}_t$ to $\boldsymbol{\mu}$, giving rise to potentially sub-optimal lower order terms. A finite time analysis with reasonably small $o(\log(1/\delta))$ terms for an optimal algorithm is desirable. In addition, while selecting one of the oracle answers to stick to has no asymptotic cost, it could have a lower order effect on the sample complexity and appear in a refined lower bound.
- Current tools in the theory of Brownian motion are insufficient to characterise the asymptotic distribution of proportions induced by tracking, even for two arms. Without tracking the Arcsine law arises, so this slightly more challenging problem holds the promise of similarly elegant results.
- Finally, the multiple answer pure exploration setting can be extended in various ways. Making \mathcal{I} continuous leads to regression problems. The parametric assumption that the arms are in one-parameter exponential families could also be relaxed.

References

- J.-Y. Audibert, S. Bubeck, and R. Munos. Best Arm Identification in Multi-armed Bandits. In *Proceedings of the 23rd Conference on Learning Theory*, 2010.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine Learning*, 47(2):235–256, 2002.
- D. A. Berry and B. Fristedt. *Bandit Problems. Sequential allocation of experiments*. Chapman and Hall, 1985.
- D. Blackwell and M. A. Girshick. *Theory of games and statistical decisions*. Wiley, 1954.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004. ISBN 0521833787.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- S. Bubeck, R. Munos, and G. Stoltz. Pure Exploration in Finitely Armed and Continuous Armed Bandits. *Theoretical Computer Science 412, 1832-1852*, 412:1832–1852, 2011.
- R. M. Castro. Adaptive sensing performance lower bounds for sparse signal detection and support estimation. *Bernoulli*, 20(4):2217–2246, 11 2014.
- L. Chen, A. Gupta, J. Li, M. Qiao, and R. Wang. Nearly optimal sampling algorithms for combinatorial pure exploration. In *Proceedings of the 2017 Conference on Learning Theory*, volume 65 of *Proceedings of Machine Learning Research*, pages 482–534. PMLR, July 2017.
- S. Chen, T. Lin, I. King, M. Lyu, and W. Chen. Combinatorial Pure Exploration of Multi-Armed Bandits. In *Advances in Neural Information Processing Systems*, 2014.
- R. Combes, S. Magureanu, and A. Proutiere. Minimal exploration in structured stochastic bandits. In *Advances in Neural Information Processing Systems*, pages 1763–1771, 2017.
- E. Even-Dar, S. Mannor, and Y. Mansour. PAC bounds for multi-armed bandit and Markov decision processes. In *Computational Learning Theory, 15th Annual Conference on Computational Learning Theory, COLT 2002, Sydney, Australia, July 8-10, 2002, Proceedings*, volume 2375 of *Lecture Notes in Computer Science*, pages 255–270. Springer, 2002. ISBN 3-540-43836-X.
- E. Even-Dar, S. Mannor, and Y. Mansour. Action Elimination and Stopping Conditions for the Multi-Armed Bandit and Reinforcement Learning Problems. *Journal of Machine Learning Research*, 7: 1079–1105, 2006.
- E. A. Feinberg, P. O. Kasyanov, and M. Voorneveld. Berge’s maximum theorem for noncompact image sets. *Journal of Mathematical Analysis and Applications*, 413(2):1040–1046, 2014.
- V. Gabillon, M. Ghavamzadeh, and A. Lazaric. Best Arm Identification: A Unified Approach to Fixed Budget and Fixed Confidence. In *Advances in Neural Information Processing Systems*, 2012.
- V. Gabillon, A. Lazaric, M. Ghavamzadeh, R. Ortner, and P. L. Bartlett. Improved learning complexity in combinatorial pure exploration bandits. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics, AISTATS 2016, Cadiz, Spain, May 9-11, 2016*, volume 51 of *JMLR Workshop and Conference Proceedings*, pages 1004–1012. JMLR.org, 2016.
- A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In *Proceedings of the 29th Conference On Learning Theory (COLT)*, 2016.
- A. Garivier and E. Kaufmann. Non-asymptotic sequential tests for overlapping hypotheses and application to near optimal arm identification in bandit models. In *arXiv 1905.03495*, 2019.
- A. Garivier, E. Kaufmann, and W. M. Koolen. Maximin action identification: A new bandit framework for games. In *Proceedings of the 29th Annual Conference on Learning Theory (COLT)*, pages 1028 – 1050, June 2016.

- R. Huang, M. M. Ajallooeian, C. Szepesvári, and M. Müller. Structured best arm identification with fixed confidence. In *International Conference on Algorithmic Learning Theory, ALT 2017, 15-17 October 2017, Kyoto University, Kyoto, Japan*, volume 76 of *Proceedings of Machine Learning Research*, pages 593–616. PMLR, 2017.
- K. Jamieson, M. Malloy, R. Nowak, and S. Bubeck. lil’UCB: an Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of the 27th Conference on Learning Theory*, 2014.
- S. Kalyanakrishnan and P. Stone. Efficient Selection in Multiple Bandit Arms: Theory and Practice. In *International Conference on Machine Learning (ICML)*, 2010.
- S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2012.
- Z. Karnin, T. Koren, and O. Somekh. Almost optimal Exploration in multi-armed bandits. In *International Conference on Machine Learning (ICML)*, 2013.
- S. Katariya, B. Kveton, C. Szepesvári, C. Vernade, and Z. Wen. Stochastic rank-1 bandits. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, AISTATS 2017, 20-22 April 2017, Fort Lauderdale, FL, USA*, volume 54 of *Proceedings of Machine Learning Research*, pages 392–401. PMLR, 2017.
- E. Kaufmann and S. Kalyanakrishnan. Information complexity in bandit subset selection. In *Proceeding of the 26th Conference On Learning Theory.*, 2013.
- E. Kaufmann and W. M. Koolen. Monte-Carlo tree search by best arm identification. In *Advances in Neural Information Processing Systems (NeurIPS) 30*, pages 4904–4913, Dec. 2017.
- E. Kaufmann, O. Cappé, and A. Garivier. On the Complexity of Best Arm Identification in Multi-Armed Bandit Models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.
- E. Kaufmann, W. M. Koolen, and A. Garivier. Sequential test for the lowest mean: From Thompson to Murphy sampling. In *Advances in Neural Information Processing Systems (NeurIPS) 31*, pages 6333–6343. Curran Associates, Inc., Dec. 2018.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6(1):4–22, 1985.
- A. Locatelli, M. Gutzeit, and A. Carpentier. An optimal algorithm for the thresholding bandit problem. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, volume 48 of *JMLR Workshop and Conference Proceedings*, pages 1690–1698. JMLR.org, 2016.
- S. Magureanu, R. Combes, and A. Proutière. Lipschitz bandits: Regret lower bound and optimal algorithms. In *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, volume 35 of *JMLR Workshop and Conference Proceedings*, pages 975–999. JMLR.org, 2014.
- O. Maron and A. Moore. The Racing algorithm: Model selection for Lazy learners. *Artificial Intelligence Review*, 11(1-5):113–131, 1997.
- D. Russo. Simple Bayesian algorithms for best arm identification. *CoRR*, abs/1602.08448, 2016.
- K. Teraoka, K. Hatano, and E. Takimoto. Efficient sampling method for Monte Carlo tree search problem. *IEICE Transactions*, 97-D(3):392–398, 2014.
- Y. Zhou, J. Li, and J. Zhu. Identify the Nash equilibrium in static games with random payoffs. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 4160–4169, International Convention Centre, Sydney, Australia, Aug. 2017. PMLR.

A Notations

Concept	Symbol
Exponential family mean parameter	$\mu \in \mathcal{O}$
Bandit model	$\boldsymbol{\mu} \in \mathcal{M} \subseteq \mathcal{O}^K$
Possible answers	\mathcal{I}
Correct answer for bandit model $\boldsymbol{\mu}$	$i^*(\boldsymbol{\mu}) : \mathcal{M} \rightarrow 2^{\mathcal{I}}$
Alternative to $i \in \mathcal{I}$	$\neg i = \{\boldsymbol{\mu} \in \mathcal{M} : i \notin i^*(\boldsymbol{\mu})\}$
K -Simplex	Δ_K
Non-negative orthant	\mathcal{Q}_+
interior, closure, convex hull of set A	$\overset{\circ}{A}, \text{cl}(A), \text{conv}(A)$
Oracle answers and weights	$i_F(\boldsymbol{\mu}) : \mathcal{M} \rightarrow 2^{\mathcal{I}}, \mathbf{w}^*(\boldsymbol{\mu}) : \mathcal{M} \rightarrow \mathcal{P}(\Delta_K)$
Bandit arm	$k \in [K]$
Number of samples of arm k at time t	$N_{t,k}$
mean of samples of arm k at time t	$\hat{\mu}_{t,k}$

B Visualisation of Multiple-Answer Pure Exploration Problems

Identification Problem	Possible answers \mathcal{I}	Correct answers $i^*(\boldsymbol{\mu}) \subseteq \mathcal{I}$	Correct $i^*(\boldsymbol{\mu})$	Oracle $i_F(\boldsymbol{\mu})$
ϵ Best Arm	$[K]$	$\{k \mid \mu_k \geq \max_{j \neq k} \mu_j - \epsilon\}$		
Thresholding Bandit	2^K	$\{k \mid \mu_k \leq \gamma\}$		
ϵ Minimum Threshold	$\{\text{lo}, \text{hi}\}$	$\{\text{lo}\}$ if $\min_k \mu_k < \gamma - \epsilon$ $\{\text{hi}\}$ if $\min_k \mu_k > \gamma + \epsilon$ $\{\text{lo}, \text{hi}\}$ o.w.		
Any Low Arm	$[K] \cup \{\text{no}\}$	$\{k \mid \mu_k \leq \gamma\}$ if $\min_k \mu_k < \gamma$ $\{\text{no}\}$ if $\min_k \mu_k > \gamma$		
Any Sign	$[K] \times \{\text{lo}, \text{hi}\}$	$\{(k, \text{lo}) \mid \mu_k \leq \gamma\} \cup \{(k, \text{hi}) \mid \mu_k \geq \gamma\}$		

Table 1: Collection of Identification Problems. The diagrams depict 2-arm instances, parameterised by the two means, with colours showing the set of correct answers: **one correct answer:** ■ $\{1\}$, ■ $\{2\}$, ■ $\{\{1\}\}$, ■ $\{\{2\}\}$, ■ $\{\{1, 2\}\}$, ■ $\{\text{lo}\}$, ■ $\{\emptyset\}/\{\text{hi}\}/\{\text{no}\}$, ■ $\{(1, \text{lo})\}$, ■ $\{(1, \text{hi})\}$, ■ $\{(2, \text{lo})\}$, ■ $\{(2, \text{hi})\}$, and **two correct answers:** ■ $\{1, 2\}$, ■ $\{\text{lo}, \text{hi}\}$, ■ $\{(1, \text{lo}), (2, \text{lo})\}$, ■ $\{(1, \text{lo}), (2, \text{hi})\}$, ■ $\{(1, \text{hi}), (2, \text{lo})\}$, ■ $\{(1, \text{hi}), (2, \text{hi})\}$.

