
A/B/n Testing with Control in the Presence of Subpopulations

Yoan Russac
CNRS, Inria, ENS
Université PSL
yoan.russac@ens.fr

Christina Katsimerou
Booking.com
christina.katsimerou@booking.com

Dennis Bohle
Booking.com
dennis.bohle@booking.com

Olivier Cappé
CNRS, Inria, ENS
Université PSL
olivier.cappe@cnrs.fr

Aurélien Garivier
UMPA, CNRS
Inria, ENS Lyon
aurelien.garivier@ens-lyon.fr

Wouter M. Koolen
Centrum Wiskunde & Informatica
wmkoolen@cwi.nl

Abstract

Motivated by A/B/n testing applications, we consider a finite set of distributions (called *arms*), one of which is treated as a *control*. We assume that the population is stratified into homogeneous subpopulations. At every time step, a subpopulation is sampled and an arm is chosen: the resulting observation is an independent draw from the arm conditioned on the subpopulation. The quality of each arm is assessed through a weighted combination of its subpopulation means. We propose a strategy for sequentially choosing one arm per time step so as to discover as fast as possible which arms, if any, have higher weighted expectation than the control. This strategy is shown to be asymptotically optimal in the following sense: if τ_δ is the first time when the strategy ensures that it is able to output the correct answer with probability at least $1 - \delta$, then $\mathbb{E}[\tau_\delta]$ grows linearly with $\log(1/\delta)$ at the exact optimal rate. This rate is identified in the paper in three different settings: (1) when the experimenter does not observe the subpopulation information, (2) when the subpopulation of each sample is observed but not chosen, and (3) when the experimenter can select the subpopulation from which each response is sampled. We illustrate the efficiency of the proposed strategy with numerical simulations on synthetic and real data collected from an A/B/n experiment.

1 Introduction

A/B/n testing is a website optimization procedure where multiple versions of the content (called "arms" below) are compared, often in order to find the one with the highest conversion rate. However, many e-commerce companies use A/B/n testing not only to deploy the best product implementation, but primarily to draw post-experiment inferences [11]. The decision-making involves, besides experiment results, factors such as the cost of scaling-up a solution, external data, or whether the implementation fits in a broader theme. In this setting, each of the arms better than the default product (which we will refer to as the "control" arm) is a contender for being deployed and the interest is not only in the best arm.

Given the control and $K \geq 1$ alternative implementations (variants), the simplest idea is to distribute the traffic uniformly among the arms; the arms that appear to be significantly better than the control at the end of the experiment are considered for deployment. While well-established, this process can be inefficient in terms of resources. Some alternatives are soon obviously worse (or better) than the control and would require fewer samples than the alternatives closer to the control. A second related shortcoming of the basic A/B/n testing approach is that setting the duration of the experiment –when done in advance– necessitates a very conservative approach by choosing a run-length that is sufficiently long to differentiate even the smallest possible changes.

To address these limitations, we consider in this work sequential testing policies that can both adjust the allocation of the samples and be stopped adaptively, in light of the data gathered during the experiment. In the terminology of multi-armed bandits, this corresponds to *pure exploration* problems (see, e.g., Chap. 33 of [15]). A pure exploration strategy will typically choose every minute (say), an allocation of traffic that favors arms for which the uncertainty is the highest. The experiment is stopped as soon as the significance is considered sufficient for every arm. Approaches have been developed in [8, 12, 10] for the identification of the single arm with the highest mean, a task called the Best Arm Identification (BAI) problem. In particular, [10] propose a strategy that is asymptotically optimal in the *fixed confidence setting*, meaning that, given a risk parameter δ , it finds the best arm with probability at least $1 - \delta$, using an expected number of samples that is hardly improvable when δ is small. Later, [18] incorporated the special role of the control arm in BAI and proposed an algorithm that declares as winning arm the one with the highest mean only if it is significantly better than the control. In this paper, we propose a solution to the problem of identifying all the arms that are better than the control, in a framework that generalizes the fixed confidence setting. In order to provide useful tools for practical A/B/n testing, we address two additional issues.

First, traditional stochastic bandit models are based on the assumption that the arm samples are i.i.d., whereas real world data streams usually show trends or some form of inhomogeneity. A particular case of interest for website optimization are the seasonal patterns caused by time-of-day or day-of-week variations. We henceforth include in our model observed covariates (e.g. the time of the day, but possibly also the country of origin, or controlled covariates like the order in which partners appear on the page, etc.) that stratify the observations into homogeneous subpopulations. We study different scenarios, depending on how much interaction is possible with these subpopulations. We provide a sample complexity analysis and an efficient algorithm in each case. In particular, we will show that using the subpopulation information efficiently can provide significant speedups of the decision-making. In the following, we will refer to the task of identifying the set of Arms that are Better than the Control in the presence of Subpopulations as the ABC-S problem.

Second, the practice of A/B/n testing often differs from a pure sequential experiment in that the experimenter cannot always fix a risk δ at the beginning and passively wait for the stopping time of the experiment without any time limitation. To address this issue, [11] proposed to define some notion of sequential "p-values" that can be monitored as the experiment progresses and used to terminate it. This notion was further used in the BAI setting in [18]. In this contribution, we elaborate on this idea by sequentially updating a suggested solution to the ABC-S problem *together with* a risk assessment for this suggestion. We show that, for any stopping time, the probability that the suggested solution is incorrect is indeed lower than the risk assessment. When the stopping time is selected as in usual fixed-confidence pure exploration, we recover the exact same guarantees but this view of the problem also provides useful results, for instance, if the experiment needs to be terminated prematurely.

Related work. Pure exploration strategies have been studied in various settings: the identification of the best arm [8, 10], the identification of the top m arms [3, 12, 9] identifying the arms that are better than a threshold [16, 4], or identifying all ϵ -good arms [17]. As far as we know, this paper is the first to consider the problem of identifying all the arms better than a control. It is also the first to consider subpopulations in pure exploration tasks. While motivated by the example of online companies, we believe that the proposed algorithms are relevant to other domains where randomized controlled trials are used for learning. An example could be clinical trials: one may wish to identify all the alternative treatments that work better than some reference medical treatment. This would permit to choose among them taking into account different characteristics (some could be cheaper, using another molecule for avoiding allergy, etc.).

Close to the notion of the control is the notion of threshold. Locatelli et al. [16] propose an algorithm for identifying all arms above a given threshold. Their algorithm samples according to the significance

of a statistical test, and shares some similarities with the present article in the Gaussian case; however, the perspective is rather different: the authors consider the *fixed-budget setting*: the total number of samples is fixed, and the goal is then to minimize the probability of returning a wrong answer at the end. Here, the index of the control arm is known but its probability distribution is not.

In our work, the quality of the different arms is assessed with a weighted combination of its subpopulations means. Minimizing the estimation error of a convex combination of means through adaptive sampling was considered in [2] with the introduction of a *stratified estimator* that will naturally appear in our analysis.

The paper is organized as follows. In Section 2, we present the mathematical model and study the information-theoretic complexity of the problem, extending the lower bound of [10] to the ABC-S setting. We show how the complexity of the problem depends on the degree of interaction that one has with the subpopulations, introducing different modes of interaction to be defined in Figure 1 below. We also consider in detail the Gaussian case which gives rise to more interpretable results. Section 3 describes how to implement the proposed strategy, which involves the numerical resolution of non-trivial optimization problems. Finally, we provide the results of numerical experiments on synthetic and real data sets in Section 4.

2 The complexity of the ABC-S problem

2.1 Mathematical framework

A problem instance consists of the following ingredients. Known to the learner are the number of arms $K \geq 1$ in addition to the designated control arm 0, the number of subpopulations J (a standard bandit being $J = 1$), and the vector $\beta \in \mathbb{R}^J$ representing the relative importance of the subpopulations for the learning objective. We further make the stochastic assumption that samples from each arm a (including the control) and subpopulation i are drawn i.i.d. from an unknown probability distribution $\nu_{a,i}$ on \mathbb{R} , whose mean we will denote by $\mu_{a,i}$. The quality of arm a is $\mu_a := \sum_{i=1}^J \beta_i \mu_{a,i}$ the combination of the means of the arms in the different populations. For $\beta \in \mathbb{R}^J$ we define the ABC-S problem as the correct identification of the set

$$\mathcal{S}_\beta(\boldsymbol{\mu}) := \left\{ a \in [K] \mid \sum_{i=1}^J \beta_i \mu_{a,i} > \sum_{i=1}^J \beta_i \mu_{0,i} \right\}.$$

At every time step t , the algorithm selects an arm A_t based on previous choices and outcomes and observes or selects (except when explicitly specified) the population type I_t . Upon the selection of the arm A_t a reward X_t is obtained. This defines a sigma-field generated by the observations up to time t denoted $\mathcal{F}_t = \sigma(I_1, X_1, \dots, I_t, X_t)$. The number of times arm a was selected for subpopulation i at time t is denoted $N_{a,i}(t) := \sum_{s=1}^t \mathbb{1}(A_s = a, I_s = i)$ and the number of draws of arm a , $N_a(t) := \sum_{s=1}^t \mathbb{1}(A_s = a)$. We define the gap with the control arm and arm a , $\Delta_a := \mu_0 - \mu_a$.

Modes of interaction We consider four modes of interaction of the learner with the bandit, as specified in Figure 1 below. In any of the three passive modes of interaction (described in Figures 1b to 1d), we assume that the subpopulation i represents a known proportion α_i of the total population, and hence that the sequence of subpopulations is drawn i.i.d. from the fixed and discrete distribution $I_t \sim \boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_J)$ with $\boldsymbol{\alpha} \in \Sigma_J := \{x \in [0, 1]^J \mid \sum_i x_i = 1\}$ the J -dimensional simplex. Here $\boldsymbol{\alpha}$ is an exogenous parameter and can differ from β which is inherent to the learning objective and is also assumed to be known. Although it is most natural in many applications to consider that $\beta = \boldsymbol{\alpha}$ (it is even necessary in the oblivious mode to make the estimation of the μ_a 's feasible), Example 1 below describes a concrete scenario in which β has negative components.

The distributions $(\nu_{a,i})_{a,i}$ are assumed to belong the same one-parameter exponential family, $\mathcal{P} := \{(\nu_\theta)_\theta : d\nu_\theta/d\xi = \exp(\theta x - b(\theta))\}$, with ξ a reference measure on \mathbb{R} and $b : \Theta \subset \mathbb{R} \mapsto \mathbb{R}$. Every probability distribution ν_θ in \mathcal{P} is entirely defined by its mean $\dot{b}(\theta)$ [1]. We may hence identify any bandit instance with its matrix of means $\boldsymbol{\mu} \in \mathbb{R}^{(K+1) \times J}$. In addition, the Kullback-Leibler divergence between two distributions ν_θ and $\nu_{\theta'}$ in \mathcal{P} may be written in the following Bregman form:

$$d(\boldsymbol{\mu}, \boldsymbol{\mu}') = \text{KL}(\nu_\theta, \nu_{\theta'}) = b(\theta') - b(\theta) - \dot{b}(\theta)(\theta' - \theta),$$

	1. See $I_t \sim \alpha$	1. Pick A_t	1. Pick A_t
1. Pick A_t and I_t	2. Pick A_t	2. See $I_t \sim \alpha$	2. Do not see $I_t \sim \alpha$
2. See $X_t \sim \nu_{A_t, I_t}$	3. See $X_t \sim \nu_{A_t, I_t}$	3. See $X_t \sim \nu_{A_t, I_t}$	3. See $X_t \sim \nu_{A_t, I_t}$
(a) <i>Active mode</i>	(b) <i>Proportional mode</i>	(c) <i>Agnostic mode</i>	(d) <i>Oblivious mode.</i>

Figure 1: Modes of Interaction between Learner and Bandit in each round. In Active mode the learner determines the subpopulation, while in the right three passive modes it is sampled from α .

where $\mu = \dot{b}(\theta)$ and $\mu' = \dot{b}(\theta')$ correspond to the means of the two distributions ν_θ and $\nu_{\theta'}$. We also use the notation $\text{kl}(p, q)$ to denote the KL divergence of two Bernoulli distributions of parameter p and q .

We define $\mathcal{L} := \{\mu : \forall a \in [K] \cup \{0\}, \forall i \in [J], \nu_{a,i} \in \mathcal{P} \text{ and } \mu_0 \neq \mu_a\}$ the set of identifiable instances where no arm has the same weighted mean as the control. At every time step, the policies we consider output a risk assessment $\hat{\delta}_t$ together with a recommendation \hat{S}_t . We focus on *safely calibrated* policies, that are defined as satisfying the following property

$$\forall \mu \in \mathcal{L}, \forall \delta \in (0, 1), \quad \mathbb{P}_\mu \left(\exists t \geq 1 : \hat{S}_t \neq \mathcal{S}_\beta(\mu) \cap \hat{\delta}_t \leq \delta \right) \leq \delta. \quad (1)$$

Finally, when fixing a level of risk δ , we consider the stopping time associated to the filtration \mathcal{F}_t , $\tau_\delta = \inf\{t \geq 0, \hat{\delta}_t \leq \delta\}$. The objective is then to minimize the expected number of rounds necessary to obtain a level of risk of at most δ . Contrary to usual δ -PAC algorithms if stopped before τ_δ , the strategy still provides guarantees on the output set following Equation 1. In particular, safely calibrated policies have a sampling rule that does not depend on any pre-specified δ , and as such they are δ -PAC for any δ .

