



HAL
open science

Prior-Dependent Allocations for Bayesian Fixed-Budget Best-Arm Identification in Structured Bandits

Nicolas Nguyen, Imad Aouali, András György, Claire Vernade

► **To cite this version:**

Nicolas Nguyen, Imad Aouali, András György, Claire Vernade. Prior-Dependent Allocations for Bayesian Fixed-Budget Best-Arm Identification in Structured Bandits. 2024. hal-04629228

HAL Id: hal-04629228

<https://hal.science/hal-04629228v1>

Preprint submitted on 28 Jun 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Prior-Dependent Allocations for Bayesian Fixed-Budget Best-Arm Identification in Structured Bandits

Nicolas Nguyen^{*1}, Imad Aouali^{*2,3}, András György⁴, and Claire Vernade¹

¹*University of Tübingen*

²*Criteo AI Lab*

³*ENSAE CREST, IP Paris*

⁴*Google Deepmind*

June 28, 2024

Abstract

We study the problem of Bayesian fixed-budget best-arm identification (BAI) in structured bandits. We propose an algorithm that uses fixed allocations based on the prior information and the structure of the environment. We provide theoretical bounds on its performance across diverse models, including the first prior-dependent upper bounds for linear and hierarchical BAI. Our key contribution is introducing new proof methods that result in tighter bounds for multi-armed BAI compared to existing methods. We extensively compare our approach to other fixed-budget BAI methods, demonstrating its consistent and robust performance in various settings. Our work improves our understanding of Bayesian fixed-budget BAI in structured bandits and highlights the effectiveness of our approach in practical scenarios.

1 Introduction

Best arm identification (BAI) addresses the challenge of finding the optimal arm in a bandit environment (Lattimore and Szepesvári, 2020), with wide-ranging applications in online advertising, drug discovery or hyperparameter tuning. BAI is commonly approached through two primary paradigms: *fixed-confidence* and *fixed-budget*. In the fixed-confidence setting (Even-Dar et al., 2006; Kaufmann et al., 2016), the objective is to find the optimal arm with a pre-specified confidence level. Conversely, fixed-budget BAI (Audibert et al., 2010; Karnin et al., 2013; Carpentier and Locatelli, 2016) involves identifying the optimal arm within a fixed number of observations. Within this fixed-budget context, two main metrics are used: the probability of error (PoE) (Audibert et al., 2010; Karnin et al., 2013; Carpentier and Locatelli, 2016)—the likelihood of incorrectly identifying the optimal arm—and the simple regret (Bubeck et al., 2009; Russo, 2016; Komiyama et al., 2023)—the expected performance disparity between the chosen and the optimal arm. We focus on PoE minimization in fixed-budget BAI.

Existing algorithms for PoE minimization in fixed-budget BAI are largely frequentist and often employ elimination strategies. Bayesian approaches have predominantly focused on the minimization of the simple regret (Komiyama et al., 2023; Azizi et al., 2023), or were studied under a frequentist lens (Hoffman et al., 2014), which do not capture the advantages of knowing informative priors. It was only recently that Atsidakou et al. (2022) introduced a Bayesian version of the well-known Sequential Halving (SH) algorithm (Karnin et al., 2013), offering a prior-dependent bound on the probability of error in multi-armed bandits (MAB), albeit under certain limiting assumptions on the prior. Their proofs are still largely influenced by frequentist approaches and come with strong constraints.

Several recent works have shed new light on adaptive methods for frequentist fixed-budget BAI. For instance, Qin (2022); Degenne (2023); Wang et al. (2023) examined whether adaptive algorithms can consistently surpass the best static algorithm for any bandit instance. Remarkably, Degenne (2023) established the

^{*}Equal contribution. Correspondence to: Nicolas Nguyen <nicolas.nguyen@uni-tuebingen.de>.

absence of such universally superior adaptive algorithms in several BAI problems, including Gaussian BAI. Wang et al. (2023) demonstrated that in Bernoulli BAI with two arms, no algorithm consistently outperforms uniform sampling. Inspired by these recent results, we develop a method based on fixed and non-adaptive allocations in the Bayesian setting. These allocations leverage both prior information and the structure of the environment. As demonstrated in our experiments in Section 5, our static but prior-informed algorithm can outperform adaptive baselines. Moreover, our proofs incorporate fully Bayesian techniques, diverging from existing works. This novel technical approach not only produces a tighter upper bound but also applies under milder assumptions.

As a motivating example, consider a scenario with three arms ($K = 3$), where the information *a priori* strongly suggests that either of the first two arms is more likely to be optimal than the third one. A pivotal question arises: what strategic approach should be employed to allocate resources (or budget) to the seemingly suboptimal third arm? Furthermore, if greater confidence is placed on the first arm compared to the second, what is the optimal budget distribution between them? This situation underlines a fundamental challenge not directly addressed by the frequentist approach, which lacks knowledge about the bandit instance prior to interaction.