C Lower bound Proofs

In this section we build up to the proof of the lower bound Theorem 1. We start with the minimax result Lemma 2.

C.1 Lemma 2: Characteristic time as the value of a game

Let $\neg j$ be the set of bandit problems for which j is not a valid answer. Let us define the characteristic time by

$$\frac{1}{T_j^*(\boldsymbol{\mu})} = D(\boldsymbol{\mu}, \neg j) = \sup_{\mathbf{w}} \inf_{\boldsymbol{\lambda} \in \neg j} \sum_k w_k d(\mu_k, \lambda_k).$$

We now view the characteristic time problem as defining a two-player zero-sum game, with the purpose of obtaining a minimax optimal mixed strategy for the inf player. We will use such a mixed strategy to construct hard learning problems for proving sample complexity lower bounds in the next section. In this section we focus on the existence of the minimax strategy. We define the *bandit complexity* game to be the semi-infinite two-player zero-sum simultaneous game where:

- MAX's pure strategies are arms $i \in [K] = \{1, \dots, K\}$,
- MIN's pure strategies are bandit models $\lambda \in \neg j \subseteq \mathcal{M}$ (we may equivalently have MIN play a point $s \in S := \{(d(\mu_1, \lambda_1), \dots, d(\mu_K, \lambda_K)) \mid \lambda \in \neg j\} \subseteq [0, \infty)^K$),
- the payoff function is $(i, \lambda) \mapsto d(\mu_i, \lambda_i)$ (or, equivalently, $(i, s) \mapsto s_i$).

By definition, $D(\mu, \neg j)$ is the optimal payoff when MAX randomises and plays first. We aim to show that a matching randomised strategy exists for when MIN plays first. That is, we want to establish a min-max theorem.

of Lemma 2. Combining (a) a standard application of Sion's minimax theorem to the bilinear function $f : \Delta \times \text{conv}(S) \rightarrow \mathbb{R}$ defined by $f(\mathbf{w}, s) = \langle \mathbf{w}, s \rangle$ and (b) the support size insight of Blackwell and Girshick [1954, Theorem 2.4.2] yields the Lemma. \square

For convenience, we will assume in the remainder that the infimum above is attained (e.g. when the convex hull of S is compact), possibly on the closure of $\neg j$. If it is not, we need to apply the below arguments to a sequence of ϵ -suboptimal \mathbb{P} and let $\epsilon \rightarrow 0$. At any rate, we assume there exist $\lambda^1, \dots, \lambda^K \in \neg j$ (or its closure) and $\mathbf{q} \in \Delta_K$ such that

$$\forall i : \sum_k q_k d(\mu_i, \lambda_i^k) \leq D(\mu, \neg j). \quad (2)$$

C.2 Consequences of the Minimax result

In this section we build up lower bounds by relating the probability of any event between two or more bandit problem. We start with a useful change of measures observation, used in [Garivier and Kaufmann, 2019] to derive a lower bound on the sample complexity of ϵ -Best Arm Identification problems.

Proposition 18. *Consider two distributions \mathbb{P} and \mathbb{Q} . Let us denote the log-likelihood ratio after n rounds by $L_n = \ln \frac{d\mathbb{P}}{d\mathbb{Q}}$. Then for any measurable event $A \in \mathcal{F}_n$ and threshold $\gamma \in \mathbb{R}$,*

$$\mathbb{Q}(A) \geq e^{-\gamma} (\mathbb{P}(A) - \mathbb{P}\{L_n > \gamma\}). \quad (3)$$

Proof.

$$\begin{aligned} \mathbb{Q}(A) &= \mathbb{E}_{\mathbb{P}}[\mathbb{I}_A e^{-L_n}] \geq \mathbb{E}_{\mathbb{P}}[\mathbb{I}_{A \cap \{L_n \leq \gamma\}} e^{-L_n}] \\ &\geq \mathbb{P}(A \cap \{L_n \leq \gamma\}) e^{-\gamma} \geq e^{-\gamma} (\mathbb{P}(A) - \mathbb{P}\{L_n > \gamma\}). \end{aligned}$$

\square

C.3 Likelihood ratio Martingales

Next we investigate the specific form of the likelihood ratio between two bandit models. Fix bandit models μ and λ , and any sampling strategy. Then after n rounds,

$$\ln \frac{d\mathbb{P}_{\mu}}{d\mathbb{P}_{\lambda}} = \sum_i N_{n,i} \text{KL}(\nu_{\mu_i, i}, \nu_{\lambda_i, i}) + M_n(\mu, \lambda)$$

where $M_n(\mu, \lambda)$ is a martingale. To see this, we write $\text{KL}(\nu_{\mu_i, i}, \nu_{\lambda_i, i}) = d(\mu, \lambda) = \phi(\mu) - \phi(\lambda) - (\mu - \lambda)\phi'(\lambda)$, where we write ϕ for the convex generator of the Bregman divergence $d(\cdot, \cdot)$. Then

$$\begin{aligned} \ln \frac{d\mathbb{P}_{\mu}}{d\mathbb{P}_{\lambda}} &= \sum_i N_{n,i} (d(\hat{\mu}_{n,i}, \lambda_i) - d(\hat{\mu}_{n,i}, \mu_i)) \\ &= \sum_i N_{n,i} \left(d(\mu_i, \lambda_i) + (\phi'(\mu_i) - \phi'(\lambda_i)) (\hat{\mu}_{n,i} - \mu_i) \right) \end{aligned}$$

hence $M_n(\mu, \lambda) = \sum_i N_{n,i} (\phi'(\mu_i) - \phi'(\lambda_i)) (\hat{\mu}_{n,i} - \mu_i)$.

C.4 Exploiting the Minimax Distribution

We now bound the probability of any event between $\boldsymbol{\mu}$ and the hard problems given by the minimax distribution.

Lemma 19. *Fix a bandit model $\boldsymbol{\mu}$ with sub-Gaussian arm distributions. Let \mathbf{q} and $\boldsymbol{\lambda}^1, \dots, \boldsymbol{\lambda}^K$ be a minimax witness from Lemma 2, and let us introduce the abbreviation $\alpha_i = \phi'(\mu_i) - \sum_k q_k \phi'(\lambda_i^k)$. Fix sample size n , and consider any event $A \in \mathcal{F}_n$. Then for any $\beta > 0$*

$$\max_{k \in [K]} \mathbb{P}_{\boldsymbol{\lambda}^k} \{A\} \geq e^{-\frac{n}{T^*(\boldsymbol{\mu})} - \beta} \left(\mathbb{P}_{\boldsymbol{\mu}} \{A\} - \exp\left(\frac{-\beta^2}{2n \max_i \alpha_i^2}\right) \right).$$

In words, if A is likely under $\boldsymbol{\mu}$ then it must also be likely under at least one $\boldsymbol{\lambda}^k$ for sample sizes $n \ll T^*(\boldsymbol{\mu})$.

Proof. Let us form the (Bayesian) mixture distribution $\mathbb{P}_{\mathbf{q}} = \sum_k q_k \mathbb{P}_{\boldsymbol{\lambda}^k}$. We have

$$L_n = -\ln \frac{d\mathbb{P}_{\mathbf{q}}}{d\mathbb{P}_{\boldsymbol{\mu}}} \leq \sum_k q_k \ln \frac{d\mathbb{P}_{\boldsymbol{\mu}}}{d\mathbb{P}_{\boldsymbol{\lambda}^k}}.$$

It follows that for any $\gamma \in \mathbb{R}$ we have

$$\{L_n > \gamma\} \subseteq \left\{ \sum_k q_k \sum_i N_{n,i} d(\mu_i, \lambda_i^k) + \sum_k q_k M_n(\boldsymbol{\mu}, \boldsymbol{\lambda}^k) > \gamma \right\}.$$

Picking $\gamma = \frac{n}{T^*(\boldsymbol{\mu})} + \beta$, we find

$$= \left\{ \sum_k q_k \sum_i N_{n,i} d(\mu_i, \lambda_i^k) + \sum_k q_k M_n(\boldsymbol{\mu}, \boldsymbol{\lambda}^k) > \frac{n}{T^*(\boldsymbol{\mu})} + \beta \right\}$$

Since $(\mathbf{w}^*, \mathbf{q})$ is a Nash equilibrium of the game and N_n/n is a mixed strategy for the first player, $\sum_k q_k \sum_i N_{n,i} d(\mu_i, \lambda_i^k) \leq n \sum_k q_k \sum_i w_i^* d(\mu_i, \lambda_i^k) = \frac{n}{T^*(\boldsymbol{\mu})}$, so we find

$$\begin{aligned} &\subseteq \left\{ \sum_k q_k M_n(\boldsymbol{\mu}, \boldsymbol{\lambda}^k) > \beta \right\} \\ &= \left\{ \sum_k q_k \sum_i N_{n,i} (\phi'(\mu_i) - \phi'(\lambda_i^k)) (\hat{\mu}_{n,i} - \mu_i) > \beta \right\} \\ &= \left\{ \sum_i N_{n,i} \alpha_i (\hat{\mu}_{n,i} - \mu_i) > \beta \right\} \end{aligned} \tag{4}$$

The above left-hand quantity is a martingale of length n . Using the sub-Gaussianity assumption, the Hoeffding-Azuma inequality gives

$$\mathbb{P}_{\boldsymbol{\mu}} \{L_n > \gamma\} \leq \exp\left(\frac{-\beta^2}{2n \max_i \alpha_i^2}\right).$$

Let A be a \mathcal{F}_n -measurable event. Combination with the change of measure argument (3) with $\max_k \mathbb{P}_{\boldsymbol{\lambda}^k} \{A\} \geq \mathbb{P}_{\mathbf{q}} \{A\}$ gives the result. \square

The sub-Gaussian assumption of the Lemma can undoubtedly be relaxed. The crucial requirement is that the n -step martingale in (4) concentrates, and hence cannot be large w.h.p.

We are now ready for the proof of Theorem 1, in which we will carefully tune the time n to which we apply the above result.