Example 1 ([Largest Profit Identification problem 13, p24]). *Consider a company choosing among K product designs the model to mass produce. Each candidate design k has an (equilibrium) sales price $\mu_{k,1}$ and production cost $\mu_{k,2}$. The goal is to find the model k with the largest profit $\mu_{k,1} - \mu_{k,2}$. Prices and costs are currently unknown, but can be adaptively sampled. Sampling the "price" subpopulation $i = 1$ is typically implemented by performing user preference studies, taking questionnaires, etc. Samples from the "cost" subpopulation $i = 2$ involve rating manufacturing facilities, forecasting material and labor costs etc. This problem is interesting both in the BAI and ABC objectives. The importance vector is here $\beta = (1, -1)$ and α has to be set by the learner.*

2.2 General form of the sample complexity

Depending on the mode of interaction from Figure 1, the learner has a set of sampling constraints to satisfy, here denoted \mathcal{C} and precisely defined in the next section. We define $\text{Alt}(\mu)$, the different problem instances where the set of arms better than the control differs from that of the instance μ . Formally, $\text{Alt}_\beta(\mu) := \{\lambda \in \mathcal{L} \mid \mathcal{S}_\beta(\lambda) \neq \mathcal{S}_\beta(\mu)\}$. This allows us to bound the sample complexity.

Theorem 1. *Let $\delta \in (0, 1)$ and $\beta \in \mathbb{R}^J$. For any strategy satisfying Equation 1 and any $\mu \in \mathcal{L}$, the expected number of rounds for the ABC-S problem for the agnostic, proportional and active mode satisfies:*

$$\mathbb{E}_\mu[\tau_\delta] \geq T^*(\mu) \text{kl}(\delta, 1 - \delta) \quad \text{and} \quad \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T^*(\mu). \quad (2)$$

where (recalling that $\lambda_a = \sum_{i=1}^J \beta_i \lambda_{a,i}$)

$$T^*(\mu)^{-1} = \sup_{w \in \mathcal{C}} \inf_{\lambda \in \text{Alt}_\beta(\mu)} \sum_{a=0}^K \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \quad (3)$$

$$= \sup_{w \in \mathcal{C}} \min_{b \neq 0} \inf_{\lambda \in \mathcal{L} : \lambda_0 = \lambda_b} \sum_{a \in \{0, b\}} \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}). \quad (4)$$

This result is established in Appendix A. T^* characterizes the difficulty of the learning problem.

2.3 Influence of the mode of interaction

We consider the four different modes governing the sampling rule as outlined in Figure 1. In the *agnostic* mode (Fig. 1c) an arm is first selected, after which the subpopulation type is observed. Mathematically, this brings the equality $\mathbb{E}_\mu[N_{a,i}(T)] = \alpha_i \mathbb{E}_\mu[N_a(T)]$ established in Lemma 2 and the independence constraint on the weights $\mathbf{w} \in \mathcal{C}_{\text{agnostic}} := \{w_{a,i} = \alpha_i u_a : (u_0, \dots, u_K) \in \Sigma_{K+1}\}$.

In the *proportional* mode (Fig. 1b), A_t is chosen based on \mathcal{F}_{t-1} and the current subpopulation I_t . Here, the constraint is that the total number of pulls of the different arms in the subpopulation i should respect the frequency of this subpopulation, i.e. $\sum_a \mathbb{E}_\mu[N_{a,i}(T)] = \alpha_i T$. This induces a marginal constraint on the weights of the form $\mathbf{w} \in \mathcal{C}_{\text{prop}} := \{\mathbf{w} \in \Sigma_{(K+1)J} \mid \forall i \leq J, \sum_a w_{a,i} = \alpha_i\}$. This result is established in Lemma 3 reported in Appendix B.

In the *active* mode (Fig. 1a), the learner has an additional degree of freedom— she can ask for any subpopulation type at any round. In that case, $\mathbf{w} \in \mathcal{C}_{\text{active}} := \Sigma_{(K+1)J}$ is unconstrained.

By remarking that $\mathcal{C}_{\text{agnostic}} \subset \mathcal{C}_{\text{prop}} \subset \mathcal{C}_{\text{active}}$, and given the optimization program (3) solved to obtain the characteristic time, one immediately gets

$$\forall \boldsymbol{\mu} \in \mathcal{L}, \quad T_{\text{active}}^*(\boldsymbol{\mu}) \leq T_{\text{proportional}}^*(\boldsymbol{\mu}) \leq T_{\text{agnostic}}^*(\boldsymbol{\mu}). \quad (5)$$

Hence, as expected, the more control/information on the subpopulation the learner has, the faster she is able to identify the set of arms that are better than the control.

To compare with the *oblivious* mode, in which the subpopulation information is not even observed, we have to assume that $\boldsymbol{\alpha} = \boldsymbol{\beta}$. In that case, the arm rewards follow a mixture distribution: $X_t | A_t = a \sim \sum_{i=1}^J \alpha_i \nu_{a,i}$. In Proposition 4 reported in Appendix B.3, we properly define the characteristic time of an oblivious safely calibrated policy and prove that the joint convexity of Kullback-Leibler divergences implies that it is larger than its agnostic counterpart. This completes the picture of the ordering of the characteristic times by showing that, when $\boldsymbol{\alpha} = \boldsymbol{\beta}$,

$$\forall \boldsymbol{\mu} \in \mathcal{L}, \quad T_{\text{active}}^*(\boldsymbol{\mu}) \leq T_{\text{proportional}}^*(\boldsymbol{\mu}) \leq T_{\text{agnostic}}^*(\boldsymbol{\mu}) \leq T_{\text{oblivious}}^*(\boldsymbol{\mu}). \quad (6)$$

Note that although we provide, in Section 3, algorithms to numerically compute the first three complexities, evaluating $T_{\text{oblivious}}^*(\boldsymbol{\mu})$ would be much harder, as the mixture distributions can no more be parameterized by their mean only. Our current techniques do not yield a general-purpose practical algorithm that is asymptotically optimal in the *oblivious* mode for the ABC-S problem. In the Bernoulli case, however, as mixtures of Bernoulli distributions are Bernoulli distribution, one can use the single-population Bernoulli approach discussed in the next paragraph. For Gaussian distributions, one can use a suboptimal approach based on the observation that location mixtures of Gaussians with bounded means are sub-Gaussian (see Appendix B.3 for details).

2.4 Single population and relationship with best arm identification

In order to illustrate the nature of the the ABC-S problem, we make a detour through the single population case, that is, when $J = 1$. Given two weights w_a, w_b and two means μ_a, μ_b , we introduce the minimum weighted transportation cost for moving the means to a common position.

$$d_{\text{mid}}(w_a, \mu_a, w_b, \mu_b) := \inf_v w_a d(\mu_a, v) + w_b d(\mu_b, v) = w_a d(\mu_a, v_{a,b}^*) + w_b d(\mu_b, v_{a,b}^*)$$

where $v_{a,b}^*$, the optimal common location, is the weighted average, i.e. $v_{a,b}^* = \frac{w_a}{w_a + w_b} \mu_a + \frac{w_b}{w_a + w_b} \mu_b$.

Constructing an instance in the alternative When identifying all the arms better than a control, there are two different ways to obtain a close-by bandit model $\boldsymbol{\lambda}$ in the alternative. The first option consists in taking an arm which does not belong to $\mathcal{S}_\beta(\boldsymbol{\mu})$ and to augment its mean on the alternative model such that it becomes above the control (or to reduce the mean of the control). Otherwise, it is possible to take an arm that is better than the control in the bandit model $\boldsymbol{\mu}$ and to shrink its mean such that it becomes lower than the control on the alternative (or augment the control). Note that the infimum over the alternative has the same expression in the two cases (see proof of Proposition 1 in Appendix A.2).

There is a priori no link between a BAI problem and an ABC one. In particular, in the BAI problem there are only $K + 1$ possible choices for the best arm while when looking for $\mathcal{S}_\beta(\boldsymbol{\mu})$ there are up to

2^K different sets to consider. Yet, the next proposition shows that the characteristic time T^* of any ABC problem with $J = 1$ subpopulation shares strong similarities with that of BAI problems.

Proposition 1. *Let $\delta \in (0, 1)$ and $\boldsymbol{\mu} \in \mathcal{L}$. For any strategy satisfying Equation 1, Equation 2 holds with*

$$T^*(\boldsymbol{\mu})^{-1} = \sup_{\boldsymbol{w} \in \Sigma_{K+1}} \inf_{\boldsymbol{\lambda} \in \text{Alt}(\boldsymbol{\beta})} \sum_{a=0}^K w_a d(\mu_a, \lambda_a) = \sup_{\boldsymbol{w} \in \Sigma_{K+1}} \min_{b \neq 0} d_{\text{mid}}(w_0, \mu_0, w_b, \mu_b).$$

The proof is reported in Appendix A.2. Note that the expression of the sample complexity is really close to the one in the BAI setting (Garivier and Kaufmann [10, Lemma 3]) except that we consider all the indices different from the control here instead of the indices different from the best arm.

2.5 The Gaussian case

In this section, we consider the Gaussian case which is of interest as the characteristic time admits a more explicit expression, making it possible to further investigate the differences between the various modes of interaction. We will state our results for the heteroscedastic case, in particular to get a closed-form proxy for the Bernoulli case, where each variance is a function of the (unknown) mean.

A/B testing When $K = 1$ (one arm and the control arm), we are considering a standard A/B test with subpopulations and one can easily prove the following result (established in Appendix C).

Proposition 2. *For any $\boldsymbol{\mu} \in \mathcal{L}$ with $K = 1$ and $\nu_{a,i} = \mathcal{N}(\mu_{a,j}, \sigma_{a,j}^2)$ one has*

1. $T_{\text{agnostic}}^*(\boldsymbol{\mu}) = \frac{2 \left(\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}} \right)^2}{\Delta_1^2}$ and $w_{a,i}^* = \frac{\alpha_i \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{a,i}^2}{\alpha_i}}}{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}}}$
2. $T_{\text{prop}}^*(\boldsymbol{\mu}) = \frac{2 \sum_{i=1}^J \frac{\beta_i^2}{\alpha_i} (\sigma_{0,i} + \sigma_{1,i})^2}{\Delta_1^2}$ and $\forall i \leq J, \forall a \in \{0, 1\}, w_{a,i}^* = \frac{\alpha_i \sigma_{a,i}}{\sigma_{0,i} + \sigma_{1,i}}$
3. $T_{\text{active}}^*(\boldsymbol{\mu}) = \frac{2 \left(\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i}) \right)^2}{\Delta_1^2}$ and $\forall i \leq J, \forall a \in \{0, 1\}, w_{a,i}^* = \frac{|\beta_i| \sigma_{a,i}}{\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}$

The optimal allocations in the *agnostic* and *proportional* cases are constrained by the proportion of the different subpopulations α , whereas, for the *active* mode, the optimal weights only depend on β . In general, the optimal weights also depend on the subpopulation variances, as is well-known in stratified sampling estimation. Note however, that when (a) the subpopulations all have a common variance σ^2 and (b) $\beta = \alpha$, then the optimal allocations and the characteristic times are equal for the *agnostic*, the *proportional* and the *active* modes. In that case, $w_{a,i}^* = \alpha_i/2$, which also corresponds to the well-known result in Gaussian A/B testing [14]. We have more generally observed that whenever the subpopulations have approximately the same variances, the *agnostic* and *proportional* modes yield very similar performances.

Weight computation in the homoscedastic case Even in scenarios where all subpopulation variances are equal to σ^2 , the *active* mode remains very attractive in the cases where $\beta \neq \alpha$. The following proposition shows that in that case, the optimal weights for the ABC-S problem can be computed efficiently.

Proposition 3 (Efficient computation in the Gaussian case). *With Gaussian distributions with a known variance σ^2 , letting $(u_0^*, \dots, u_K^*) = \arg\max_{\boldsymbol{u} \in \Sigma_{K+1}} \min_{b \neq 0} \frac{\Delta_b^2}{2 \left(\frac{1}{u_0} + \frac{1}{u_b} \right)}$, the optimal weights for the active mode satisfy*

$$\forall a \in \{0, \dots, K\}, \forall i \leq J, w_{a,i}^* = u_a^* \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|}.$$

If, in addition $\alpha = \beta$, the above also holds for the agnostic and the proportional modes.

The interesting part of Proposition 3 is that computing (u_0^*, \dots, u_K^*) can be done efficiently using Theorem 5 from [10]. The optimal weights of the ABC-S problem can be deduced from u^* without any further calculation.

3 Algorithms

To obtain our algorithms, we instantiate the Track-and-Stop algorithm template to our ABC-S problem. Garivier and Kaufmann [10] introduced Track-and-Stop and proved its asymptotic optimality in the BAI setting. Asymptotic optimality for general partition identification problems was subsequently established by Kaufmann and Koolen [13, Theorem 23] under the assumption of continuity of the oracle weights $\boldsymbol{\mu} \mapsto \boldsymbol{w}^*(\boldsymbol{\mu})$. Degenne and Koolen [6] show that the continuity assumption holds for all single-answer problems, in the upper-hemicontinuity sense, which they show implies asymptotic optimality of the Track-and-Stop (T-a-S) algorithm. These results directly apply to our ABC-S problem. Degenne et al. [7] interpret T-a-S as a noisy sequential equilibrium computation for the max-min problem from the lower bound (e.g. Equation 3) and develop computationally attractive variants including lazy iterative solution of the \boldsymbol{w}^* problem, and optimistic gradients instead of forced exploration.

The details of our implementation are given in Appendix F. In short, we use the simple standard $\Theta(\sqrt{t})$ forced exploration rounds, a mode/subpopulation aware upgrade of the D-tracking scheme [10] (which is empirically superior to C-tracking) and we approximately and incrementally compute the oracle weights using the *AdaHedge vs Best Response* iterative saddle point solver from [7]. We use one single learner, instead of one per possible answer, as advocated in [7, Section 4]. Note that we are not affected by the non-convergence of D-tracking from [7, Appendix E], as our problem has a unique \boldsymbol{w}^* because it is strictly concave in \boldsymbol{w} (see Appendix E).

The sampling rule The high level overview of the algorithm is as follows. We are given the number of arms K and subpopulations J , the exponential family, the mode of interaction, the subpopulation importance coefficients β and, for passive modes, their natural frequencies α . The algorithm then proceeds in rounds $t = 1, 2, \dots$. Each round t , it calculates the empirical frequencies $\hat{\boldsymbol{\mu}}_t \in \mathbb{R}^{(K+1) \times J}$ given by $\hat{\mu}_{a,i}(t) = \frac{1}{N_{a,i}(t)} \sum_{s=1}^t X_s \mathbf{1}\{A_s = a, I_s = i\}$. It then computes (a suitable approximation of) the maximiser (i.e. the oracle policy) $\boldsymbol{w}_t = \boldsymbol{w}^*(\hat{\boldsymbol{\mu}}_t) \in \Sigma_{(K+1) \times J}$ of problem (2). In the active mode, we “D-track” \boldsymbol{w}_t , i.e. we sample $(A_t, I_t) \in \operatorname{argmax}_{a,i} N_{a,i}(t-1) - t\boldsymbol{w}_t(a, i)$. In the proportional mode, the subpopulation I_t is given and we “D-track” the conditional distribution of \boldsymbol{w}_t on arms given the subpopulation, i.e. $A_t \in \operatorname{argmax}_a N_{a,I_t}(t-1) - t\alpha_{I_t} \boldsymbol{w}_t(a|I_t)$, where $\boldsymbol{w}_t(a, i) = \alpha_i \boldsymbol{w}_t(a|i)$. In the agnostic mode we “D-track” the marginal distribution of \boldsymbol{w}_t on arms, i.e. $A_t \in \operatorname{argmax}_a N_a(t-1) - t\boldsymbol{w}_t(a)$. For each mode, this sampling strategy ensures that $N_{a,i}(t) \approx t\boldsymbol{w}_t(a, i) \approx t\boldsymbol{w}_{a,i}^*(\boldsymbol{\mu})$, thus driving down the reported level of confidence as quickly as possible given the lower bound from Theorem 1.

The recommendation Concluding each round, we recommend $\mathcal{S}_\beta(\hat{\boldsymbol{\mu}}_t)$ at confidence level $\hat{\delta}(t) = \min\{\delta \in (0, 1) | \Lambda(t) \geq \beta(t, \delta)\}$ obtained by inverting the threshold $\beta(t, \delta)$ at the GLR statistic

$$\Lambda(t) = \min_{b \neq 0} \inf_{\boldsymbol{\lambda} \in \mathcal{L}: \lambda_0 = \lambda_b} \sum_{a \in \{0, b\}} \sum_{i=1}^J N_{a,i}(t) d(\hat{\mu}_{a,i}(t), \lambda_{a,i}). \quad (7)$$

The threshold For the sharpest theoretically supported thresholds we refer to [13]. Namely, an ABC-S problem with K -arms and J -subpopulations has 2^K answers, and its *rank* [13, Definition 22] is $2J$, as can be read off from (4). By [13, Proposition 23] we have validity for $\beta(t, \delta) = 6J \ln \ln t + \ln \frac{1}{\delta} + K + 2J \cdot O(\ln \ln \frac{1}{\delta})$. In practice, we follow [10] and use instead the heavily stylized $\ln((1 + \ln t)/\delta)$ that omits several union bounds.

Theorem 2. *For every mode, Subpopulation Track-and-Stop is safely calibrated (Equation 1). Moreover, Subpopulation Track-and-Stop is asymptotically optimal and matches the lower bound from Theorem 1, in the sense that*

$$\text{for every bandit } \boldsymbol{\mu} \in \mathcal{L}, \lim_{\delta \rightarrow 0} \frac{\mathbb{E}[\tau_\delta]}{\ln(1/\delta)} = T^*(\boldsymbol{\mu}).$$

We include the proof in Appendix E.

4 Experiments

4.1 Simulations

We conduct numerical experiments to evaluate the proposed algorithms, focusing on Bernoulli bandit models, which are ubiquitous in practical applications.

In our experiments, in addition to our T-a-S algorithms with the various interaction modes, we include two more sampling rules for comparison: (1) uniform sampling as a baseline, and (2) the experimentally efficient *Best Challenger* heuristic inspired by [10], adapted to the ABC problem and denoted BC-ABC in the sequel. BC [10] for the BAI problem samples in every round the empirical best arm \hat{a}_t or its best challenger, i.e. the arm $\hat{c}_t \neq \hat{a}_t$ at which the GLR statistic (Equation 7) reaches its minimum. Our BC-ABC adaptation samples in every round the control arm or the arm that yields the minimum GLR statistic $\Lambda(t)$, in the agnostic interaction mode (since $\Lambda(t)$ is subpopulation independent). For clearer comparison between the sampling strategies, all algorithms use the Chernoff stopping criterion [10] to determine either when to stop or output the risk assessment at a given time. We also opted for sampling rules independent from the confidence parameter δ , because we are aiming for safely calibrated policies.

We first illustrate the fact that the T-a-S algorithm provides a correct –but rather conservative– assessment of the risk of its decision whatever the time it is stopped at. To do so, we generated 1000 bandit instances uniformly at random from $[0, 1]$ with $K = 2$ arms. For each instance, we recorded the first time a certain risk assessment level is reached and the correctness of the algorithm’s recommendation at that point. We map to each risk assessment level the proportion of errors across all instances. We chose two stopping rates that are not supported by theory but are recommended in practice [10]. Figure 2 (Left) illustrates the isotonic curve fitted on our observations and suggests that even the most lenient stopping threshold $\ln((\ln(t) + 1)/\delta)$ results in much lower empirical probability of error than the risk assessment. In the following, we use the stopping threshold $\ln((\ln(t) + 1)/\delta)$.

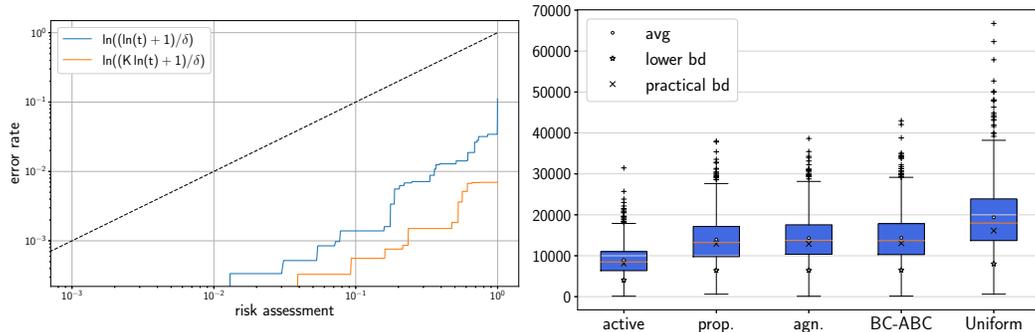


Figure 2: (Left) Risk assessment calibration on a log-log scale. (Right) Stopping time boxplot for $\mu = [0.1 \ 0.4 \ 0.3; 0.2 \ 0.5 \ 0.2; 0.5 \ 0.1 \ 0.1] \in [0, 1]^{(K+1) \times J}$ when $\beta = [1/3, 1/3, 1/3]$, $\alpha = [0.4, 0.5, 0.1]$ with Bernoulli distributions.

In our second experiment ¹, we generated 3000 Bernoulli bandit instances with $K = 2$ and a random number of subpopulations J between 2 and 10. Each subpopulation-arm’s mean $\mu_{a,i}$ is drawn uniformly at random from $[0, 1]$, and the subpopulation frequency vector α is drawn from a Dirichlet(10) distribution. Table 1 reports the average stopping time of each algorithm across all bandit instances. On average, the T-a-S algorithms at all modes stop at similar times, and all adaptive sampling methods terminate faster than uniform sampling.

Table 1: Average stopping time. Description in text.

T-a-S (active)	T-a-S (proportional)	T-a-S (agnostic)	BC-ABC	Uniform
14871	15231	15444	15279	21586

¹Code at https://gitlab.com/ckatsimerou/abc_s_public

To better understand the role of β and α , we ran the algorithms on a specific model (see Figure 2, Right) with $\alpha \neq \beta$. In this case, the optimal proportions are constrained by the frequencies of the subpopulation for passive interaction modes. The expected number of samples needed to identify the ABC-S solution is lower for the active policy, which has an additional degree of freedom in its sampling strategy. The *proportional* interaction mode and the *agnostic* interaction modes perform similarly. As expected, all the proposed strategies outperform the uniform sampling rule. We contrast the stopping time with the lower bound $\text{kl}(\delta, 1 - \delta)T^*(\mu)$, and with a more practical version, which indicates, approximately, the first time at which the GLR statistic crosses the threshold, i.e. solving $t = \ln((\ln(t) + 1)/\delta)T^*(\mu)$, as was done in [7]. All adaptive algorithms perform well on this instance, with their average runtime being very close to their respective practical bound.

4.2 Application to A/B/n experiment

We evaluate the algorithms on data collected from an actual A/B/n experiment, which compares different copies of a component of the webpage, in order to identify the ones better than the default copy. The metric of interest is whether the visitor clicked at least once during the experiment to the next page after getting exposed to one of the variants. For this setup we considered $K = 2$ copies competing against the control, with each copy being treated as an arm. Due to global traffic, the data exhibits strong seasonality patterns within a day, as seen in Figure 3a, in which every point corresponds to click-through rate per six hours (quarter of day) for 12 consecutive days. We treat the $J = 4$ seasons as i.i.d. subpopulations. Within each season we shuffled the data to eliminate the weekly trend.

The summary statistics of the dataset, together with the characteristic times and the optimal weights for each T-a-S mode can be found in Appendix G. Note that the small gaps between the arm means makes this practical ABC-S problem much harder than the synthetically generated examples.

We tested all algorithms described in Section 4.1. Each algorithm terminates when it reaches for the first time $\hat{\delta}_t \leq 0.1$ or outputs a risk assessment on the recommendation if it runs out of samples, which in this experiment occurs after $1.4 \cdot 10^7$ observations. Here, we weigh the importance β of each season equally to its observed frequency α . Doing so, we do not expect large performance discrepancies between the different T-a-S interaction modes, which is confirmed by their characteristic times (Appendix G). The observations from Fig. 3b are similar to the results from the numerical simulations: adaptive sampling achieves lower sample complexity over uniform sampling and T-a-S for the active interaction mode terminates faster than for the passive modes. All algorithms yield the correct recommendation, but not with the same risk assessment. All T-a-S algorithms terminated within the available sample size, BC-ABC almost terminated and output a risk assessment slightly above 0.1 and uniform’s risk assessment was 0.67. Of course, when viewing seasonality as a subpopulation, the *active* mode is unrealistic, but it is still informative to see that it can be very economical in hard problems in which sampling the subpopulations actively is an option. In this instance, *proportional*, *agnostic* and *oblivious* modes terminated at similar times. However, we would recommend using the *proportional* mode, given that we expect it to never perform worse than the other passive modes on average. One should not be surprised by the curve for the uniform sampling, this policy was stopped before convergence because it ran out of samples.

Lastly, here we assumed that seasons occur in i.i.d. fashion, but in reality there is temporal dependence between them. This imposes extra constraints on the optimal weights and increases the sample complexity. However, we do not expect this to be detrimental for cases in which seasons alternate frequently and full cycles are observed often, as was the case with our example.