C.5 Proof of Theorem 1

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 which is returned. Since the algorithm is δ -PAC by hypothesis, $\mathbb{P}_\mu(\tau_\delta \leq T, \hat{i}_\delta \notin i^*(\mu)) \leq \delta$. So

$$\begin{aligned} \mathbb{P}_\mu(\tau_\delta \leq T) &= \sum_i \mathbb{P}_\mu(\tau_\delta \leq T, \hat{i}_\delta = i) \\ &\leq \delta + \sum_{i \in i^*(\mu)} \mathbb{P}_\mu(\tau_\delta \leq T, \hat{i}_\delta = i). \end{aligned}$$

For $i \in i^*(\mu)$, fix a minimax strategy $\lambda^1, \dots, \lambda^K \in \neg i$ and $\mathbf{q} \in \Delta_K$ as given by Lemma 2. Then by Lemma 19, for any $\beta > 0$

$$\begin{aligned} \mathbb{P}_\mu(\tau_\delta \leq T, \hat{i}_\delta = i) &\leq \exp\left(\frac{T}{T_i^*(\mu)} + \beta\right) \max_k \mathbb{P}_{\lambda^k}(\tau_\delta \leq T, \hat{i}_\delta = i) + \exp\left(\frac{-\beta^2}{2T \max_i \alpha_i^2}\right) \\ &\leq \delta \exp\left(\frac{T}{T_i^*(\mu)} + \beta\right) + \exp\left(\frac{-\beta^2}{2T \max_i \alpha_i^2}\right) \end{aligned}$$

Let $\alpha^2 = \max_i \alpha_i^2$. For $\eta \in (0, 1)$, $T \leq (1 - \eta)T_i^*(\mu) \log(1/\delta)$ and $\beta = \frac{\eta}{2\sqrt{1-\eta}} \sqrt{\frac{T}{T_i^*(\mu)} \log(1/\delta)}$,

$$\begin{aligned} \mathbb{P}_\mu(\tau_\delta \leq T, \hat{i}_\delta = i) &\leq \delta \exp\left(\frac{T}{T_i^*(\mu)} + \frac{\eta}{2\sqrt{1-\eta}} \sqrt{\frac{T}{T_i^*(\mu)} \log(1/\delta)}\right) + \exp\left(\frac{-\eta^2 \log(1/\delta)}{8(1-\eta)T_i^*(\mu)\alpha^2}\right) \\ &\leq \delta \exp\left((1 - \eta/2) \log \frac{1}{\delta}\right) + \exp\left(\frac{-\eta^2 \log(1/\delta)}{8(1-\eta)T_i^*(\mu)\alpha^2}\right) \\ &= \delta^{\eta/2} + \delta^{\eta^2/(8(1-\eta)T_i^*(\mu)\alpha^2)} \\ &\xrightarrow{\delta \rightarrow 0} 0 \end{aligned}$$

Suppose that $T = \min_{i \in i^*(\mu)} (1 - \eta)T_i^*(\mu) \log(1/\delta)$ for some $\eta \in (0, 1)$. Then we must have $\lim_{\delta \rightarrow 0} \mathbb{P}_\mu(\tau_\delta \leq T) = 0$ and therefore $\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\log(1/\delta)} \geq \lim_{\delta \rightarrow 0} \frac{T}{\log(1/\delta)} (1 - \mathbb{P}_\mu(\tau_\delta \leq T)) = (1 - \eta) \min_{i \in i^*(\mu)} T_i^*(\mu)$. Letting η go to zero, we obtain that

$$\liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\log(1/\delta)} \geq \min_{i \in i^*(\mu)} T_i^*(\mu).$$

D Vanilla Track and Stop Fails for Multiple Answers

We argue that Track and Stop in general does not ensure convergence of N_t/t to $\mathbf{w}^*(\mu)$ when that set is not convex. We illustrate our claim on the *Any Half-Space* problem, a generalisation of *Any Sign* from Table 1. Given $n \in \mathbb{N}$ hyperplanes of \mathbb{R}^K passing through 0, parametrized for $m \in [n]$ by normal vectors $\mathbf{u}_m \in \mathbb{R}^K$ with $\|\mathbf{u}_m\|_1 = 1$, the algorithm has to return $(m, s) \in [n] \times \{-1, 1\}$ such that $s\mu^\top \mathbf{u}_m \geq 0$. i.e. it must return any of the half-spaces in which μ lies. See Figure 3. The arms have Gaussian distributions with variance 1.

This problem is chosen for the simplicity of its \mathbf{w}^* mapping. Indeed $\mathbf{w}^*(\mu, \neg(m, s)) = \{\mathbf{u}_m\}$ if $s\mu^\top \mathbf{u}_m \geq 0$ and $\mathbf{w}^*(\mu, \neg(m, s)) = \Delta_K$ otherwise. The distance to that alternative is $D(\mu, \neg(m, s)) = \mathbb{I}\{s\mu^\top \mathbf{u}_m \geq 0\}(\mu^\top \mathbf{u}_m)^2$. The optimal weights set $\mathbf{w}^*(\mu)$ is the union of those $\{\mathbf{u}_m\}$ for which the distance is the greatest, and can be non-convex.

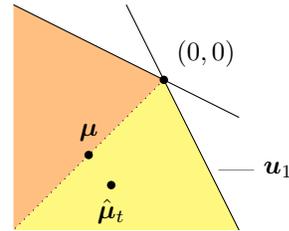


Figure 3: *Any Half-Space* problem. The oracle answers $i_F(\mu)$ are $\square \{\mathbf{u}_1\}$ and $\square \{\mathbf{u}_2\}$ (and both on the diagonal).

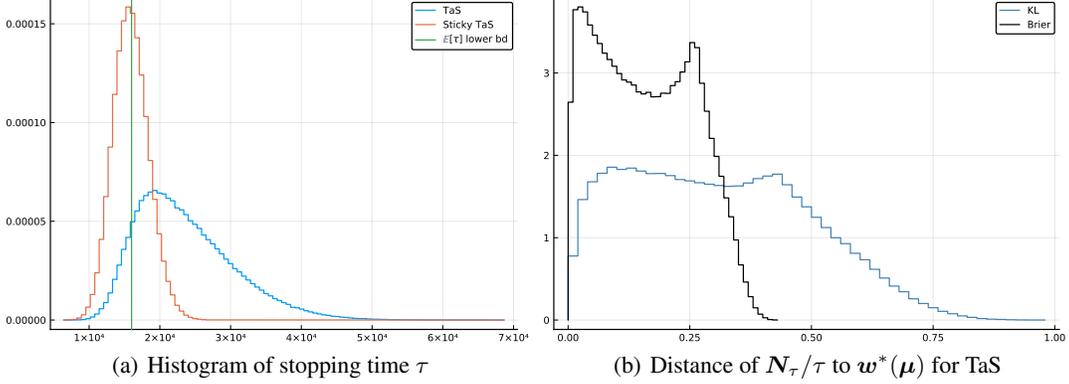


Figure 2: Suboptimality of Track-and-Stop with C -tracking on a $K = 10$ arm instance of the *Any Half-Space* problem, where each $\mu_k = -1/10$. We run the algorithm with an excessively small $\delta = e^{-80}$ to focus on the asymptotic regime, using $500K$ repetitions. *Left*: empirical distribution of the stopping time of Track-and-Stop and of Sticky Track-and-Stop, with respective means 24K and 16K. *Right*: the reason for the suboptimality is that the sampling proportions of Track-and-Stop do not converge to the oracle weights.

For $a \in [0, 1]$, let $K = 2$, $n = 2$, $\mathbf{u}_1 = (a, 1 - a)$, $\mathbf{u}_2 = (1 - a, a)$ and $\boldsymbol{\mu} = (\mu_0, \mu_0)$ for some $\mu_0 \in \mathbb{R}$. Suppose that after stage t_0 of Track and Stop, $\hat{\boldsymbol{\mu}}_t$ verifies that $\hat{\boldsymbol{\mu}}_t^\top \mathbf{u}_m$ has same sign as $\boldsymbol{\mu}^\top \mathbf{u}_m$ for both $m \in \{1, 2\}$ (in expectation this happens except on at most a finite number of stages). Then $\mathbf{w}^*(\hat{\boldsymbol{\mu}}_t) = \{\mathbf{u}_1\}$ iff $\hat{\mu}_{t,1} > \hat{\mu}_{t,2}$ and $\mathbf{w}^*(\hat{\boldsymbol{\mu}}_t) = \{\mathbf{u}_2\}$ iff $\hat{\mu}_{t,1} < \hat{\mu}_{t,2}$. The case $\hat{\mu}_{t,1} = \hat{\mu}_{t,2}$ has probability 0, hence we ignore it.

C -tracking ensures that N_t/t is close to $\sum_{s=1}^t \mathbf{w}_s$ for $\mathbf{w}_s \in \mathbf{w}^*(\hat{\boldsymbol{\mu}}_s)$. Calling $T_1(t)$ the number of stages up to t for which $\hat{\mu}_{t,1} > \hat{\mu}_{t,2}$ and neglecting the first t_0 stages, $N_t/t \approx T_1(t)\mathbf{u}_1 + (1 - T_1(t))\mathbf{u}_2$. In the nomenclature of random walks, $T_1(t)$ is the occupation time of the region below the diagonal on Figure 3.

In order for Track and Stop to be optimal, N_t/t need to be close to $\mathbf{w}^*(\boldsymbol{\mu}) = \{\mathbf{u}_1, \mathbf{u}_2\}$ at $t_{opt} = \log(1/\delta)/D(\boldsymbol{\mu})$. For $\mathbf{u}_1 \neq \mathbf{u}_2$ this means that the distribution of $T_1(t_{opt})/t_{opt}$ must be concentrated on $\{0, 1\}$. If the limit distribution of $T_1(t)/t$ (assuming it exists) for $t \rightarrow +\infty$ has mass in $(0, 1)$, Track and Stop likely has suboptimal asymptotic sample complexity.

In the case of $a = 1/2$, $\mathbf{u}_1 = \mathbf{u}_2 = (1/2, 1/2)$, for t even $N_{t,1} = N_{t,2}$ and $T_1(t) = \#\{s \leq t : \sum_{u=1}^{s/2} (X_{2u}^{(1)} - X_{2u+1}^{(2)}) > 0\}$ is the occupation time of \mathbb{R}^+ for a Gaussian random walk. Its limit distribution is the Arcsine distribution. But in that case N_t/t is always optimal. Experimentally, when $a \neq 1/2$ (hence $\mathbf{u}_1 \neq \mathbf{u}_2$ and N_t/t not always optimal), we observe that the limit distribution for $T_1(t)/t$ is not Arcsine, but has mass in $(0, 1)$ for $a \in (0, 1)$. See Figure 4.

Figure 2 displays the stopping time of Track and Stop on such an hyperplane problem for $K = n = 10$ and shows that Track and Stop is empirically suboptimal and that N_τ/τ , proportions of pulls at the stopping time, is not concentrated near $\mathbf{w}^*(\boldsymbol{\mu})$.

E Failure of D-tracking

We illustrate the suboptimality of Track-and-Stop with D-tracking in a single-answer problem. On a 3-arms problem with Gaussian distributions with same variance, the algorithm must answer the two following queries:

- What is the sign of μ_1 ?
- Is there $m \in \{1, 2\}$ such that $\boldsymbol{\mu}^\top \mathbf{u}^{(m)} \geq 0$? For some $a \in (0, 1)$, $\mathbf{u}^{(1)} = (0, a, 1 - a)$ and $\mathbf{u}^{(2)} = (0, 1 - a, a)$.

The possible answers are $\{(+, \text{yes}), (+, \text{no}), (-, \text{yes}), (-, \text{no})\}$, and there is only one correct answer.

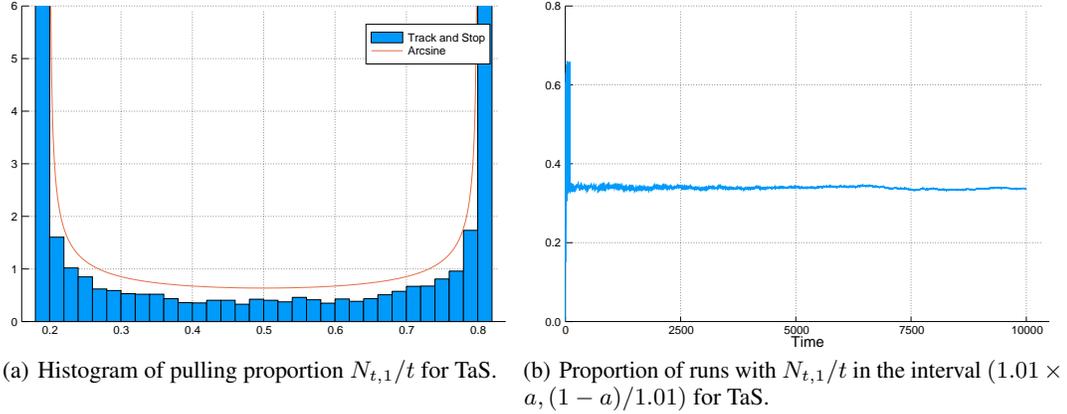


Figure 4: Pulling proportions of Track-and-Stop with C -tracking on the *Any Half-Space* problem with $K = 2$ and $a = 1/5$. Both figures use data from the same 10000 runs of Track-and-Stop. *Left*: histogram at $T = 10000$ of $N_{t,1}/t$, normalized such that the total area of the bars is 1. The values of the leftmost and rightmost bars are ≈ 15 . The probability distribution function of the Arcsine distribution is shown for comparison. *Right*: evolution over time of the mass in the interval $(1.01 \times a, \frac{1-a}{1.01})$ (non-extremal bars on the left).

The divergence from μ to an alternative where the sign is flipped is $D_1(\mu) = \frac{1}{2}\mu_1^2$. If $\mu^\top \mathbf{u}^{(m)} < 0$ for both $\mathbf{u}^{(m)}$, then the divergence from μ to an alternative where the second answer is yes is $D_{2,3} = \max(\frac{1}{2}(\mu^\top \mathbf{u}^{(1)})^2, \frac{1}{2}(\mu^\top \mathbf{u}^{(2)})^2)$. Let $D_2(\mu)$ and $D_3(\mu)$ be the two terms in that maximum. In that case, the oracle weights $\mathbf{w}^*(\mu)$ are

- $\{\mathbf{w}^{(1)}(\mu)\}$ if $D_2 > D_3$, where $\mathbf{w}^{(1)}(\mu) = (\frac{D_2}{D_1+D_2}, a\frac{D_1}{D_1+D_2}, (1-a)\frac{D_1}{D_1+D_2})$,
- $\{\mathbf{w}^{(2)}(\mu)\}$ if $D_2 < D_3$, where $\mathbf{w}^{(2)}(\mu) = (\frac{D_3}{D_1+D_3}, (1-a)\frac{D_1}{D_1+D_3}, a\frac{D_1}{D_1+D_3})$,
- the convex hull of $\{\mathbf{w}^{(1)}(\mu), \mathbf{w}^{(2)}(\mu)\}$ if $D_2 = D_3$.

Let $\mu = (-\mu_0, -\mu_0, -\mu_0)$ for $\mu_0 > 0$. Then $\mathbf{w}^*(\mu) = \text{conv}\{(\frac{1}{2}, \frac{1}{2}a, \frac{1}{2}(1-a)), (\frac{1}{2}, \frac{1}{2}(1-a), \frac{1}{2}a)\}$.

If $a \in \{0, 1\}$, then these oracle weights correspond to $\mathbf{w}^{(1)}$ and $\mathbf{w}^{(2)}$ defined below Remark 8.

Track-and-Stop with D-tracking will track $\mathbf{w}^{(1)}(\hat{\mu}_t)$ or $\mathbf{w}^{(2)}(\hat{\mu}_t)$ depending on which side of the hyperplane $D_2 = D_3$ the empirical mean $\hat{\mu}_t$ lies. As in Appendix D, we do not know the distribution of the tracked vector \mathbf{w}_t , but if $\hat{\mu}_t$ crosses that boundary often enough, then as explained below Remark 8, D-tracking will get N_t/t outside of the convex hull of $\{\mathbf{w}^{(1)}(\mu), \mathbf{w}^{(2)}(\mu)\}$ and be suboptimal.

We verify experimentally that suboptimality in Figure 5. For these experiments, the Gaussians have variance $1/4$, $\mu_0 = 1/5$, $a = 1/10$. The fixed optimal sampling strategy samples $\text{argmin}_k N_{t,k} - t(\frac{1}{2}, \frac{1}{2}a, \frac{1}{2}(1-a))_k$.

F Continuity Proofs

We first introduce the necessary notions used in the modification of Berge's theorem we will apply, following [Feinberg et al., 2014].

Definition 20. For a function $f : U \rightarrow \mathbb{R}$ with U a non-empty subset of a topological space, define the level sets

$$L_f(y, U) = \{x \in U : f(x) \leq y\},$$

$$L_f^<(y, U) = \{x \in U : f(x) < y\}.$$

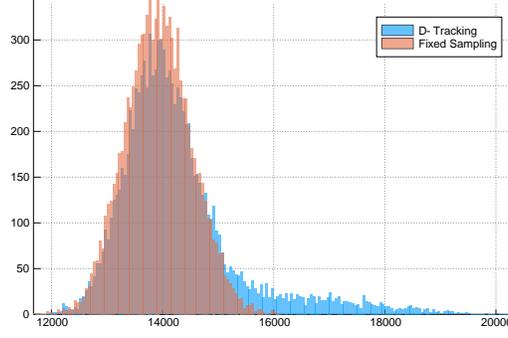


Figure 5: Histogram of the stopping time of Track-and-Stop with D-tracking and of fixed optimal sampling. The data is comprised of 10000 runs of each algorithm, with $\delta = e^{-60}$.

A function f is *lower semi-continuous* on U if all the level sets $L_f(y, U)$ are closed. It is *inf-compact* on U if all these level sets are compact. It is *upper semi-continuous* if all the strict level sets $L_f^<(y, U)$ are open.

Let \mathbb{X} and \mathbb{Y} be Hausdorff topological spaces. Let $u : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ be a function, $\Phi : \mathbb{X} \rightarrow \mathbb{S}(\mathbb{Y})$ be a set-valued function, where $\mathbb{S}(\mathbb{Y})$ is the set of non-empty subsets of \mathbb{Y} . The objects of study are

$$v(x) = \inf_{y \in \Phi(x)} u(x, y),$$

$$\Phi^*(x) = \{y \in \Phi(x) : u(x, y) = v(x)\}.$$

For $U \subset \mathbb{X}$, let the graph of Φ restricted to U be $Gr_U(\Phi) = \{(x, y) \in U \times \mathbb{Y} : y \in \Phi(x)\}$.

Definition 21. A function $u : \mathbb{X} \times \mathbb{Y} \rightarrow \overline{\mathbb{R}}$ is called \mathbb{K} -inf-compact on $Gr_{\mathbb{X}}(\Phi)$ if for all non-empty compact subset C of \mathbb{X} , u is inf-compact on $Gr_C(\Phi)$.

We will use two versions of Berge's theorem. The first one restricts Φ to be compact-valued. The second one removes that hypothesis on Φ at the price of hypotheses on u . Denote by $\mathbb{K}(\mathbb{X})$ the subset of $\mathbb{S}(\mathbb{X})$ containing non-empty compact subsets of \mathbb{X} .

Theorem 22 (Berge's theorem). *Let \mathbb{X} and \mathbb{Y} be Hausdorff topological spaces. Assume that*

- $\Phi : \mathbb{X} \rightarrow \mathbb{K}(\mathbb{X})$ is continuous (i.e. both lower hemicontinuous and upper hemicontinuous),
- $u : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ is continuous.

Then the uncton $v : \mathbb{X} \rightarrow \mathbb{R}$ is continuous and the solution multifunction $\Phi^ : \mathbb{X} \rightarrow \mathbb{S}(\mathbb{Y})$ is upper hemicontinuous and compact valued.*

Theorem 23 (Feinberg et al. 2014). *Assume that*

- \mathbb{X} is compactly generated,
- $\Phi : \mathbb{X} \rightarrow \mathbb{S}(\mathbb{Y})$ is lower hemicontinuous,
- $u : \mathbb{X} \times \mathbb{Y} \rightarrow \mathbb{R}$ is \mathbb{K} -inf-compact and upper semi-continuous on $Gr_{\mathbb{X}}(\Phi)$.

Then the function $v : \mathbb{X} \rightarrow \mathbb{R}$ is continuous and the solution multifunction $\Phi^ : \mathbb{X} \rightarrow \mathbb{S}(\mathbb{Y})$ is upper hemicontinuous and compact valued.*

Theorem 4 is cut into several successive lemma, whose proofs together prove the theorem. The first three lemmas prove the continuity of $(w, \mu) \mapsto D(w, \mu, \neg i)$, first in the case where $\neg i$ is compact, then in the general case.

Lemma 24. *Set $i \in \mathcal{I}$. If $\neg i$ is compact, then $(w, \mu) \mapsto D(w, \mu, \neg i)$ is jointly continuous on $\Delta_K \times \mathcal{M}$ and $\lambda^*(w, \mu)$ is non-empty, upper hemicontinuous and compact valued.*

Proof. We apply Theorem 22 to

- $\mathbb{X} = \Delta_K \times \mathcal{M}$,
- $\mathbb{Y} = \neg i$,
- $\Phi(\boldsymbol{\mu}) = \neg i$,
- $u((\boldsymbol{w}, \boldsymbol{\mu}), \boldsymbol{\lambda}) = D(\boldsymbol{w}, \boldsymbol{\mu}, \boldsymbol{\lambda})$.

Φ is compact-valued, non-empty and continuous (since it is constant). u is continuous. The hypotheses are verified and the theorem gives the wanted result. \square

Let $\mathcal{Q}_+ = \{\boldsymbol{w} \in \mathbb{R}^K : \forall k \in [K], w_k \geq 0\}$, $\overset{\circ}{\mathcal{Q}}_+$ be its interior and for $\varepsilon > 0$, $\mathcal{Q}_+^\varepsilon = \{\boldsymbol{w} \in \mathbb{R}^K : \forall k \in [K], w_k \geq \varepsilon\}$.

Lemma 25. *The function $(\boldsymbol{w}, \boldsymbol{\mu}) \mapsto D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ is jointly continuous on $\overset{\circ}{\mathcal{Q}}_+ \times \mathcal{M}$. On the same set, $\boldsymbol{\lambda}^*(\boldsymbol{w}, \boldsymbol{\mu})$ is upper hemicontinuous, non-empty and compact valued. In particular, the same properties hold on $\overset{\circ}{\Delta}_K \times \mathcal{M}$.*

Proof. Let $\varepsilon > 0$. We prove the result for $(\boldsymbol{w}, \boldsymbol{\mu}) \in \mathcal{Q}_+^\varepsilon \times \mathcal{O}^K$ (and note that $\mathcal{M} \subseteq \mathcal{O}^K$).