5 Conclusion

In this work, we considered the pure exploration task of identifying all the arms that are better than a control arm in the presence of subpopulations (ABC-S). We design asymptotically optimal policies for this problem under different assumptions on the mode of interaction between the learner and the bandit. We observed that the *active* mode, in which the learner decides which subpopulation it samples, may significantly reduce decision times. On the other hand, the other modes, in which the learner has to respect the natural proportions of the different subpopulations (i.e., in *proportional* and *agnostic* modes) produce more modest effects, except when the subpopulations differ significantly in

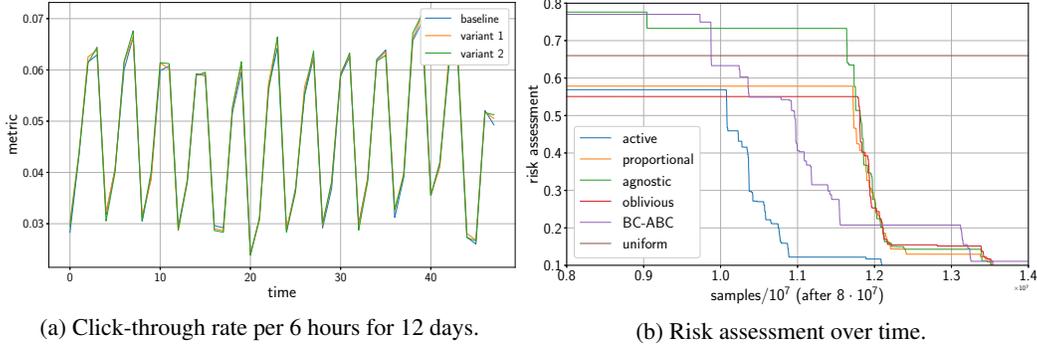


Figure 3: Real data and results.

variances. Finally, we proposed a natural way to provide anytime decisions with risk guarantees in the Track-and-Stop framework.

6 Potential Societal Impact

The contributions presented in this work are mostly related to methods and, as such, do not have a direct expected societal impact. This being said, a potential concern that will need to be addressed more carefully in subsequent applications of these methods is the use of subpopulation information, which could be exploited to target specific user behaviour or characteristics. In the use case considered in Section 4.2, the subpopulations correspond to time slots that are used to model seasonality in the user responses, which does not raise any specific ethical concern. However, in cases where the subpopulations are formed using characteristics of individual users, the impact needs to be assessed more thoroughly. Note that in such cases, restricting to one of the more conservative modes of interaction (i.e. agnostic or even oblivious) may become necessary in order to prevent undue use of population-dependent information.

Acknowledgment

The authors would like to thank the anonymous reviewers whose comments and questions helped improve the clarity of this manuscript. A. Garivier acknowledges the support of the Project IDEX-LYON of the University of Lyon, in the framework of the Programme Investissements d’Avenir (ANR-16-IDEX-0005), and Chaire SeqALO (ANR-20-CHIA-0020).

References

- [1] O. Cappé, A. Garivier, O.-A. Maillard, R. Munos, G. Stoltz, et al. Kullback–leibler upper confidence bounds for optimal sequential allocation. *Annals of Statistics*, 41(3):1516–1541, 2013.
- [2] A. Carpentier and R. Munos. Finite-time analysis of stratified sampling for monte carlo. In *NIPS-Twenty-Fifth Annual Conference on Neural Information Processing Systems*, 2011.
- [3] L. Chen, J. Li, and M. Qiao. Nearly instance optimal sample complexity bounds for top-k arm selection. In *Artificial Intelligence and Statistics*, pages 101–110. PMLR, 2017.
- [4] J. Cheshire, P. Menard, and A. Carpentier. The influence of shape constraints on the thresholding bandit problem. In *Conference on Learning Theory*, pages 1228–1275. PMLR, 2020.
- [5] S. de Rooij, T. van Erven, P. Grünwald, and W. M. Koolen. Follow the leader if you can, Hedge if you must. *Journal of Machine Learning Research*, 15:1281–1316, Apr. 2014.
- [6] R. Degenne and W. M. Koolen. Pure exploration with multiple correct answers. In *Advances in Neural Information Processing Systems (NeurIPS) 32*, pages 14591–14600. Dec. 2019.
- [7] R. Degenne, W. M. Koolen, and P. Ménard. Non-asymptotic pure exploration by solving games. In *Advances in Neural Information Processing Systems (NeurIPS) 32*, pages 14492–14501. Dec. 2019.
- [8] E. Even-Dar, S. Mannor, Y. Mansour, and S. Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- [9] V. Gabillon, M. Ghavamzadeh, and A. Lazaric. Best arm identification: A unified approach to fixed budget and fixed confidence. In *NIPS-Twenty-Sixth Annual Conference on Neural Information Processing Systems*, 2012.
- [10] A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In *Conference on Learning Theory*, pages 998–1027. PMLR, 2016.
- [11] R. Johari, L. Pekelis, and D. J. Walsh. Always valid inference: Bringing sequential analysis to A/B testing. *arXiv preprint arXiv:1512.04922*, 2015.
- [12] S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pages 655–662, 2012.
- [13] E. Kaufmann and W. M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. Preprint, Oct. 2018.
- [14] E. Kaufmann, O. Cappé, and A. Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.
- [15] T. Lattimore and C. Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020. doi: 10.1017/9781108571401.
- [16] A. Locatelli, M. Gutzeit, and A. Carpentier. An optimal algorithm for the thresholding bandit problem. In *International Conference on Machine Learning*, pages 1690–1698. PMLR, 2016.
- [17] B. Mason, L. Jain, A. Tripathy, and R. Nowak. Finding all $\{\epsilon\}$ -good arms in stochastic bandits. *Advances in Neural Information Processing Systems*, 2020.
- [18] F. Yang, A. Ramdas, K. Jamieson, and M. J. Wainwright. A framework for Multi-A(rmed)/B(andit) testing with online FDR control. In *Advances in Neural Information Processing Systems*, 2017.

A General form of the characteristic time

A.1 Proof of Theorem 1

Theorem 1. Let $\delta \in (0, 1)$ and $\beta \in \mathbb{R}^J$. For any strategy satisfying Equation 1 and any $\mu \in \mathcal{L}$, the expected number of rounds for the ABC-S problem for the agnostic, proportional and active mode satisfies:

$$\mathbb{E}_\mu[\tau_\delta] \geq T^*(\mu) \text{kl}(\delta, 1 - \delta) \quad \text{and} \quad \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T^*(\mu). \quad (2)$$

where (recalling that $\lambda_a = \sum_{i=1}^J \beta_i \lambda_{a,i}$)

$$T^*(\mu)^{-1} = \sup_{\mathbf{w} \in \mathcal{C}} \inf_{\lambda \in \text{Alt}_\beta(\mu)} \sum_{a=0}^K \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \quad (3)$$

$$= \sup_{\mathbf{w} \in \mathcal{C}} \min_{b \neq 0} \inf_{\lambda \in \mathcal{L}: \lambda_0 = \lambda_b} \sum_{a \in \{0, b\}} \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}). \quad (4)$$

Proof. Using the transportation lemma from [14] and recalling that $N_{a,i}(t)$ is the number of draws of arm a in subpopulation i up to time t , we have for any safely calibrated policies

$$\forall \lambda \in \text{Alt}_\beta(\mu), \quad \sum_{i=1}^J \sum_{a=0}^K \mathbb{E}_\mu[N_{a,i}(\tau_\delta)] d(\mu_{a,i}, \lambda_{a,i}) \geq \text{kl}(\delta, 1 - \delta).$$

Therefore,

$$\begin{aligned} \text{kl}(\delta, 1 - \delta) &\leq \inf_{\lambda \in \text{Alt}_\beta(\mu)} \sum_{a=0}^K \sum_{i=1}^J \mathbb{E}_\mu[N_{a,i}(\tau_\delta)] d(\mu_{a,i}, \lambda_{a,i}) \\ &= \mathbb{E}_\mu[\tau_\delta] \inf_{\lambda \in \text{Alt}_\beta(\mu)} \sum_{a=0}^K \sum_{i=1}^J \frac{\mathbb{E}_\mu[N_{a,i}(\tau_\delta)]}{\mathbb{E}_\mu[\tau_\delta]} d(\mu_{a,i}, \lambda_{a,i}) \\ &\leq \mathbb{E}_\mu[\tau_\delta] \sup_{\mathbf{w} \in \mathcal{C}} \inf_{\lambda \in \text{Alt}_\beta(\mu)} \sum_{a=0}^K \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}). \end{aligned}$$

In the last inequality, we used the fact that the normalized expected numbers of draws satisfy the set of constraints defined by $\mathcal{C} \subset \Sigma_{(K+1)J}$. Using $\text{kl}(\delta, 1 - \delta) \sim \ln(1/\delta)$ when δ tends to 0 gives the first result.

We denote $\Lambda(\mathbf{w}, \lambda, \mu) := \sum_{a=0}^K \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i})$. To obtain the second result, we will simplify the expression of $T^*(\mu)^{-1}$. Using that the KL divergences and the weights are positive, for λ to be in the alternative, one of the two following conditions need to be met: (1) there exists $a \in \mathcal{S}_\beta(\mu)$ such that $\lambda_a < \lambda_0$. (2) there exists $a \in \mathcal{S}_\beta^-(\mu) := \{a \in [K] \mid \mu_a < \mu_0\}$ such that $\lambda_a > \lambda_0$.

For this reason, one has

$$\inf_{\lambda \in \text{Alt}_\beta(\mu)} \Lambda(\mathbf{w}, \lambda, \mu) = \min \left(\min_{a \in \mathcal{S}_\beta(\mu)} \inf_{\lambda: \lambda_a < \lambda_0} \Lambda(\mathbf{w}, \lambda, \mu), \min_{a \in \mathcal{S}_\beta^-(\mu)} \inf_{\lambda: \lambda_a > \lambda_0} \Lambda(\mathbf{w}, \lambda, \mu) \right).$$

We obtain the desired result by remarking that the inner optimization programs \inf_λ are each achieved on the boundary (the constraint being satisfied with equality) where they coincide, and that $\{1, \dots, K\} = \mathcal{S}_\beta(\mu) \cup \mathcal{S}_\beta^-(\mu)$. \square

A.2 Proof of Proposition 1

In the particular case when $J = 1$, the expression of the characteristic time can be simplified.

Proposition 1. Let $\delta \in (0, 1)$ and $\boldsymbol{\mu} \in \mathcal{L}$. For any strategy satisfying Equation 1, Equation 2 holds with

$$T^*(\boldsymbol{\mu})^{-1} = \sup_{\boldsymbol{w} \in \Sigma_{K+1}} \inf_{\boldsymbol{\lambda} \in \text{Alt}_\beta(\boldsymbol{\mu})} \sum_{a=0}^K w_a d(\mu_a, \lambda_a) = \sup_{\boldsymbol{w} \in \Sigma_{K+1}} \min_{b \neq 0} d_{\text{mid}}(w_0, \mu_0, w_b, \mu_b).$$

Proof. The first part of the proof can be obtained using similar argument than for Theorem 1. The missing part is the simplification of the expression of $T^*(\boldsymbol{\mu})$.

We denote $\Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) := \sum_{a=0}^K w_a d(\mu_a, \lambda_a)$. Following the reasoning from the proof of Theorem 1 one of the two following conditions needs to be met: (1) there exists $a \in \mathcal{S}_\beta(\boldsymbol{\mu})$ such that $\lambda_a < \lambda_0$. (2) there exists $a \in \mathcal{S}_\beta^-(\boldsymbol{\mu}) := \{a \in [K] \mid \mu_a < \mu_0\}$ such that $\lambda_a > \lambda_0$.

For this reason, one has

$$\inf_{\boldsymbol{\lambda} \in \text{Alt}_\beta(\boldsymbol{\mu})} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \min \left(\min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a < \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}), \min_{a \in \mathcal{S}_\beta^-(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a > \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \right).$$

In this simpler case, it is possible to obtain an explicit formula for this infimum. We start from

$$T^*(\boldsymbol{\mu})^{-1} = \sup_{\boldsymbol{w} \in \Sigma_{K+1}} \min \left(\min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a < \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}), \min_{a \in \mathcal{S}_\beta^-(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a > \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) \right).$$

Let us focus on the case, $\lambda_a < \lambda_0$ and fix an index $a \in \mathcal{S}_\beta(\boldsymbol{\mu})$. Λ is always smaller when all the λ_b for $b \neq 0$ and $b \neq a$ coincides with μ_b . This gives,

$$\min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a < \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a \leq \lambda_0} w_0 d(\mu_0, \lambda_0) + w_a d(\mu_a, \lambda_a).$$

We consider the Lagrangian function, $L(\lambda_0, \lambda_a, q) = w_0 d(\mu_0, \lambda_0) + w_a d(\mu_a, \lambda_a) + q(\lambda_a - \lambda_0)$. Differentiating with respect to λ_0 and λ_a brings the condition

$$\lambda_0^* = \lambda_a^* = \lambda_{a,0}^* = \operatorname{argmin}_{\lambda} w_0 d(\mu_0, \lambda) + w_a d(\mu_a, \lambda) = \frac{w_0}{w_0 + w_a} \mu_0 + \frac{w_a}{w_0 + w_a} \mu_a.$$

Recalling, $d_{\text{mid}}(w_a, \mu_a, w_b, \mu_b) := \inf_v w_a d(\mu_a, v) + w_b d(\mu_b, v)$ one has,

$$\min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a < \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \min_{a \in \mathcal{S}_\beta(\boldsymbol{\mu})} d_{\text{mid}}(w_0, \mu_0, w_a, \mu_a). \quad (8)$$

Solving the optimization program for $a \in \mathcal{S}_\beta^-(\boldsymbol{\mu})$ and under the constraint $\lambda_a > \lambda_0$, gives the exact same set of constraints and optimal solution, i.e.

$$\min_{a \in \mathcal{S}_\beta^-(\boldsymbol{\mu})} \inf_{\boldsymbol{\lambda}: \lambda_a > \lambda_0} \Lambda(\boldsymbol{w}, \boldsymbol{\lambda}, \boldsymbol{\mu}) = \min_{a \in \mathcal{S}_\beta^-(\boldsymbol{\mu})} d_{\text{mid}}(w_0, \mu_0, w_a, \mu_a). \quad (9)$$

Bringing Equation 8 and Equation 9 together and remarking that $[K] = \mathcal{S}_\beta(\boldsymbol{\mu}) \cup \mathcal{S}_\beta^-(\boldsymbol{\mu})$ gives the announced result. \square

A.3 Link between ABC and BAI

In the particular case, of Gaussian distributions with a known variance σ^2 , for any $\boldsymbol{\mu} \in \mathcal{L}$, one can easily create a BAI instance with the same characteristic time as $T^*(\boldsymbol{\mu})$.