We will apply Theorem 23 to

- $\mathbb{X} = \mathcal{Q}_+^\varepsilon \times \mathcal{O}^K$,
- $\mathbb{Y} = \mathcal{O}^K$,
- $\Phi((\boldsymbol{w}, \boldsymbol{\mu})) = \neg i$,
- $u((\boldsymbol{w}, \boldsymbol{\mu}), \boldsymbol{\lambda}) = D(\boldsymbol{w}, \boldsymbol{\mu}, \boldsymbol{\lambda})$,
- $v(\boldsymbol{w}, \boldsymbol{\mu}) = D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$.

We now verify the hypothesis of the theorem. First, \mathbb{X} is compactly generated since it is a metric space. Secondly, Φ is lower hemicontinuous since it is constant.

The function u is continuous, hence upper semi-continuous. It remains to check that u is \mathbb{K} -inf-compact on $Gr_{\mathbb{X}}(\Phi)$.

Let C be a non-empty compact subset of $\mathcal{Q}_+^\varepsilon \times \mathcal{O}^K$. We need to prove that u is inf-compact on $Gr_C(\Phi) = C \times \neg i$. The level sets $L_u(y, C \times \neg i)$ for $y \in \mathbb{R}$ are closed by continuity of u . Indeed they are the reverse image of a closed set $[-\infty, y]$ by a continuous function, hence they are closed in $(\mathcal{Q}_+^\varepsilon \times \mathcal{O}^K) \times \mathcal{O}^K$. We only need to prove that they are bounded.

Set $y \in \mathbb{R}$. Let $\mu_k^+ = \sup_{(\boldsymbol{w}, \boldsymbol{\mu}) \in C} \mu_k$, finite since C is compact. Define μ_k^- in a similar way with an infimum. For $j \in [K]$, $\lim d(\mu_j, \lambda_j) = +\infty$ when λ_j approaches the boundaries of the open interval \mathcal{O} . Then for all $k \in [K]$, there exists λ_k^+ such that $\lambda > \lambda_k^+ \Rightarrow \forall (\boldsymbol{w}, \boldsymbol{\mu}) \in C, d(\mu_k, \lambda) > y/\varepsilon$. Define λ_k^- in a similar way. For $\boldsymbol{\lambda} \notin [\lambda_1^-, \lambda_1^+] \times \dots \times [\lambda_K^-, \lambda_K^+]$ there exists a $k \in [K]$ such that $d(\mu_k, \lambda_k) > y/\varepsilon$ for all $(\boldsymbol{w}, \boldsymbol{\mu}) \in C$. Since $w_k \geq \varepsilon$ for all k , for all $(\boldsymbol{w}, \boldsymbol{\mu}) \in C$, $D(\boldsymbol{w}, \boldsymbol{\mu}, \boldsymbol{\lambda}) = \sum_{k=1}^K w_k d(\mu_k, \lambda_k) > y$. The level set $L_u(y, C \times \neg i)$ is bounded.

We have verified the hypotheses of the theorem for compact subsets of $\mathcal{Q}_+^\varepsilon \times \mathcal{O}^K$ and obtain that $v(\boldsymbol{w}, \boldsymbol{\mu}) = D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ is continuous as a function of $(\boldsymbol{w}, \boldsymbol{\mu})$ on that set. On that same set, the function giving the points realizing the infimum $\boldsymbol{\lambda}^*(\boldsymbol{w}, \boldsymbol{\mu})$ is upper hemicontinuous, non-empty and compact valued. \square

For a projection P on a subset of coordinates $S \subseteq [K]$ and $\boldsymbol{w} \in \mathcal{Q}_+$, denote the projected vector by $P\boldsymbol{w}$. The same proof as the one of the previous lemma applied to the projected spaces gives the following corollary.

Corollary 26. *Let P be a projection on a subset of coordinates $S \subseteq [K]$. Then the function $(\boldsymbol{u}, \boldsymbol{\mu}) \mapsto D(\boldsymbol{u}, \boldsymbol{\mu}, P\neg i)$ is continuous on $P\overset{\circ}{\mathcal{Q}}_+ \times \mathcal{M}$.*

We now extend the continuity of $D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ on all of $\Delta_K \times \mathcal{M}$.

Lemma 27. *The function $(\boldsymbol{w}, \boldsymbol{\mu}) \mapsto D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ is continuous on $\Delta_K \times \mathcal{M}$.*

Proof. Let $(\boldsymbol{w}, \boldsymbol{\mu}) \in \Delta_K \times \mathcal{M}$ and $\boldsymbol{\lambda}^\varepsilon$ be such that $D(\boldsymbol{w}, \boldsymbol{\mu}, \boldsymbol{\lambda}) \leq D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i) + \varepsilon$, which exists by definition of $D(\boldsymbol{w}, \boldsymbol{\mu}, \neg i)$ as an infimum.

The function $(\mathbf{w}, \boldsymbol{\mu}) \mapsto D(\mathbf{w}, \boldsymbol{\mu}, \boldsymbol{\lambda}^\varepsilon)$ is continuous. Hence, there exists ε' such that for $\|\mathbf{w}' - \mathbf{w}\|_\infty \leq \varepsilon'$ and $\|\boldsymbol{\mu}' - \boldsymbol{\mu}\|_\infty \leq \varepsilon'$, $D(\mathbf{w}', \boldsymbol{\mu}', \boldsymbol{\lambda}^\varepsilon) \leq D(\mathbf{w}, \boldsymbol{\mu}, \boldsymbol{\lambda}^\varepsilon) + \varepsilon$. For $(\mathbf{w}', \boldsymbol{\mu}')$ in such a neighbourhood of $(\mathbf{w}, \boldsymbol{\mu})$,

$$D(\mathbf{w}', \boldsymbol{\mu}', \neg i) \leq D(\mathbf{w}', \boldsymbol{\mu}', \boldsymbol{\lambda}^\varepsilon) \leq D(\mathbf{w}, \boldsymbol{\mu}, \boldsymbol{\lambda}^\varepsilon) + \varepsilon \leq D(\mathbf{w}, \boldsymbol{\mu}, \neg i) + 2\varepsilon.$$

We proved that $(\mathbf{w}, \boldsymbol{\mu}) \mapsto D(\mathbf{w}, \boldsymbol{\mu}, \neg i)$ is upper semi-continuous.

Now let $\varepsilon > 0$ be such that $\min_{k: w_k > 0} w_k \geq 2\varepsilon$ and $(\mathbf{w}', \boldsymbol{\mu}')$ be in an ε neighbourhood of $(\mathbf{w}, \boldsymbol{\mu})$. Then $w_k > 0 \Rightarrow w'_k > \varepsilon$. For $\mathbf{u} \in \Delta_K$, denote by $P\mathbf{u}$ its projection on the coordinates for which $w_k > 0$.

$$\begin{aligned} D(\mathbf{w}', \boldsymbol{\mu}', \neg i) &\geq D(P\mathbf{w}', \boldsymbol{\mu}', P\neg i), \\ D(\mathbf{w}, \boldsymbol{\mu}, \neg i) &= D(P\mathbf{w}, \boldsymbol{\mu}, P\neg i). \end{aligned}$$

By the previous lemma, $(P\mathbf{w}', \boldsymbol{\mu}') \mapsto D(P\mathbf{w}', \boldsymbol{\mu}', P\neg i)$ is continuous on $PQ_+^\varepsilon \times \mathcal{M}$. Hence for $(\mathbf{w}', \boldsymbol{\mu}')$ in a small enough neighbourhood of $(\mathbf{w}, \boldsymbol{\mu})$, $D(P\mathbf{w}', \boldsymbol{\mu}', P\neg i) \geq D(P\mathbf{w}, \boldsymbol{\mu}, P\neg i) - \varepsilon$. In that neighbourhood,

$$D(\mathbf{w}', \boldsymbol{\mu}', \neg i) \geq D(P\mathbf{w}', \boldsymbol{\mu}', P\neg i) \geq D(P\mathbf{w}, \boldsymbol{\mu}, P\neg i) - \varepsilon \geq D(\mathbf{w}, \boldsymbol{\mu}, \neg i) - \varepsilon.$$

This proves the lower semi-continuity of $(\mathbf{w}, \boldsymbol{\mu}) \mapsto D(\mathbf{w}, \boldsymbol{\mu}, \neg i)$. We now have both lower and upper semi-continuity: that function is continuous. \square

We proved continuity of the function of $\Delta_K \times \mathcal{M}$ but upper hemi-continuity of $\boldsymbol{\lambda}^*$ only for $\mathbf{w} \in \overset{\circ}{\Delta}_K$. We now show an example where $\boldsymbol{\lambda}^*$ is empty at a \mathbf{w} on the boundary of the simplex.

Let $\boldsymbol{\mu} = (0, 0)$, $\neg i = \{(\sqrt{x}, \sqrt{1+1/x}) : x \in \mathbb{R}^+\}$ and $d(\mu_i, \lambda_i) = (\mu_i - \lambda_i)^2$. Then

$$\inf_{\boldsymbol{\lambda}} [w_1 d(\mu_1, \lambda_1) + w_2 d(\mu_2, \lambda_2)] = \inf_x (w_1 x + (1 - w_1)(1 + \frac{1}{x})) = 1 - w_1 + 2\sqrt{w_1(1 - w_1)},$$

with minimum attained for $w_1 \in (0, 1)$ at $x = \sqrt{\frac{1-w_1}{w_1}}$. Hence for $w_1 > 0$, $\boldsymbol{\lambda}^*(\mathbf{w}, (0, 0)) = \{(\sqrt{\frac{1-w_1}{w_1}}, \sqrt{\frac{w_1}{1-w_1}})\}$, non-empty and compact. If $w_1 = 0$ then the infimum is not attained and $\boldsymbol{\lambda}^*$ is empty.

Lemma 28. *For all $i \in \mathcal{I}$, $D(\boldsymbol{\mu}, \neg i)$ is a continuous function of $\boldsymbol{\mu}$ on \mathcal{M} and $\mathbf{w}^*(\boldsymbol{\mu}, \neg i)$ is upper hemicontinuous and compact-valued.*

Proof. We apply Theorem 22 to

- $\mathbb{X} = \mathcal{M}$,
- $\mathbb{Y} = \Delta_K$,
- $\Phi(\boldsymbol{\mu}) = \Delta_K$,
- $u(\boldsymbol{\mu}, \mathbf{w}) = D(\mathbf{w}, \boldsymbol{\mu}, \neg i)$.

Φ is compact-valued, non-empty and continuous (since it is constant). u is continuous by Lemma 27. The hypotheses are verified and the theorem gives the wanted result. \square

Lemma 29. *For all $i \in \mathcal{I}$, $D(\boldsymbol{\mu})$ is a continuous function of $\boldsymbol{\mu}$ on \mathcal{M} and $\mathbf{w}^*(\boldsymbol{\mu})$ is upper hemicontinuous and compact-valued.*

Proof. We apply Theorem 22 to

- $\mathbb{X} = \mathcal{M}$,
- $\mathbb{Y} = \Delta_K$,
- $\Phi(\boldsymbol{\mu}) = \Delta_K$,
- $u(\boldsymbol{\mu}, \mathbf{w}) = \max_{i \in \mathcal{I}} D(\mathbf{w}, \boldsymbol{\mu}, \neg i)$.

Φ is compact-valued, non-empty and continuous (since it is constant). u is continuous since it is a finite maximum of continuous functions. The hypotheses are verified and the theorem gives the wanted result. \square

Lemma 30. *$i_F(\boldsymbol{\mu})$ is upper hemicontinuous and compact valued.*

Proof. We apply Theorem 22 to

- $\mathbb{X} = \mathcal{M}$,
- $\mathbb{Y} = \{1, \dots, K\}$,
- $\Phi(\boldsymbol{\mu}) = \{1, \dots, K\}$,
- $u(\boldsymbol{\mu}, \boldsymbol{w}) = D(\boldsymbol{\mu}, -i)$.

Φ is compact-valued, non-empty and continuous (since it is constant). u is continuous Lemma 28. The hypotheses are verified and the theorem gives the wanted result. \square

G Algorithm Analysis

G.1 Proof of Lemma 14

Lemma 31 (Lemma 19 of Garivier and Kaufmann 2016). *There exists two constants B and C (that depend on μ and ξ) such that*

$$\mathbb{P}_{\boldsymbol{\mu}}(\mathcal{E}'_T{}^c) \leq BT e^{-C\sqrt{h(T)}}.$$

For $h(T) = \sqrt{T}$, $\sum_{T=1}^{+\infty} \mathbb{P}_{\boldsymbol{\mu}}(\mathcal{E}'_T{}^c)$ is then finite.

Lemma 32 (Magureanu et al. 2014). *Let $\beta(t, \delta) = \log(Ct^2/\delta)$ with C a constant verifying the inequality $C \geq e \sum_{t=1}^{\infty} (\frac{e}{K})^K \frac{(\log^2(Ct^2) \log(t))^K}{t^2}$. Then*

$$\mathbb{P}_{\boldsymbol{\mu}} \left\{ \exists t \in \mathbb{N}, \sum_{k=1}^K N_{t,k} d(\hat{\boldsymbol{\mu}}_{t,k}, \boldsymbol{\mu}_k) > \beta(t, \delta) \right\} \leq \delta.$$

Let $f(t) = \exp(\beta(t, 1/t^5)) = Ct^{10}$ in the definition of the confidence ellipsoid \mathcal{C}_t (see Algorithm 1). Then

$$\sum_{T=T_0}^{+\infty} \mathbb{P}_{\boldsymbol{\mu}}(\mathcal{E}_T^c) \leq \sum_{T=1}^{+\infty} \sum_{t=\sqrt{T}}^T \frac{1}{t^5} < +\infty.$$

G.2 Proof of Lemma 15

Set $T > T_1$ and suppose that $\mathcal{E}_T \cap \mathcal{E}'_T$ is true. For $t \geq \eta(T)$, if $\tau_{\delta} > t$ then $tC_{\varepsilon, \xi}^*(\boldsymbol{\mu}) \leq D(N_t, \hat{\boldsymbol{\mu}}_t, -i_{\boldsymbol{\mu}}) \leq \beta(t, \delta)$, hence $t \leq \beta(t, \delta)/C_{\varepsilon, \xi}^*(\boldsymbol{\mu})$.

$$\begin{aligned} \min(\tau_{\delta}, T) &\leq \lceil \eta(T) \rceil + \sum_{t=\lceil \eta(T) \rceil+1}^T \mathbb{I}\{\tau_{\delta} > t-1\} \\ &\leq \lceil \eta(T) \rceil + \sum_{t=\lceil \eta(T) \rceil+1}^T \mathbb{I}\left\{t \leq 1 + \frac{\beta(t, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}\right\} \\ &\leq \lceil \eta(T) \rceil + \sum_{t=\lceil \eta(T) \rceil+1}^T \mathbb{I}\left\{t \leq 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}\right\} \\ &\leq \max\left(\lceil \eta(T) \rceil, 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}\right). \end{aligned}$$

Suppose that $\tau_{\delta} > T$. Then $T \leq \max\left(\lceil \eta(T) \rceil, 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}\right)$. But by hypothesis, $\eta(T) < T-1$, such that that inequality implies $T \leq 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})}$. We conclude that

$$\tau_{\delta} \leq \inf \left\{ T : T > 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}^*(\boldsymbol{\mu})} \right\}.$$

G.3 Continuity Results and Proof of Lemma 6

Lemma 33. *Let $\varepsilon > 0$ and $A \subseteq \Delta_K$ be a convex set and let $\mathbf{w}_1, \dots, \mathbf{w}_t \in \Delta_K$ be such that for all $s \in [t]$, $\inf_{\mathbf{w} \in A} \|\mathbf{w}_s - \mathbf{w}\|_\infty \leq \varepsilon$. Then $\inf_{\mathbf{w} \in A} \left\| \frac{1}{t} \sum_{s=1}^t \mathbf{w}_s - \mathbf{w} \right\|_\infty \leq \varepsilon$.*

Proof. For $s \in [t]$, let $\mathbf{w}_s^* \in A$ be such that $\|\mathbf{w}_s - \mathbf{w}_s^*\|_\infty \leq \varepsilon$. Then $\frac{1}{t} \sum_{s=1}^t \mathbf{w}_s^* \in A$ by convexity and

$$\left\| \frac{1}{t} \sum_{s=1}^t \mathbf{w}_s - \frac{1}{t} \sum_{s=1}^t \mathbf{w}_s^* \right\|_\infty \leq \frac{1}{t} \sum_{s=1}^t \|\mathbf{w}_s - \mathbf{w}_s^*\|_\infty \leq \varepsilon.$$

□

Proof of Lemma 6. Let $\varepsilon > 0$. By Theorem 4, \mathbf{w}^* is upper hemicontinuous: there exists $\xi > 0$ such that if $\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi$ then for all $\mathbf{w}_t \in \mathbf{w}^*(\hat{\boldsymbol{\mu}}_t)$, $\inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu})} \|\mathbf{w}_t - \mathbf{w}\|_\infty \leq \varepsilon$.

For $\xi > 0$, there exists t_ξ such that for $t \geq t_\xi$, $\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi$ by hypothesis. Then for $\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu})$,

$$\left\| \frac{1}{t} \sum_{s=1}^t \mathbf{w}_s - \mathbf{w} \right\|_\infty \leq \frac{t_\xi}{t} + \frac{t - t_\xi}{t} \left\| \frac{1}{t - t_\xi} \sum_{s=t_\xi}^t \mathbf{w}_s - \mathbf{w} \right\|_\infty.$$

Taking infimums and using the convexity of $\mathbf{w}^*(\boldsymbol{\mu})$ to apply Lemma 33,

$$\inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu})} \left\| \frac{1}{t} \sum_{s=1}^t \mathbf{w}_s - \mathbf{w} \right\|_\infty \leq \frac{t_\xi}{t} + \varepsilon \leq 2\varepsilon \text{ for } t \geq t_\xi/\varepsilon.$$

□

G.4 Proof of Lemma 16

Under \mathcal{E}_T , for $t \geq h(T)$, $\boldsymbol{\mu} \in \mathcal{C}_t$. Hence $i_F(\boldsymbol{\mu}) \subseteq I_t$.

For $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathcal{M}$, let $ch(\boldsymbol{\mu}, \boldsymbol{\mu}') = \inf_{\lambda \in \mathbb{R}^K} \sum_{k=1}^K (d(\lambda_k, \mu_k) + d(\lambda_k, \mu'_k))$. ch is a semi-metric on \mathcal{M} .

Suppose that the event \mathcal{E}_T holds and set $t > h(T)$. Then on one hand $\boldsymbol{\mu} \in \mathcal{C}_t$, such that $D(N_{t-1}, \hat{\boldsymbol{\mu}}_{t-1}, \boldsymbol{\mu}) \leq \log f(t-1)$. On the other hand, by definition every point $\boldsymbol{\mu}' \in \mathcal{C}_t$ verifies $D(N_{t-1}, \hat{\boldsymbol{\mu}}_{t-1}, \boldsymbol{\mu}') \leq \log f(t-1)$. We obtain in particular that

$$\sum_{k=1}^K N_{t-1,k} (d(\hat{\mu}_{t-1,k}, \mu_k) + d(\hat{\mu}_{t-1,k}, \mu'_k)) \leq 2 \log f(t-1).$$

By hypotheses, $N_{t-1,k} \geq n(t-1)$. We obtain that $ch(\boldsymbol{\mu}, \boldsymbol{\mu}') \leq \frac{2 \log f(t-1)}{n(t-1)}$ for all $\boldsymbol{\mu}' \in \mathcal{C}_t$.

By upper hemicontinuity of $i_F(\boldsymbol{\mu})$, there exists $\varepsilon > 0$ such that $\|\boldsymbol{\mu} - \boldsymbol{\mu}'\|_\infty \leq \varepsilon \Rightarrow i_F(\boldsymbol{\mu}') \subseteq i_F(\boldsymbol{\mu})$. There exists $\Delta > 0$ such that $ch(\boldsymbol{\mu}, \boldsymbol{\mu}') \leq \Delta \Rightarrow i_F(\boldsymbol{\mu}') \subseteq i_F(\boldsymbol{\mu})$.

For such a Δ and $T_\Delta = \inf\{t \in \mathbb{N} : \frac{2 \log f(t)}{n(t)} \leq \Delta\}$, if $t \geq \max(h(T), T_\Delta)$, then $i_F(\boldsymbol{\mu}') \subseteq i_F(\boldsymbol{\mu})$ for all $\boldsymbol{\mu}' \in \mathcal{C}_t$. hence $I_t = \bigcup_{\boldsymbol{\mu}' \in \mathcal{C}_t} i_F(\boldsymbol{\mu}') \subseteq i_F(\boldsymbol{\mu})$.