Lemma 1. Let $\boldsymbol{\mu} \in \mathcal{L}$ where all the arms are Gaussian distributions with known variance σ^2 . Let μ_0 be the mean of the control arm. We define $\tilde{\boldsymbol{\mu}}$ as follows

$$\tilde{\mu}_k = \begin{cases} 2\mu_0 - \mu_k & \text{if } \mu_k > \mu_0 \\ \mu_k & \text{otherwise.} \end{cases}$$

By denoting $T_{\text{BAI}}^*(\boldsymbol{\mu})$ the characteristic time for the BAI problem with a bandit instance $\boldsymbol{\mu}$, then one has

$$T^*(\boldsymbol{\mu}) = T_{\text{BAI}}^*(\tilde{\boldsymbol{\mu}}).$$

Proof. In the particular case of Gaussian distributions with known variance σ^2 , easy calculation brings

$$T^*(\boldsymbol{\mu})^{-1} = \sup_{w \in \Sigma_{K+1}} \min_{b \neq 0} d_{\text{mid}}(w_0, \mu_0, w_b, \mu_b) = \min_{b \neq 0} \frac{(\mu_0 - \mu_b)^2}{2} ((1 - \alpha_b)^2 w_0 + \alpha_b^2 w_b)$$

with $\alpha_b = \frac{w_0}{w_b + w_0}$.

On the instance $\tilde{\boldsymbol{\mu}}$, first note that $\tilde{\mu}_0$ is the best arm by construction. For an index k such that $\mu_k > \mu_0$, one has

$$\tilde{\mu}_0 - \tilde{\mu}_k = \mu_0 - (2\mu_0 - \mu_k) = \mu_k - \mu_0 > 0.$$

For this reason, $\forall k \in [K], \tilde{\mu}_0 > \tilde{\mu}_k$. Using Lemma 3 from [10], by defining $I_\gamma(\mu_1, \mu_2) := \gamma d(\mu_1, \gamma\mu_1 + (1-\gamma)\mu_2) + (1-\gamma)d(\mu_2, \gamma\mu_1 + (1-\gamma)\mu_2)$, one has

$$T_{\text{BAI}}^*(\tilde{\boldsymbol{\mu}}) = \sup_{w \in \Sigma_{K+1}} \min_{b \neq 0} (w_0 + w_b) I_{\frac{w_0}{w_0 + w_b}}(\mu_0, \mu_b).$$

Furthermore, letting $\alpha_b = w_0/(w_0 + w_b)$:

$$\begin{aligned} (w_0 + w_b) I_{\alpha_b}(\mu_0, \mu_b) &= (w_0 + w_b) \left(\alpha_b \frac{(1 - \alpha_b)^2 (\mu_0 - \mu_b)^2}{2\sigma^2} + (1 - \alpha_b) \frac{\alpha_b^2 (\mu_0 - \mu_b)^2}{2\sigma^2} \right) \\ &= \frac{(\mu_0 - \mu_b)^2}{2\sigma^2} ((1 - \alpha_b)^2 w_0 + \alpha_b^2 w_b). \end{aligned}$$

Plugging this in the expression of $T_{\text{BAI}}^*(\tilde{\boldsymbol{\mu}})$ gives the announced result. \square

B Results for specific modes of interaction

B.1 Agnostic mode

Lemma 2. For any agnostic policy where A_t is chosen knowing \mathcal{F}_{t-1} but independently from I_t , when defining $N_{a,j}(t) = \sum_{s=1}^t \mathbb{1}(A_s = a \cap I_s = j)$ and $N_a(t) = \sum_{s=1}^t \mathbb{1}(A_s = a)$, then

$$\forall a \in \{0, \dots, K\}, \forall j \in \{1, \dots, J\}, \forall t \geq 1, \quad \mathbb{E}_{\boldsymbol{\mu}}[N_{a,j}(t)] = \alpha_j \mathbb{E}_{\boldsymbol{\mu}}[N_a(t)]$$

Proof.

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\mu}}[N_{a,j}(t)] &= \sum_{s=1}^t \mathbb{P}(A_s = a \cap I_s = j) = \sum_{s=1}^t \mathbb{P}_{\boldsymbol{\mu}}(A_s = a | I_s = j) \mathbb{P}(I_s = j) \\ &= \sum_{s=1}^t \alpha_j \mathbb{P}_{\boldsymbol{\mu}}(A_s = a | I_s = j) = \sum_{s=1}^t \alpha_j \mathbb{P}_{\boldsymbol{\mu}}(A_s = a) \\ &= \alpha_j \mathbb{E}_{\boldsymbol{\mu}}[N_a(t)], \end{aligned}$$

where in the third equality, we have used that the action A_t is selected independently from the population indicator I_t . \square

B.2 Proportional mode

Lemma 3. For any proportional policy where A_t is chosen knowing \mathcal{F}_{t-1} and I_t , when defining $N_{a,j}(t) = \sum_{s=1}^t \mathbb{1}(A_s = a \cap I_s = j)$ and $N_a(t) = \sum_{s=1}^t \mathbb{1}(A_s = a)$, then

$$\forall j \in \{1, \dots, J\}, \forall t \geq 1, \quad \sum_{a=0}^K \mathbb{E}_{\boldsymbol{\mu}}[N_{a,j}(t)] = \alpha_j t.$$

Proof.

$$\begin{aligned} \sum_{a=0}^K \mathbb{E}_{\boldsymbol{\mu}}[N_{a,j}(t)] &= \sum_{s=1}^t \sum_{a=0}^K \mathbb{E}_{\boldsymbol{\mu}}[\mathbb{1}(I_s = j) \mathbb{1}(A_s = a)] = \sum_{s=1}^t \mathbb{E}_{\boldsymbol{\mu}} \left[\mathbb{1}(I_s = j) \sum_{a=0}^K \mathbb{1}(A_s = a) \right] \\ &= \sum_{s=1}^t \mathbb{P}_{\boldsymbol{\mu}}(I_s = j) = \alpha_j t. \end{aligned}$$

\square

B.3 Oblivious mode

In the oblivious mode, the subpopulations can not be observed by the learner. In this case, we have

1. $\mathbb{E}[X_t | A_t = a] = \sum_{i=1}^J \alpha_i \mu_{a,i}$,
2. $X_t | A_t = a \sim \sum_{i=1}^J \alpha_i \nu_{a,i}$.

While with observable subpopulations the distributions are entirely characterized by their means, this is no longer the case with mixture distributions. In particular, this requires defining a different alternative.

$$\text{Alt}(\nu) := \{ \nu' : \forall a, \nu'_a = \sum_{i=1}^J \alpha_i \nu'_{a,i} \text{ with } \nu'_{a,i} \in \mathcal{P} \text{ and } \mathcal{S}(\nu') \neq \mathcal{S}(\nu) \}.$$

Proposition 4. *Let $\delta \in (0, 1)$ and $\beta \in \mathbb{R}^J$. For any oblivious strategy satisfying Equation 1 and any $\mu \in \mathcal{L}$, the characteristic time satisfies*

$$\mathbb{E}_\mu[\tau_\delta] \geq T_{\text{oblivious}}^*(\mu) \text{kl}(\delta, 1 - \delta) \quad \text{and} \quad \liminf_{\delta \rightarrow 0} \frac{\mathbb{E}_\mu[\tau_\delta]}{\ln(1/\delta)} \geq T_{\text{oblivious}}^*(\mu).$$

where

$$T_{\text{oblivious}}^*(\mu)^{-1} = \sup_{\mathbf{w} \in \Sigma_{K+1}} \inf_{\nu' \in \text{Alt}(\nu)} \sum_{a=0}^K w_a \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right). \quad (10)$$

Furthermore,

$$\forall \mu \in \mathcal{L}, \quad T_{\text{oblivious}}^*(\mu) \geq T_{\text{agnostic}}^*(\mu).$$

Proof. Using the transportation lemma from [14] we have for any safely calibrated oblivious policy

$$\forall \nu' \in \text{Alt}(\nu), \quad \sum_{a=0}^K \mathbb{E}_\mu[N_a(\tau_\delta)] \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right) \geq \text{kl}(\delta, 1 - \delta).$$

Therefore,

$$\begin{aligned} \text{kl}(\delta, 1 - \delta) &\leq \inf_{\nu' \in \text{Alt}(\nu)} \sum_{a=0}^K \mathbb{E}_\mu[N_a(\tau_\delta)] \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right) \\ &= \mathbb{E}_\mu[\tau_\delta] \inf_{\nu' \in \text{Alt}(\nu)} \sum_{a=0}^K \frac{\mathbb{E}_\mu[N_a(\tau_\delta)]}{\mathbb{E}_\mu[\tau_\delta]} \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right) \\ &\leq \mathbb{E}_\mu[\tau_\delta] \sup_{\mathbf{w} \in \Sigma_{K+1}} \inf_{\nu' \in \text{Alt}(\nu)} \sum_{a=0}^K w_a \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right). \end{aligned}$$

Using $\text{kl}(\delta, 1 - \delta) \sim \ln(1/\delta)$ when δ tends to 0 gives the first result.

Using the joint convexity of the KL divergence one gets

$$\text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right) \leq \sum_{i=1}^J \alpha_i \text{KL}(\nu_{a,i}, \nu'_{a,i}).$$

Assuming that the mean of $\nu'_{a,i} = \lambda_{a,i}$ and recalling that for distributions in \mathcal{P} , one has $\text{KL}(\nu_{a,i}, \nu'_{a,i}) = d(\mu_{a,i}, \lambda_{a,i})$, we deduce,

$$\begin{aligned} T_{\text{oblivious}}^*(\mu)^{-1} &= \sup_{\mathbf{w} \in \Sigma_{K+1}} \inf_{\nu' \in \text{Alt}(\nu)} \sum_{a=0}^K w_a \text{KL} \left(\sum_{i=1}^J \alpha_i \nu_{a,i}, \sum_{i=1}^J \alpha_i \nu'_{a,i} \right) \\ &\leq \sup_{\mathbf{w} \in \Sigma_{K+1}} \inf_{\lambda \in \text{Alt}(\mu)} \sum_{a=0}^K \sum_{i=1}^J \alpha_i w_a d(\mu_{a,i}, \lambda_{a,i}) \\ &= T_{\text{agnostic}}^*(\mu)^{-1}. \end{aligned}$$

□

Except in the case of Bernoulli distributions –where the mixture is also a Bernoulli distribution–, finding a strategy that matches $T_{\text{oblivious}}^*(\boldsymbol{\mu})^{-1}$ is a hard task. However, one may use the following lemma to treat the mixture in a sub-optimal way, based on the fact that it exhibits sub-gaussian behavior.

Lemma 4 (Sub-gaussianity of mixture). *For each $\mu \in \mathbb{R}$, assume that ν_μ is a distribution on \mathbb{R} with mean $\mathbb{E}_{X \sim \nu_\mu}[X] = \mu$ that is σ^2 -sub-Gaussian, meaning that $\mathbb{E}_{X \sim \nu_\mu}[e^{\lambda(X-\mu)}] \leq e^{\sigma^2 \lambda^2/2}$ for any $\lambda \in \mathbb{R}$. Further let $\alpha(\mu)$ be a prior on μ with mean m that is itself η^2 -sub-Gaussian, meaning that $\mathbb{E}_{\mu \sim \alpha}[e^{\lambda(\mu-m)}] \leq e^{\lambda^2 \eta^2/2}$. Then the mixture distribution $Q = \mathbb{E}_{\mu \sim \alpha}[\nu_\mu]$ is $\sigma^2 + \eta^2$ sub-Gaussian.*

Proof. The mixture distribution obviously has mean $\mathbb{E}_{X \sim Q}[X] = m$ and

$$\begin{aligned} \mathbb{E}_{X \sim Q}[e^{\lambda(X-m)}] &= \mathbb{E}_{\mu \sim \alpha} \left[e^{\lambda(\mu-m)} \mathbb{E}_{X \sim \nu_\mu}[e^{\lambda(X-\mu)}] \right] \\ &\leq \mathbb{E}_{\mu \sim \alpha} \left[e^{\lambda(\mu-m)} \right] e^{\sigma^2 \lambda^2/2} \\ &\leq e^{(\sigma^2 + \eta^2) \lambda^2/2}. \end{aligned}$$

□

In particular, if α is supported on $[\pm M]$, then α is M^2 sub-Gaussian, and hence Q is $(\sigma^2 + M^2)$ sub-Gaussian.

C Optimal allocations in the Gaussian case for $K = 1$ (A/B testing)

Lemma 5. *When $K = 1$ with Gaussian distributions such that $\nu_{a,i} = \mathcal{N}(\mu_{a,i}, \sigma_{a,i}^2)$ the following holds*

$$\inf_{\lambda: \lambda_0 = \lambda_1} \sum_{i=1}^J w_{0,i} d(\mu_{0,i}, \lambda_{0,i}) + \sum_{i=1}^J w_{1,i} d(\mu_{1,i}, \lambda_{1,i}) = \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)}.$$

Proof. One has for $b \in \{0, 1\}$,

$$d(\mu_{b,i}, \lambda_{b,i}) = \frac{(\lambda_{b,i} - \mu_{b,i})^2}{2\sigma_{b,i}^2}.$$

Using the result from Theorem 1 for the case $K = 1$, the following holds

$$\inf_{\boldsymbol{\lambda} \in \text{Alt}_\beta(\boldsymbol{\mu})} \sum_{a=0}^1 \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) = \min_{\boldsymbol{\lambda} \in \mathcal{L}: \lambda_0 = \lambda_1} \sum_{a=0}^1 \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}).$$

We introduce

$$L(\lambda_0, \lambda_1, q) = \sum_{i=1}^J w_{0,i} \frac{(\lambda_{0,i} - \mu_{0,i})^2}{2\sigma_{0,i}^2} + \sum_{i=1}^J w_{1,i} \frac{(\lambda_{1,i} - \mu_{1,i})^2}{2\sigma_{1,i}^2} + q \left(\sum_{i=1}^J \beta_i (\lambda_{0,i} - \lambda_{1,i}) \right).$$

One has,

$$\min_{\boldsymbol{\lambda} \in \mathcal{L}: \lambda_0 = \lambda_1} \sum_{a=0}^1 \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) = \sup_{q \in \mathbb{R}} \inf_{\boldsymbol{\lambda} \in \mathcal{L}} L(\lambda_0, \lambda_1, q).$$

Differentiating with respect to $\lambda_{0,i}$ and $\lambda_{1,i}$ brings the conditions

$$\lambda_{0,i} = \mu_{0,i} - \frac{q\beta_i\sigma_{0,i}^2}{w_{0,i}} \quad \text{and} \quad \lambda_{1,i} = \mu_{1,i} + \frac{q\beta_i\sigma_{1,i}^2}{w_{1,i}}.$$

Plugging these values back in L gives the function

$$f(q) = -\frac{q^2}{2} \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right) + q \sum_{i=1}^J \beta_i (\mu_{0,i} - \mu_{1,i}).$$

Easy calculations show that the maximum of the function f is attained for

$$q^* = \frac{\sum_{i=1}^J \beta_i (\mu_{0,i} - \mu_{1,i})}{\sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)}.$$

Plugging this value back in the expression of f ,

$$f(q^*) = \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)}.$$

□

Proposition 2. For any $\mu \in \mathcal{L}$ with $K = 1$ and $\nu_{a,i} = \mathcal{N}(\mu_{a,j}, \sigma_{a,j}^2)$ one has

1. $T_{\text{agnostic}}^*(\mu) = \frac{2 \left(\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}} \right)^2}{\Delta_1^2}$ and $w_{a,i}^* = \frac{\alpha_i \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{a,i}^2}{\alpha_i}}}{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}}}$
2. $T_{\text{prop}}^*(\mu) = \frac{2 \sum_{i=1}^J \frac{\beta_i^2 (\sigma_{0,i} + \sigma_{1,i})^2}{\alpha_i}}{\Delta_1^2}$ and $\forall i \leq J, \forall a \in \{0, 1\}, w_{a,i}^* = \frac{\alpha_i \sigma_{a,i}}{\sigma_{0,i} + \sigma_{1,i}}$
3. $T_{\text{active}}^*(\mu) = \frac{2 \left(\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i}) \right)^2}{\Delta_1^2}$ and $\forall i \leq J, \forall a \in \{0, 1\}, w_{a,i}^* = \frac{|\beta_i| \sigma_{a,i}}{\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}$

Proof.

Agnostic mode From the Lemma 2, we have

$$\begin{aligned} T_{\text{agnostic}}^*(\mu)^{-1} &= \sup_{\mathbf{w} \in \mathcal{C}_{\text{agnostic}}} \inf_{\lambda \in \text{Alt}_{\beta}(\mu)} \sum_{a=0}^1 \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \\ &= \sup_{\mathbf{w} \in \mathcal{C}_{\text{agnostic}}} \inf_{\lambda: \lambda_0 = \lambda_1} \sum_{a=0}^1 \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}) \quad (\text{Theorem 1}) \\ &= \sup_{\mathbf{w} \in \mathcal{C}_{\text{agnostic}}} \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)} \quad (\text{Lemma 5}). \end{aligned}$$

$\mathbf{w} \in \mathcal{C}_{\text{agnostic}}$ implies $w_{a,i} = \alpha_i u_a$ with $(u_0, \dots, u_K) \in \Sigma_{K+1}$. For this reason,

$$\mathbf{w}^* = \underset{\mathbf{u}: u_0 + u_1 = 1}{\text{argmin}} \sum_{i=1}^J \frac{\beta_i^2}{\alpha_i} \left(\frac{\sigma_{0,i}^2}{u_0} + \frac{\sigma_{1,i}^2}{u_1} \right).$$

We let $c_a := \sum_{i=1}^J \frac{\beta_i^2 \sigma_{a,i}^2}{\alpha_i}$ for $a \in \{0, 1\}$. Plugging $u_1 = 1 - u_0$ in the previous expression and differentiating with respect to u_0 brings the condition

$$u_0^2 + 2u_0 \frac{c_0}{c_1 - c_0} - \frac{c_0}{c_1 - c_0}.$$

Solving this polynomial and using that $\mathbf{u} \in \Sigma_2$ gives the unique solution

$$u_0^* = \frac{\sqrt{c_0}}{\sqrt{c_0} + \sqrt{c_1}}.$$

Implying,

$$w_{0,i}^* = \alpha_i \frac{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}}}{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}}} \quad \text{and} \quad w_{1,i}^* = \alpha_i \frac{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}}}{\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}}}.$$

With those values,

$$T_{\text{agnostic}}^*(\boldsymbol{\mu}) = \frac{2 \left(\sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{0,i}^2}{\alpha_i}} + \sqrt{\sum_{i=1}^J \frac{\beta_i^2 \sigma_{1,i}^2}{\alpha_i}} \right)^2}{\Delta_1^2}.$$

Proportional mode Following the same line of proof, gives

$$T_{\text{prop}}^*(\boldsymbol{\mu})^{-1} = \sup_{\mathbf{w} \in \mathcal{C}_{\text{prop}}} \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)}.$$

The main difference is now on the constraints on the weights. In the proportional mode, following Lemma 3, $\forall i \leq J, \sum_{a=0}^1 w_{a,i} = \alpha_i$. We consider the Lagrangian function:

$$L(w_0, w_1, q_1, \dots, q_J) = \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right) + \sum_{i=1}^J q_i \left(\sum_{a \in \{0,1\}} w_{a,i} - \alpha_i \right).$$

Differentiating with respect to $w_{0,i}$ and $w_{1,i}$ gives the constraints:

$$\frac{-\beta_i^2 \sigma_{0,i}^2}{w_{0,i}^2} + q_i = 0 \quad \text{and} \quad \frac{-\beta_i^2 \sigma_{1,i}^2}{w_{1,i}^2} + q_i = 0.$$

From which we can deduce

$$\frac{w_{0,i}}{\sigma_{0,i}} = \frac{w_{1,i}}{\sigma_{1,i}}.$$

From $w_{0,i} + w_{1,i} = \alpha_i$, we deduce,

$$q_i^* = \frac{\beta_i^2 (\sigma_{0,i} + \sigma_{1,i})^2}{\alpha_i^2}.$$

Plugging this value in the first constraint gives

$$w_{0,i}^* = \alpha_i \frac{\sigma_{0,i}}{\sigma_{0,i} + \sigma_{1,i}} \quad \text{and} \quad w_{1,i}^* = \alpha_i \frac{\sigma_{1,i}}{\sigma_{0,i} + \sigma_{1,i}}.$$

Using those weights,

$$T_{\text{prop}}^*(\boldsymbol{\mu}) = \frac{2 \sum_{i=1}^J \frac{\beta_i^2}{\alpha_i} (\sigma_{0,i} + \sigma_{1,i})^2}{\Delta_1^2}.$$

Active mode Following the proof of Proposition 2, one has

$$T_{\text{active}}^*(\boldsymbol{\mu})^{-1} = \sup_{\mathbf{w} \in \Sigma_{(K+1)J}} \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma_{0,i}^2}{w_{0,i}} + \frac{\sigma_{1,i}^2}{w_{1,i}} \right)}. \quad (11)$$

Using the constraint $\mathbf{w} \in \Sigma_{(K+1)J}$, one gets

$$w_{1,J} = 1 - \sum_{a=\{0,1\}} \sum_{i=1}^{J-1} w_{a,i} - w_{0,J}. \quad (12)$$

We need to minimize the function (where $w_{1,J}$ has been replaced by the expression from Equation 12)

$$f(\mathbf{w}) = \sum_{a=\{0,1\}} \sum_{i=1}^{J-1} \beta_i^2 \frac{\sigma_{a,j}^2}{w_{a,j}} + \beta_J^2 \frac{\sigma_{0,J}^2}{w_{0,J}} + \beta_J^2 \frac{\sigma_{1,J}^2}{1 - \sum_{a=0}^1 \sum_{i=1}^{J-1} w_{a,i} - w_{0,J}}.$$

For $i \leq J - 1$, taking the derivative with respect to $w_{0,i}$ and $w_{1,i}$ gives the following constraints

$$\begin{aligned}\beta_i^2 \sigma_{0,i}^2 \left(1 - \sum_{a=0}^1 \sum_{i=1}^{J-1} w_{a,i} - w_{0,J} \right)^2 &= \beta_J^2 \sigma_{1,J}^2 w_{0,i}^2, \\ \beta_J^2 \sigma_{1,i}^2 \left(1 - \sum_{a=0}^1 \sum_{i=1}^{J-1} w_{a,i} - w_{0,J} \right)^2 &= \beta_J^2 \sigma_{1,J}^2 w_{1,i}^2.\end{aligned}$$

From which we deduce

$$\forall i \leq J - 1, \quad \frac{w_{0,i}}{\sigma_{0,i}} = \frac{w_{1,i}}{\sigma_{1,i}}. \quad (13)$$

Differentiating with respect to $w_{1,J}$ gives

$$\sigma_{0,J} \left(1 - \sum_{a=0}^1 \sum_{i=1}^{J-1} w_{a,i} - w_{0,J} \right) = \sigma_{1,J} w_{0,J}.$$

Rearranging and using Equation 13 gives,

$$w_{0,J} = \frac{\sigma_{0,J}}{\sigma_{0,J} + \sigma_{1,J}} - \sum_{i=1}^{J-1} \frac{w_{0,i}}{\sigma_{0,i}} \frac{\sigma_{0,i} + \sigma_{1,i}}{\sigma_{0,J} + \sigma_{1,J}} \sigma_{0,J}. \quad (14)$$

Using Equation 13 and Equation 14, we define the function

$$g(w_{0,1}, \dots, w_{0,J-1}) = \sum_{i=1}^{J-1} \beta_i^2 \frac{\sigma_{0,i}}{w_{0,i}} (\sigma_{0,i} + \sigma_{1,i}) + \frac{\beta_J^2 (\sigma_{0,J} + \sigma_{1,J})^2}{1 - \sum_{i=1}^{J-1} \frac{w_{0,i} (\sigma_{0,i} + \sigma_{1,i})}{\sigma_{0,i}}}.$$

Differentiating with respect to $w_{0,i}$ for $i \leq J - 1$ brings

$$\forall i \leq J - 1, \quad \frac{|\beta_i| \sigma_{0,i}}{\sigma_{0,J} + \sigma_{1,J}} \left(1 - \sum_{i=1}^{J-1} \frac{w_{0,i}}{\sigma_{0,i}} (\sigma_{0,i} + \sigma_{1,i}) \right) = |\beta_J| w_{0,i}. \quad (15)$$

Multiplying both sides of this equation by $(\sigma_{0,i} + \sigma_{1,i})/\sigma_{0,i}$ and summing for $i \leq J - 1$,

$$\sum_{i=1}^{J-1} \frac{w_{0,i}}{\sigma_{0,i}} (\sigma_{0,i} + \sigma_{1,i}) = \frac{\sum_{i=1}^{J-1} |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}{\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}.$$

Plugging this value back in Equation 15 one has:

$$\forall i \leq J - 1, \quad w_{0,i} = \frac{|\beta_i| \sigma_{0,i}}{\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}.$$

From Equation 13, we deduce,

$$\forall i \leq J - 1, \quad w_{1,i} = \frac{|\beta_i| \sigma_{1,i}}{\sum_{i=1}^J |\beta_i| (\sigma_{0,i} + \sigma_{1,i})}.$$

We obtain the value of $w_{0,J}$ using Equation 14 and that of $w_{1,J}$ using Equation 12. Plugging those weights in the expression given by Equation 11 yields the characteristic time. \square

D General form of the optimal allocation in the Gaussian case

Proposition 3 (Efficient computation in the Gaussian case). *With Gaussian distributions with a known variance σ^2 , letting $(u_0^*, \dots, u_K^*) = \operatorname{argmax}_{u \in \Sigma_{K+1}} \min_{b \neq 0} \frac{\Delta_b^2}{2 \left(\frac{1}{u_0} + \frac{1}{u_b} \right)}$, the optimal weights for the active mode satisfy*

$$\forall a \in \{0, \dots, K\}, \quad \forall i \leq J, \quad w_{a,i}^* = u_a^* \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|}.$$

If, in addition $\alpha = \beta$, the above also holds for the agnostic and the proportional modes.