G.5 Proof of Lemma 17

Lemma 34 (Garivier and Kaufmann 2016). *For all $t \geq 1$ and $k \in [K]$, the C-tracking rule ensures that $N_{t,k} \geq \sqrt{t + K^2} - 2K$ and that*

$$\left\| N_t - \sum_{s=0}^{t-1} \mathbf{w}_s \right\|_\infty \leq K(1 + \sqrt{t}).$$

Lemma 35. Suppose that there exists $T_I \in \mathbb{N}$ such that for $T \geq T_I$, $\mathbf{w}_t \in \mathbf{w}^*(\hat{\boldsymbol{\mu}}_t, \neg i_\mu)$. Then for T such that $h(T) \geq T_I$, it holds that on \mathcal{E}'_T C-Tracking verifies

$$\forall t \geq 4 \frac{K^2}{\varepsilon^2} + 3 \frac{h(T)}{\varepsilon}, \quad \inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \left\| \frac{N(t)}{t} - \mathbf{w} \right\|_\infty \leq 3\varepsilon.$$

Proof. Suppose that T verifies $h(T) \geq T_I$. Using Lemma 34, for $t \geq h(T)$ one can write for all $\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)$,

$$\begin{aligned} \left\| \frac{N_t}{t} - \mathbf{w} \right\|_\infty &\leq \left\| \frac{N_t}{t} - \frac{1}{t} \sum_{s=0}^{t-1} \mathbf{w}_s \right\|_\infty + \left\| \frac{1}{t} \sum_{s=0}^{t-1} \mathbf{w}_s - \mathbf{w} \right\|_\infty \\ &\leq \frac{K(1 + \sqrt{t})}{t} + \left\| \frac{1}{t} \sum_{s=0}^{t-1} \mathbf{w}_s - \mathbf{w} \right\|_\infty \\ &\leq \frac{2K}{\sqrt{t}} + \frac{h(T)}{t} + \left\| \frac{1}{t} \sum_{s=h(T)}^{t-1} (\mathbf{w}_s - \mathbf{w}) \right\|_\infty. \end{aligned}$$

The definition of event \mathcal{E}'_T uses $\xi > 0$ such that if $\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi$ then for all $\mathbf{w}_t \in \mathbf{w}^*(\hat{\boldsymbol{\mu}}_t, \neg i_\mu)$, $\inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \|\mathbf{w}_t - \mathbf{w}\|_\infty \leq \varepsilon$. Under that event, for $t \geq h(T)$,

$$\forall \mathbf{w}_t \in \mathbf{w}^*(\hat{\boldsymbol{\mu}}_t, \neg i_\mu), \quad \inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \|\mathbf{w}_t - \mathbf{w}\|_\infty \leq \varepsilon, \quad \|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi,$$

The convexity of $\mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)$ ensures that $\inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \left\| \frac{1}{t} \sum_{s=h(T)}^T \mathbf{w}_s - \mathbf{w} \right\|_\infty \leq \varepsilon$ as well by Lemma 33. Hence, taking infimums and using the hypothesis that the event \mathcal{E}'_T holds,

$$\begin{aligned} \inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \left\| \frac{N_t}{t} - \mathbf{w} \right\|_\infty &\leq \frac{2K}{\sqrt{t}} + \frac{h(T)}{t} + \inf_{\mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}, \neg i_\mu)} \left\| \frac{1}{t} \sum_{s=h(T)}^{t-1} (\mathbf{w}_s - \mathbf{w}) \right\|_\infty \\ &\leq \frac{2K}{\sqrt{t}} + \frac{h(T)}{t} + \varepsilon. \end{aligned}$$

For $t \geq 2 \frac{K^2}{\varepsilon^2} \left(1 + \varepsilon \frac{h(T)}{2K^2} + \sqrt{1 + \varepsilon \frac{h(T)}{K^2}} \right)$, the right-hand-side is smaller than 2ε . In particular, this is also true for $t \geq 4 \frac{K^2}{\varepsilon^2} + 3 \frac{h(T)}{\varepsilon}$. \square

Proof of Lemma 17. Let $T \in \mathbb{N}$ be such that $h(T) \geq T_\Delta$. Under \mathcal{E}_T , for $t \geq h(T)$ the set I_t is constant and equal to $i_F(\boldsymbol{\mu})$ by Lemma 16. For this stage on, $i_t = i_\mu$ is constant and $\mathbf{w}^*(\boldsymbol{\mu}, \neg i_t) \subseteq \mathbf{w}^*(\boldsymbol{\mu})$. We can hence take $T_I = h(T)$ in Lemma 35 and we get the wanted result. \square

G.6 Proof of the empirical complexity of Track and Stop

We prove Theorem 7 and Theorem 9.

Lemma 16 depends only on the amount of forced exploration, thus it is valid for Track and Stop. After some $T_\Delta > 0$, $I_t = i_F(\boldsymbol{\mu})$. Since $\hat{\boldsymbol{\mu}}_t \in I_t$, $i_F(\hat{\boldsymbol{\mu}}_t) \subseteq i_F(\boldsymbol{\mu})$. The alternative selected by Track and Stop to compute \mathbf{w}_t will be in $i_F(\boldsymbol{\mu})$.

We modify the event \mathcal{E}'_T used for Sticky Track and Stop into $\mathcal{E}'_T = \bigcap_{t=h(T)}^T \{\|\hat{\boldsymbol{\mu}}_t - \boldsymbol{\mu}\|_\infty \leq \xi\}$ where ξ is such that

$$\|\boldsymbol{\mu}' - \boldsymbol{\mu}\|_\infty \leq \xi \Rightarrow \forall \mathbf{w}' \in \mathbf{w}^*(\boldsymbol{\mu}') \exists \mathbf{w} \in \mathbf{w}^*(\boldsymbol{\mu}), \|\mathbf{w}' - \mathbf{w}\|_\infty \leq \varepsilon.$$

The difference is that we use the upper hemicontinuity of $\mathbf{w}^*(\boldsymbol{\mu})$ instead of $\mathbf{w}^*(\boldsymbol{\mu}, \neg i)$ for some $i \in \mathcal{I}$.

From that point on, we proceed as in the proof of the sample complexity of Sticky Track and Stop, except that N_t/t will not necessarily converge to $\mathbf{w}^*(\boldsymbol{\mu}, i_{\boldsymbol{\mu}})$, but to $\text{conv}(\mathbf{w}^*(\boldsymbol{\mu}))$, convex hull of $\mathbf{w}^*(\boldsymbol{\mu})$. Lemma 17 is true for Track and Stop with the adapted \mathcal{E}'_T if $\mathbf{w}^*(\boldsymbol{\mu}, i_{\boldsymbol{\mu}})$ is replaced by $\text{conv}(\mathbf{w}^*(\boldsymbol{\mu}))$. The analogue of $C_{\varepsilon, \xi}^*(\boldsymbol{\mu})$ is

$$C_{\varepsilon, \xi}(\boldsymbol{\mu}) = \inf_{\substack{\boldsymbol{\mu}': \|\boldsymbol{\mu}' - \boldsymbol{\mu}\|_{\infty} \leq \xi \\ \mathbf{w}': \inf_{\mathbf{w} \in \text{conv}(\mathbf{w}^*(\boldsymbol{\mu}))} \|\mathbf{w}' - \mathbf{w}\|_{\infty} \leq 3\varepsilon}} D(\mathbf{w}', \boldsymbol{\mu}').$$

Let T_{Δ} be defined as in Lemma 16. Let T be such that $h(T) \geq T_{\Delta}$. Let $\eta(T) = 4\frac{K^2}{\varepsilon^2} + 3\frac{h(T)}{\varepsilon}$. By Lemma 17 changed as explained, for $t \geq \eta(T)$, if $\mathcal{E}_T \cap \mathcal{E}'_T$ then $D(N_t, \hat{\boldsymbol{\mu}}_t) \geq tC_{\varepsilon, \xi}(\boldsymbol{\mu})$.

We now apply Lemma 15. $\eta(T) < T - 1$ if $h(T) < \frac{\varepsilon}{3}(T - 1) - \frac{4}{3}\frac{K^2}{\varepsilon}$. For $h(T) = \sqrt{T}$ and T bigger than a constant T_{η} depending on K and ε , this is true. Then under $\mathcal{E}_T \cap \mathcal{E}'_T$, the hypotheses of Lemma 15 are verified with $T_1 = h^{-1}(\max(T_{\Delta}, T_{\eta}))$.

We obtain that the hypotheses of Lemma 13 are verified for

$$T_0 = \max(T_1, \inf\{T : 1 + \frac{\beta(T, \delta)}{C_{\varepsilon, \xi}(\boldsymbol{\mu})} \leq T\}).$$

Note that $\lim_{\delta \rightarrow 0} T_0 / \log(1/\delta) = 1/C_{\varepsilon, \xi}(\boldsymbol{\mu})$. Taking $\varepsilon \rightarrow 0$ (hence $\xi \rightarrow 0$ as well), we obtain $\lim_{\delta \rightarrow 0} \frac{\mathbb{E}_{\boldsymbol{\mu}}[\tau_{\delta}]}{\log(1/\delta)} = \frac{1}{\lim_{\varepsilon \rightarrow 0} C_{\varepsilon, \xi}(\boldsymbol{\mu})}$.

Finally, $\lim_{\delta \rightarrow 0} C_{\varepsilon, \xi}(\boldsymbol{\mu}) = \inf_{\mathbf{w} \in \text{conv}(\mathbf{w}^*(\boldsymbol{\mu}))} D(\mathbf{w}, \boldsymbol{\mu})$. This proves Theorem 9.

If $i_F(\boldsymbol{\mu})$ is a singleton, then $\mathbf{w}^*(\boldsymbol{\mu})$ is convex and $\text{conv}(\mathbf{w}^*(\boldsymbol{\mu})) = \mathbf{w}^*(\boldsymbol{\mu})$, leading to the observation that $\inf_{\mathbf{w} \in \text{conv}(\mathbf{w}^*(\boldsymbol{\mu}))} D(\mathbf{w}, \boldsymbol{\mu}) = D(\boldsymbol{\mu})$. In that case, Track-and-Stop is asymptotically optimal: Theorem 7 is proved.

H Divergences

An important building block in pure exploration algorithms is the largest weighted distance from $\boldsymbol{\mu}$ to the closest point $\boldsymbol{\lambda}$ in some set of alternatives,

$$D(\boldsymbol{\mu}, \Lambda) = \sup_{\mathbf{w} \in \Delta} \inf_{\boldsymbol{\lambda} \in \Lambda} \sum_k w_k d(\mu_k, \lambda_k)$$

In this section we compute a few of these distances in closed form to get a feeling for their behaviour. We do it for the Gaussian divergence $d(\mu, \lambda) = \frac{1}{2}(\mu - \lambda)^2$.

H.1 Hyper-planes and Half-spaces

Lemma 36. *When $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^K \mid \langle \mathbf{a}, \boldsymbol{\lambda} \rangle = b\}$ is a hyper-plane, we find*

$$D(\boldsymbol{\mu}, \Lambda) = \frac{1}{2} \left(\frac{\langle \mathbf{a}, \boldsymbol{\mu} \rangle - b}{\sum_{i=1}^d |a_i|} \right)^2 \quad \text{and} \quad w_i^*(\boldsymbol{\mu}) = \frac{|a_i|}{\sum_{i=1}^d |a_i|}.$$

Note that the same result holds for the half-space $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^K \mid \langle \mathbf{a}, \boldsymbol{\lambda} \rangle \geq b\}$ when $\boldsymbol{\mu} \notin \Lambda$, i.e. $\langle \mathbf{a}, \boldsymbol{\mu} \rangle < b$. If $\boldsymbol{\mu} \in \Lambda$ then $D(\boldsymbol{\mu}, \Lambda) = 0$. The Lemma implies in particular that for Best Arm Identification with $K = 2$ arms, corresponding to $\mathbf{a} = (-1, +1)$ and $b = 0$, the optimal weights \mathbf{w}^* are uniform, as was shown by Kaufmann et al. [2016]. For the ε -BAI variant of the problem we set $b = \pm\varepsilon$, so here \mathbf{w}^* is also uniform.