Proof. From Lemma 5, when the distribution are Gaussian with a known variance σ^2 one has

$$T_{\text{active}}^*(\boldsymbol{\mu})^{-1} = \sup_{\mathbf{w} \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_1^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right)}.$$

Using the same continuity argument than in [10], we know that the supremum of \mathbf{w} is attained and is indeed a maximum. Let

$$\Lambda_b(v, w) := \frac{\Delta_b^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{v_i} + \frac{\sigma^2}{w_i} \right)}.$$

Then

$$\begin{aligned} \max_{\mathbf{w} \in \Sigma_{(K+1)J}} \min_{b \neq 0} \Lambda_b(w_0, w_b) &= \max_{\substack{u \in \Sigma_{K+1} \\ \forall a, \sum_{i=1}^J w_{a,i} = u_a}} \min_{b \neq 0} \Lambda_b(w_0, w_b) \\ &= \max_{u \in \Sigma_{K+1}} \max_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \min_{b \neq 0} \Lambda_b(w_0, w_b) \\ &\leq \max_{u \in \Sigma_{K+1}} \min_{b \neq 0} \max_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \Lambda_b(w_0, w_b) \quad (\text{Max-min inequality}). \end{aligned}$$

Let $b \neq 0$,

$$\max_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \Lambda_b(w_0, w_b) = \max_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \frac{\Delta_b^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right)}.$$

Equivalently, we are interested in

$$\min_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right).$$

We introduce the associated Lagrangian function

$$f(w, q) = \sum_{i=1}^J \beta_i^2 \left(\frac{1}{w_{0,i}} + \frac{1}{w_{b,i}} \right) + \sum_{a=0}^K q_a \left(\sum_{i=1}^J w_{a,i} - u_a \right).$$

Taking the derivative with respect to $w_{0,i}$ and $w_{b,i}$ for the different values of i yields

$$w_{0,i} = \frac{|\beta_i|}{\sqrt{q_0}} \quad \text{and} \quad w_{b,i} = \frac{|\beta_i|}{\sqrt{q_b}}.$$

Summing over i implies that

$$\sqrt{q_0} = \frac{\sum_{i=1}^J |\beta_i|}{u_0} \quad \text{and} \quad \sqrt{q_b} = \frac{\sum_{i=1}^J |\beta_i|}{u_b},$$

and plugging the above in the expression of the weights yields

$$w_{0,i} = \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|} u_0 \quad \text{and} \quad w_{b,i} = \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|} u_b.$$

In particular,

$$\max_{\substack{w \in \Sigma_{(K+1)J} \\ \forall a, \sum_i w_{a,i} = u_a}} \frac{\Delta_b^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right)} = \frac{\Delta_b^2}{2\sigma^2 \left(\sum_{i=1}^J |\beta_i| \right)^2 \left(\frac{1}{u_0} + \frac{1}{u_b} \right)}, \quad (16)$$

yielding

$$\max_{\mathbf{w} \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_b^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right)} \leq \frac{1}{2\sigma^2 \left(\sum_{i=1}^J |\beta_i| \right)^2} \max_{u \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_b^2}{\left(\frac{1}{u_0} + \frac{1}{u_b} \right)}.$$

On the other hand, letting $w_{a,i} = \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|} u_a$ with $\sum_a u_a = 1$ we have

$$\frac{1}{2\sigma^2 \left(\sum_{i=1}^J |\beta_i| \right)^2} \max_{u \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_b^2}{\left(\frac{1}{u_0} + \frac{1}{u_b} \right)} \leq \max_{w \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_b^2}{2 \sum_{i=1}^J \beta_i^2 \left(\frac{\sigma^2}{w_{0,i}} + \frac{\sigma^2}{w_{b,i}} \right)},$$

showing that the two optimization programs are equivalent and that when denoting

$$(u_0^*, \dots, u_K^*) = \operatorname{argmax}_{u \in \Sigma_{(K+1)J}} \min_{b \neq 0} \frac{\Delta_b^2}{\left(\frac{1}{u_0} + \frac{1}{u_b} \right)},$$

one has

$$\forall a \in \{0, \dots, K\}, \forall i \leq J, w_{a,i}^* = u_a^* \frac{|\beta_i|}{\sum_{i=1}^J |\beta_i|}.$$

This corresponds to the optimal allocation strategy in the *active* mode. Recalling that when $\alpha = \beta$, the optimal weights for the *active* mode satisfy both $\mathcal{C}_{\text{prop}}$ and $\mathcal{C}_{\text{agnostic}}$ completes the proof. \square

E Asymptotic Optimality: Proof of Theorem 2

In this section we show that T-a-S with C-tracking and a certain threshold $\beta(t, \delta)$ is safely calibrated and asymptotically optimal. This is an important sanity check to validate our approach theoretically. Note that for the experimental validation we have explored a practically appealing variant of this algorithm: we employ an iterative scheme to approximate $w^*(\hat{\mu}(t))$, use D-tracking, and stylise the threshold.

Safe calibration follows from the definition of the recommendation rule (we report the answer $S_\beta(\hat{\mu}(t))$ at the empirical estimate $\hat{\mu}(t)$ of the bandit instance), together with the computation of the risk assessment $\hat{\delta}_t$. It does not depend on the sampling rule. Our confidence level $\hat{\delta}_t$ is obtained by inverting the threshold $\beta(t, \delta)$ at the GLR statistic (7). Safe calibration then follows from an anytime-valid GLR deviation inequality with boundary $\beta(t, \delta)$. We refer to [13, Proposition 23] for a boundary that is, in case of the ABC-S problem, of order $\ln \frac{1}{\delta} + K + 2J \cdot O(\ln \ln \frac{t}{\delta})$.

It remains to argue that the T-a-S sampling rule converges to the oracle weights. The original T-a-S proof for the BAI problem is due to [10, Theorem 14]. An upgrade to any single-answer problem, including our ABC-S, is due to [6]. For active mode, their theorem applies directly, while for agnostic mode it applies with the pair (I_t, X_t) regarded as the observation. We get:

Theorem 3 ([6, Theorems 7 and 10]). *For all ABC-S instances $\mu \in \mathcal{L}$ in active mode and agnostic mode, Track-and-Stop with C-tracking and stopping threshold $\beta(t, \delta) = \ln(t^2/\delta) + O(1)$ is δ -correct with asymptotically optimal sample complexity.*

In *proportional* mode, we have the additional constraint that the learner chooses its arm in response to seeing (but not controlling) the subpopulation I_t . Still, the tracking convergence result [6, Lemma 6] goes through, upon observing that the empirical distribution of I_t converges to α by the law of the large numbers, and hence our conditional tracking (see ‘‘sampling rule’’ in Section 3) adds the right conditional to the right marginal. All in all, the computed joint weights converge to the joint $w_{\text{prop}}^*(\mu)$, and tracking makes the sampling proportions also converge there.

We conclude with a remark on our use of D-tracking. Recall that D-tracking is the idea of advancing $N_a(t)$ towards t times the most current oracle weights, i.e. $tw_a^*(\hat{\mu}(t))$, while C-tracking makes $N_a(t)$ advance towards the sum of encountered oracle weights, i.e. $\sum_{s=1}^t w_a^*(\hat{\mu}(s))$. As argued in [7, Appendix E], D-tracking can fail to make $N_a(t)/t$ converge to $w_a^*(\mu)$. However, this requires that the maximiser of the lower bound problem is not unique at μ (as we are maximising a concave function, the set of maximisers is always convex). Here we argue that such a situation does not occur for the ABC-S problem. To see why, we argue that the lower bound objective, as a function of w , is strictly concave. It suffices to show this for the *active* mode problem, as the problems for the other modes are further constrained maximisation problems of the same objective.

Lemma 6. *Fix a bandit instance $\mu \in \mathcal{L}$. Let $\lambda \mapsto d(\mu_{a,j}, \lambda)$ be a strongly convex function for each arm a and subpopulation j . Then for the ABC-S problem with β such that $\beta_j \neq 0$ for all j , the oracle weights $w^*(\mu)$ are unique.*

Proof. Let $\mathbf{w}^*(\boldsymbol{\mu})$ be any oracle weights at $\boldsymbol{\mu}$. We will show the lower bound objective (4) is strictly concave as a function of \mathbf{w} around $\mathbf{w}^*(\boldsymbol{\mu})$, so that $\mathbf{w}^*(\boldsymbol{\mu})$ was in fact unique. For each $k > 0$, let $\boldsymbol{\lambda}^k$ be the minimiser in $\text{Alt}^k(\boldsymbol{\mu})$ of the weighted divergence in (4).

We perform a second-order Taylor expansion of the inner objective around $\boldsymbol{\lambda}^k$, which is a good approximation near $\boldsymbol{\lambda}^k$ (which is, after all, what matters when reasoning about \mathbf{w} near $\mathbf{w}^*(\boldsymbol{\mu})$). To this end, let us abbreviate the divergences, and their first and second derivatives in their second argument by $d_{a,j}^k := d(\mu_{a,j}, \lambda_{a,j}^k)$, $g_{a,j}^k := d'(\mu_{a,j}, \lambda_{a,j}^k)$ and $h_{a,j}^k := d''(\mu_{a,j}, \lambda_{a,j}^k)$, which all depend on $\boldsymbol{\lambda}^k$. A second-order Taylor expansion of the inner objective of (4) around $\boldsymbol{\lambda}^k$ yields

$$\inf_{\boldsymbol{\lambda} \in \text{Alt}^k} \sum_{a,j} w_{a,j} d(\mu_{a,j}, \lambda_{a,j}) \approx \sum_{a \in \{0,k\}, j} w_{a,j} \left(d_{a,j}^k - \frac{(g_{a,j}^k)^2}{2h_{a,j}^k} \right) + \frac{\left(\sum_j \beta_j \left(\frac{g_{0,j}^k}{h_{0,j}^k} - \frac{g_{k,j}^k}{h_{k,j}^k} \right) \right)^2}{2 \sum_{a \in \{0,k\}, j} \frac{\beta_j^2}{w_{a,j} h_{a,j}^k}}$$

where the optimiser is given by

$$\lambda_{a,j} = \lambda_{a,j}^k - \frac{g_{a,j}^k}{h_{a,j}^k} + \frac{\beta_j (\delta_{a=0} - \delta_{a=k}) \sum_j \beta_j \left(\frac{g_{0,j}^k}{h_{0,j}^k} - \frac{g_{k,j}^k}{h_{k,j}^k} \right)}{w_{a,j} h_{a,j}^k \sum_{a \in \{0,k\}, j} \frac{\beta_j^2}{w_{a,j} h_{a,j}^k}}.$$

Due to the last term, each of these is a strictly concave function of $w_{a,j}$ for $a \in \{0, k\}$ and all $j \leq J$ (here we use $\beta_j \neq 0$ and strong convexity $h_{a,j}^k > 0$).