Proof. We have

$$D(\boldsymbol{\mu}, \Lambda) = \sup_{\mathbf{w} \in \Delta} \inf_{\boldsymbol{\lambda}: \langle \mathbf{a}, \boldsymbol{\lambda} \rangle = b} \sum_{i=1}^d w_i \frac{(\mu_i - \lambda_i)^2}{2}$$

Introducing Lagrange multiplier θ , we find

$$\sup_{\mathbf{w} \in \Delta} \sup_{\theta} \inf_{\boldsymbol{\lambda} \in \mathbb{R}^d} \sum_{i=1}^d w_i \frac{(\mu_i - \lambda_i)^2}{2} - \theta (\langle \mathbf{a}, \boldsymbol{\lambda} \rangle - b)$$

Plugging in the solution $\lambda_i = \mu_i + \frac{\theta a_i}{w_i}$ results in

$$\sup_{\mathbf{w} \in \Delta} \sup_{\theta} -\theta^2 \sum_{i=1}^d \frac{a_i^2}{2w_i} - \theta (\langle \mathbf{a}, \boldsymbol{\mu} \rangle - b)$$

Now solving for θ results in $\theta = \frac{-\langle \mathbf{a}, \boldsymbol{\mu} \rangle - b}{\sum_{i=1}^d \frac{a_i^2}{w_i}}$ and objective function value

$$\sup_{\mathbf{w} \in \Delta} \frac{(\langle \mathbf{a}, \boldsymbol{\mu} \rangle - b)^2}{2 \sum_{i=1}^d \frac{a_i^2}{w_i}}$$

Further solving for \mathbf{w} tells us that $w_i = \frac{|a_i|}{\sum_{i=1}^d |a_i|}$ and hence the value is

$$\frac{1}{2} \left(\frac{\langle \mathbf{a}, \boldsymbol{\mu} \rangle - b}{\sum_{i=1}^d |a_i|} \right)^2.$$

□

H.2 Minimum Threshold

The following two lemmas appear as [Kaufmann et al., 2018, Lemma 1] for general divergences $d(\mu, \lambda)$.

Lemma 37. *Let $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^K \mid \min_k \lambda_k \leq \gamma\}$. Then when $\boldsymbol{\mu} \notin \Lambda$,*

$$D(\boldsymbol{\mu}, \Lambda) = \frac{1}{\sum_k \frac{1}{d(\mu_k, \gamma)}} \quad \text{where} \quad w_k^*(\boldsymbol{\mu}) = \frac{1}{\sum_j \frac{1}{d(\mu_j, \gamma)}}$$

Lemma 38. *Let $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^K \mid \min_k \lambda_k \geq \gamma\}$. Then when $\boldsymbol{\mu} \notin \Lambda$,*

$$D(\boldsymbol{\mu}, \Lambda) = d(\min_k \mu_k, \gamma) \quad \text{where} \quad \mathbf{w}^*(\boldsymbol{\mu}) = \mathbf{1}_{k=\operatorname{argmin}_j \mu_j}.$$

For the version of the problem with slack ϵ , we can simply replace γ by the appropriate $\gamma \pm \epsilon$.

H.3 Sphere

We now consider the distance to the sphere both from within and from the outside.

Lemma 39. *Let $\Lambda = \{\boldsymbol{\lambda} \in \mathbb{R}^K \mid \|\boldsymbol{\lambda}\| = 1\}$. Consider any $\boldsymbol{\mu} \in \mathbb{R}^K$. Then*

$$D(\boldsymbol{\mu}, \Lambda) = \frac{1}{2K^2} \left(\sqrt{K(1 - \|\boldsymbol{\mu}\|^2) + \|\boldsymbol{\mu}\|_1^2} - \|\boldsymbol{\mu}\|_1 \right)^2$$

$$\text{and} \quad w_k^*(\boldsymbol{\mu}) = \frac{1}{K} + \frac{|\mu_k| - \frac{1}{K} \|\boldsymbol{\mu}\|_1}{\sqrt{K(1 - \|\boldsymbol{\mu}\|^2) + \|\boldsymbol{\mu}\|_1^2}},$$

provided that

$$\left(\|\boldsymbol{\mu}\|_1 - K \min_k |\mu_k| \right)^2 \leq \|\boldsymbol{\mu}\|_1^2 - K(\|\boldsymbol{\mu}\|^2 - 1) \quad (5)$$

Note that the proviso is always satisfied when $\|\boldsymbol{\mu}\| \leq 1$. When $\|\boldsymbol{\mu}\| > 1$ it depends. When the proviso is not satisfied, boundary conditions are active. In that case a pairwise swapping argument shows that $w_k^*(\boldsymbol{\mu}) = 0$ for the k of minimal $|\mu_k|$. The rest of the solution is found by removing k , and solving the remaining problem of size $K - 1$.

Proof. We need to find

$$D(\boldsymbol{\mu}, \Lambda) = \max_{\boldsymbol{w} \in \Delta} \min_{\boldsymbol{\lambda}: \|\boldsymbol{\lambda}\|=1} \frac{1}{2} \sum_k w_k (\mu_k - \lambda_k)^2$$

As strong duality holds for the inner problem [Boyd and Vandenberghe, 2004, Appendix B], we may introduce a Lagrange multiplier θ for the constraint, and write

$$\max_{\boldsymbol{w} \in \Delta, \theta} \min_{\boldsymbol{\lambda} \in \mathbb{R}^K} \frac{1}{2} \sum_k w_k (\mu_k - \lambda_k)^2 + \frac{\theta}{2} (1 - \|\boldsymbol{\lambda}\|^2).$$

The innermost problem is unbounded in $\boldsymbol{\lambda}$ unless $\min_k w_k \geq \theta$, so we add this as an outer constraint. Then the minimiser is found at $\lambda_k = \frac{\mu_k}{1 - \frac{\theta}{w_k}}$, and by substituting that in, the problem simplifies to

$$\max_{\substack{\boldsymbol{w} \in \Delta \\ \theta \leq \min_k w_k}} -\frac{1}{2} \sum_k \frac{\mu_k^2}{\frac{1}{\theta} - \frac{1}{w_k}} + \frac{\theta}{2} = \max_{\substack{\boldsymbol{w} \in \Delta \\ \theta \leq \min_k w_k}} -\frac{1}{2} \sum_k \frac{\mu_k^2 \theta^2}{w_k - \theta} + \frac{\theta}{2} (1 - \|\boldsymbol{\mu}\|^2).$$

As a point of interpretation, note that we will find $\theta > 0$ when $\|\boldsymbol{\mu}\| < 1$, and $\theta < 0$ for $\|\boldsymbol{\mu}\| > 1$. Next we solve for \boldsymbol{w} , enforcing unit sum (but delaying non-negativity). With Lagrange multiplier c , we need to have

$$c = \frac{1}{2} \frac{\mu_k^2}{\left(\frac{w_k}{\theta} - 1\right)^2} \quad \text{resulting in} \quad w_k = \theta \left(1 + \sqrt{\frac{\mu_k^2}{2c}}\right).$$

Solving for the normalisation results in

$$c = \frac{1}{2} \left(\frac{\theta}{1 - \theta K \|\boldsymbol{\mu}\|_1}\right)^2 \quad \text{whence} \quad w_k = \theta + (1 - \theta K) \frac{|\mu_k|}{\|\boldsymbol{\mu}\|_1}.$$

Plugging this in, it remains to solve

$$\max_{\theta} -\frac{\theta^2}{2(1 - \theta K)} \|\boldsymbol{\mu}\|_1^2 + \frac{\theta}{2} (1 - \|\boldsymbol{\mu}\|^2).$$

This concave problem is bounded by non-negativity of the right-hand side of (5). Cancelling the derivative results in a quadratic equation, with the single feasible solution

$$\theta = \frac{1}{K} \left(1 - \frac{\|\boldsymbol{\mu}\|_1}{\sqrt{K(1 - \|\boldsymbol{\mu}\|^2) + \|\boldsymbol{\mu}\|_1^2}}\right).$$

Filling this in yields the value and weights of the Lemma. Finally, we need to check for negativity in the weights. Using the above expression for the weights, have $\min_k w_k \geq 0$ if

$$-\frac{1}{\frac{\|\boldsymbol{\mu}\|_1}{\min_k |\mu_k|} - K} \leq \theta \quad \text{i.e.} \quad \frac{\|\boldsymbol{\mu}\|_1}{\sqrt{\|\boldsymbol{\mu}\|_1^2 - K(\|\boldsymbol{\mu}\|^2 - 1)}} \leq \frac{1}{1 - K \frac{\min_k |\mu_k|}{\|\boldsymbol{\mu}\|_1}}$$

which we can further reorganise to (5), as required. \square

H.4 Composition of two independent problems

We consider the case where we seek to answer two independent queries on disjoint sets of arms. Let A, B be a partition of $[K]$. Suppose that the structure of the problem and the answers decompose according to this partition, i.e. $\mathcal{M} = \mathcal{M}^A \times \mathcal{M}^B$, $\mathcal{I} = \mathcal{I}^A \times \mathcal{I}^B$ and $i^*(\boldsymbol{\mu}) = i_A^*(\boldsymbol{\mu}^A) \times i_B^*(\boldsymbol{\mu}^B)$. Then we can also write all alternative sets $\neg i$ as $\neg i_A \times \mathcal{I}^B \cup \mathcal{I}^A \times \neg i_B$. It then holds that for all $\boldsymbol{w} \in \boldsymbol{w}^*(\boldsymbol{\mu}, \neg i)$,

$$\begin{aligned} \sum_{k \in A} w_k &= \frac{\frac{1}{D(\boldsymbol{\mu}^A, \neg i_A)}}{\frac{1}{D(\boldsymbol{\mu}^A, \neg i_A)} + \frac{1}{D(\boldsymbol{\mu}^B, \neg i_B)}}, \\ \sum_{k \in B} w_k &= \frac{\frac{1}{D(\boldsymbol{\mu}^B, \neg i_B)}}{\frac{1}{D(\boldsymbol{\mu}^A, \neg i_A)} + \frac{1}{D(\boldsymbol{\mu}^B, \neg i_B)}}, \\ \frac{1}{D(\boldsymbol{\mu}, \neg i)} &= \frac{1}{D(\boldsymbol{\mu}^A, \neg i_A)} + \frac{1}{D(\boldsymbol{\mu}^B, \neg i_B)}. \end{aligned}$$

Since the sample complexity is proportional to $1/D$ we obtain the natural conclusion that the number of samples needed to solve two independent queries is the sum of the samples needed by each query.