Now we still need to consider the $\max_{\mathbf{w} \in \Sigma_{(K+1) \times J}} \min_{k > 0}$ problem. Let's convexify this for the inside finite min, and min-max swap to get a problem of the form $\min_{\mathbf{q} \in \Sigma_K} \max_{\mathbf{w} \in \Sigma_{(K+1) \times J}}$. Fixing the minimax outer strategy for \mathbf{q} , we find that \mathbf{w} is the maximiser of the strictly concave function

$$\mathbf{w} \mapsto \sum_{k > 0} q_k \left(\sum_{a \in \{0,k\}, j} w_{a,j} \left(d_{a,j}^k - \frac{(g_{a,j}^k)^2}{2h_{a,j}^k} \right) + \frac{\left(\sum_j \beta_j \left(\frac{g_{0,j}^k}{h_{0,j}^k} - \frac{g_{k,j}^k}{h_{k,j}^k} \right) \right)^2}{2 \sum_{a \in \{0,k\}, j} \frac{\beta_j^2}{w_{a,j} h_{a,j}^k}} \right)$$

To complete the argument, we argue that $q_k > 0$ for all $k > 0$, or, equivalently, that at \mathbf{w}^* the $\min_{k > 0}$ are all equalised. For if not, we can move mass from $w_{k,j}$ for the higher $k > 0$ to $w_{k',j}$ for the lower k' and increase the objective value. This then proves that $\mathbf{w}^*(\boldsymbol{\mu})$ is unique, as the objective function is bounded above by a strictly concave function itself maximised at $\mathbf{w} = \mathbf{w}^*(\boldsymbol{\mu})$. \square

F Algorithm Details

In this section we go into more details on the algorithm for each mode. Let us start with some notation. Let $\beta(\delta, t)$ be a threshold function. We denote the inverse of $\beta(t, \delta)$ in its second argument by

$$\beta^{-1}(t, \Lambda) = \min \{ \delta \in (0, 1) \mid \Lambda \geq \beta(t, \delta) \}.$$

We extend the definition of the GLR statistic to sample frequencies \mathbf{w} and bandit $\boldsymbol{\mu}$ by

$$\Lambda(\mathbf{w}, \boldsymbol{\mu}) := \min_{b \neq 0} \inf_{\boldsymbol{\lambda} \in \mathcal{L}: \lambda_0 = \lambda_b} \sum_{a \in \{0,b\}} \sum_{i=1}^J w_{a,i} d(\mu_{a,i}, \lambda_{a,i}),$$

so that the original definition (7) is $\Lambda(t) = \Lambda(\mathbf{N}(t)/t, \hat{\boldsymbol{\mu}}(t))$. For any $\boldsymbol{\mu}$, we denote by $\nabla_{\mathbf{w}} \Lambda(\mathbf{w}, \boldsymbol{\mu})$ any sub-gradient of $\mathbf{w} \mapsto \Lambda(\mathbf{w}, \boldsymbol{\mu})$. We can obtain one such a sub-gradient by letting $(b, \boldsymbol{\lambda})$ be any minimiser of $\Lambda(\mathbf{w}, \boldsymbol{\mu})$, and constructing the vector with entry (a, i) given by

$$(a, i) \mapsto \begin{cases} d(\mu_{a,i}, \lambda_{a,i}) & \text{if } a \in \{0, b\} \\ 0 & \text{otherwise} \end{cases}.$$

Our algorithms will make use of an online learning method (called \mathcal{A} below) for linear losses defined on the simplex. This online learning task is known as the Hedge or Experts setting in the literature.

We will make use of AdaHedge [5], as it adapts automatically to the range of the losses and does not require tuning. Our methods for the active, proportional and agnostic modes are displayed as Algorithms 1, 2 and 3. Each algorithm consists of a Forced Exploration part, which serves to ensure that the empirical estimate of the bandit model converges, i.e. $\hat{\boldsymbol{\mu}}(t) \rightarrow \boldsymbol{\mu}$. By forcing exploration sub-linearly often, the main term in the sample complexity is unaffected asymptotically. Each algorithm further makes use of online learning to compute $\boldsymbol{w}^*(\boldsymbol{\mu})$. In the notation of this section, we have

$$\boldsymbol{w}^*(\boldsymbol{\mu}) = \operatorname{argmax}_{\boldsymbol{w} \in \mathcal{C}} \Lambda(\boldsymbol{w}, \boldsymbol{\mu}) .$$

Our approach to learning $\boldsymbol{w}^*(\boldsymbol{\mu})$ is to perform gradient steps on the plug-in loss function $\boldsymbol{w} \mapsto -\Lambda(\boldsymbol{w}, \hat{\boldsymbol{\mu}}(t))$. It is in the convex domain $\mathcal{C} \subseteq \Sigma_{(K+1) \times J}$ that we see the main difference between the three modes. Recall from Section 2.3 that in the active mode \boldsymbol{w} is not constrained further, in the proportional mode the subpopulation marginal of \boldsymbol{w} must equal $\boldsymbol{\alpha}$, i.e. $\langle \mathbf{1}, \boldsymbol{w} \rangle = \boldsymbol{\alpha}$, and in the agnostic mode \boldsymbol{w} must be the independent product $\boldsymbol{w} = \boldsymbol{v}\boldsymbol{\alpha}$ of some arm marginal $\boldsymbol{v} \in \Sigma_{K+1}$ and the subpopulation frequencies $\boldsymbol{\alpha}$. We hence need to design online learners for each of the three \mathcal{C} . In the active case, we have one learner \mathcal{A} that learns the full joint $\boldsymbol{w}^*(a, j)$ directly, in the proportional case we use one learner \mathcal{A}_j for each subpopulation $j \in [J]$ to learn the conditional distribution $\boldsymbol{w}^*(a|j)$, and in the agnostic case we again use one learner to learn the common marginal $\boldsymbol{w}^*(a)$. This difference is reflected in the loss function used in each mode, and hence in the gradient that is fed to each learner. In the active case we use the full $(K+1) \times J$ gradients

$$\boldsymbol{\ell}_t^{\text{active}} := -\nabla_{\boldsymbol{w}} \Lambda(\boldsymbol{w}_t, \hat{\boldsymbol{\mu}}(t)) .$$

In the proportional case we have $\boldsymbol{w}(a, i) = \boldsymbol{w}(a|i)\alpha_i$, and by the chain rule we hence have gradients

$$\boldsymbol{\ell}_t^{i, \text{proportional}} := -\nabla_{\boldsymbol{w}(a|i)} \Lambda(\boldsymbol{w}_t, \hat{\boldsymbol{\mu}}(t)) = -\alpha_i \nabla_{\boldsymbol{w}} \Lambda(\boldsymbol{w}_t, \hat{\boldsymbol{\mu}}(t)) \boldsymbol{e}_i .$$

Finally, in the agnostic case we have $\boldsymbol{w}(a, i) = \boldsymbol{w}(a)\alpha_i$, and again by the chain rule we have

$$\boldsymbol{\ell}_t^{\text{agnostic}} := -\nabla_{\boldsymbol{w}(a)} \Lambda(\boldsymbol{w}_t, \hat{\boldsymbol{\mu}}(t)) = -\nabla_{\boldsymbol{w}} \Lambda(\boldsymbol{w}_t, \hat{\boldsymbol{\mu}}(t)) \boldsymbol{\alpha} .$$

Run Time In each of the three modes, our algorithms evaluate Λ for the confidence in the recommendation, compute one sub-gradient of Λ for the loss function, and spend $O(K \times J)$ time bookkeeping. Evaluation and sub-gradient computation for Λ boil down to solving a convex minimisation problem with an equality constraint. We use Newton's method with backtracking line search to find the minimiser given b . Each Newton iteration takes $O(J^2)$ time (recall that only 2 arms are involved), and we never needed more than 40. Doing this K times for the explicit minimum over b yields a total per iteration run time of $O(KJ^2)$.

Algorithm 1 Algorithm for Active Mode.

Require: Online learner \mathcal{A} for $(K + 1) \times J$ experts.

```
for  $t = 1, 2, \dots$  do
  if any pair  $(a, i)$  has  $N_{a,i}(t-1) \leq \sqrt{t}$  then
    Pick  $A_t, I_t$  any such pair
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    ▷ Forced Exploration
  else
    Get  $w_t$  from online learner  $\mathcal{A}$ 
    Pick  $(A_t, I_t) \in \operatorname{argmin}_{a,i} N_{a,i}(t-1) - tw_t(a, i)$ 
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    Send loss vector  $\ell_t = -\nabla_w \Lambda(w_t, \hat{\mu}(t))$  to  $\mathcal{A}$ 
    ▷ Direct Tracking
  end if
  Recommend  $\hat{S}_t = \mathcal{S}_\beta(\hat{\mu}(t))$  at confidence  $\delta_t = \beta^{-1}(t, \Lambda(N(t)/t, \hat{\mu}(t)))$ .
end for
```

Algorithm 2 Algorithm for Proportional Mode.

Require: J online learners $\mathcal{A}^{(1)}, \dots, \mathcal{A}^{(J)}$ for $(K + 1)$ experts each.

```
for  $t = 1, 2, \dots$  do
  See  $I_t \sim \alpha$ .
  if any arm  $a$  has  $N_{a, I_t}(t-1) \leq \sqrt{\sum_a N_{a, I_t}(t-1)}$  then
    Pick  $A_t$  any such arm
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    ▷ Forced Exploration
  else
    Get  $w_t^{(j)}$  from each online learner  $\mathcal{A}^{(j)}$ 
    Pick  $A_t \in \operatorname{argmin}_a N_{a, I_t}(t-1) - tw_t^{(I_t)}(a)$ 
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    ▷ Direct Tracking
    For  $j \in [J]$ , send loss vector  $\ell_t^{(j)} = -\alpha_j \nabla_w \Lambda([\alpha_1 w_t^{(1)} \dots \alpha_J w_t^{(J)}], \hat{\mu}(t)) e_j$  to  $\mathcal{A}^{(j)}$ 
  end if
  Recommend  $\hat{S}_t = \mathcal{S}_\beta(\hat{\mu}(t))$  at confidence  $\delta_t = \beta^{-1}(t, \Lambda(N(t)/t, \hat{\mu}(t)))$ .
end for
```

Algorithm 3 Algorithm for Agnostic Mode.

Require: Online learner \mathcal{A} for $(K + 1)$ experts.

```
for  $t = 1, 2, \dots$  do
  if any arm  $a$  has  $\sum_{j=1}^J N_{a,j}(t-1) \leq \sqrt{t}$  then
    Pick  $A_t$  any such arm
    See  $I_t \sim \alpha$ .
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    ▷ Forced Exploration
  else
    Get  $w_t$  from online learner  $\mathcal{A}$ 
    Pick  $A_t \in \operatorname{argmin}_a \sum_{j=1}^J N_{a,j}(t-1) - tw_t(a)$ 
    Obtain sample  $X_t$  from  $\nu_{A_t, I_t}$ .
    See  $I_t \sim \alpha$ .
    Send loss vector  $\ell_t = -\nabla_w \Lambda(w_t \alpha^\top, \hat{\mu}(t)) \alpha$  to  $\mathcal{A}$ .
    ▷ Direct Tracking
  end if
  Recommend  $\hat{S}_t = \mathcal{S}_\beta(\hat{\mu}(t))$  at confidence  $\delta_t = \beta^{-1}(t, \Lambda(N(t)/t, \hat{\mu}(t)))$ .
end for
```

G Details of A/B/n experiment

	1	2	3	4
0	0.0296	0.0372	0.0588	0.0620
1	0.0300	0.0373	0.0596	0.0626
2	0.0295	0.0373	0.0591	0.0630

(a) Estimated click probabilities for the different options and seasons

	1	2	3	4
0	0.0719	0.1222	0.1311	0.1238
1	0.0215	0.0475	0.0340	0.0352
2	0.0614	0.1060	0.1377	0.1078

(c) w^* for *active*, $T^* = 3.98 \cdot 10^6$

	1	2	3	4
0	0.0814	0.1269	0.1289	0.0889
1	0.0230	0.0355	0.0457	0.0277
2	0.0914	0.1326	0.1067	0.1113

(e) w^* for *proportional*, $T^* = 4.06 \cdot 10^6$

	1
0	0.44482
1	0.11111
2	0.44406

(g) w^* for *agnostic*, $T^* = 4.61 \cdot 10^6$

	1
0	0.44480
1	0.11111
2	0.44409

(i) w^* for *oblivious*, $T^* = 4.63 \cdot 10^6$

	1
0	0.500
1	0.076
2	0.424

(k) Sampling proportions BC-ABC

	1
0	0.1958
1	0.2950
2	0.2813
3	0.2279

(b) frequency and importance vector ($\alpha = \beta$)

	1	2	3	4
0	0.0740	0.1246	0.1476	0.1226
1	0.0108	0.0179	0.0214	0.0178
2	0.0727	0.1228	0.1460	0.1218

(d) Sampling proportions *active*

	1	2	3	4
0	0.0912	0.1374	0.1307	0.1056
1	0.0148	0.0223	0.0214	0.0173
2	0.0898	0.1353	0.1293	0.1048

(f) Sampling proportions *proportional*

	1
0	0.4648
1	0.0766
2	0.4587

(h) Sampling proportions *agnostic*

	1
0	0.4647
1	0.0764
2	0.4589

(j) Sampling proportions *oblivious*

Figure 4: Summary of oracle weights and sampling proportions for A/B/n experiment