Relation to prior works. While Atsidakou et al. (2022) considers prior information, their method does not exploit it to its full potential. To maintain adaptivity, they impose restrictive assumptions on the prior. Moreover, their results are only valid for a specific budget allocation, while ours are applicable for any fixed and non-adaptive allocation rule. This facilitates the creation of ad-hoc allocation strategies, informed and guided by our theoretical results. Additionally, a particularly relevant application of prior information is found in structured bandit problems, such as linear bandits (Abbasi-Yadkori et al., 2011; Hoffman et al., 2014; Katz-Samuels et al., 2020; Azizi et al., 2021) and hierarchical bandits (Hong et al., 2022b; Aouali et al., 2023), where arm rewards are determined by underlying *latent* parameters. Our approach captures the structure of these problems as reflected in the prior, leading to more efficient exploration thanks to arm correlations. This aspect of our work also extends beyond the scope of Atsidakou et al. (2022), which primarily addressed MAB.

Contributions. **1)** We present and analyze Prior-Informed BAI, PI-BAI, a fixed-budget BAI algorithm that leverages prior information for efficient exploration. Our main contributions include establishing upper prior-dependent bounds on its expected PoE in multi-armed, linear, and hierarchical bandits. Specifically, in the MAB setting, our upper bound is smaller and is valid under milder assumptions on the prior. **2)** The proof techniques developed for PI-BAI provide a fully Bayesian perspective, significantly diverging from existing methodologies that rely on frequentist proofs. This allows a more comprehensive framework for understanding and analyzing Bayesian BAI algorithms, while also enabling us to relax previously held assumptions. **3)** Our algorithms and proof techniques are naturally applicable to structured problems, such as linear and hierarchical bandits, leading to the first Bayesian BAI algorithm with a prior-dependent PoE bound in these settings. **4)** We empirically evaluate PI-BAI and its variants in various numerical setups. Our experiments on synthetic and real-world data show the generality of PI-BAI and highlight its good performance.

2 Background

Notation. Let Δ_K be the K -simplex and $\Delta_K^\dagger = \{\omega \in \Delta_K; \omega_i > 0, \forall i \in [K]\}$. For any positive-definite matrix $A \in \mathbb{R}^{d \times d}$ and vector $v \in \mathbb{R}^d$, we define $\|v\|_A = \sqrt{v^\top A v}$. Also, $\lambda_1(A)$ and $\lambda_d(A)$ denote the maximum and minimum eigenvalues of A , respectively. We denote by e_i the i -th vector of the canonical basis.

We consider a scenario involving K arms. In each round $t \in [n]$, the agent selects an arm $A_t \in [K]$, and then receives a stochastic reward $Y_t \sim P(\cdot; \theta, A_t)$, where θ is the *unknown* parameter vector and $P(\cdot; \theta, i)$ is the *known* reward distribution of arm i , given θ . We denote by $r(i; \theta) = \mathbb{E}_{Y \sim P(\cdot; \theta, i)} [Y]$ the mean reward of arm i , given θ . We adopt the Bayesian view where θ is assumed to be sampled from a *known* prior distribution P_0 . Given a bandit instance characterized by $\theta \sim P_0$, the goal is to find the (unique) optimal arm $i_*(\theta) = \operatorname{argmax}_{i \in [K]} r(i; \theta)$ by interacting with the bandit instance for a fixed-budget of n rounds. These interactions are summarized by *the history* $H_n = \{(A_t, Y_t)\}_{t \in [n]}$, and we let $J_n \in [K]$ be the arm selected by

the agent after n rounds. In this Bayesian setting, [Atsidakou et al. \(2022\)](#) introduced the expected PoE as

$$\mathcal{P}_n = \mathbb{E} [\mathbb{P}(J_n \neq i_*(\theta) \mid \theta)], \quad (1)$$

a Bayesian metric that corresponds to the average PoE across all bandit instances sampled from the prior, $\theta \sim P_0$. This is different from the frequentist counterpart where the performance is assessed for a single instance θ .

2.1 Multi-Armed Bandit

In this setting, each component of $\theta = (\theta_i)_{i \in [K]}$ is sampled independently from the prior distribution. We focus on the Gaussian case where $\theta_i \sim \mathcal{N}(\mu_{0,i}, \sigma_{0,i}^2)$, with $\mu_{0,i}$ and $\sigma_{0,i}^2$ being the *known* prior reward mean and variance for arm i . Then, given θ , the reward distribution of arm i is $\mathcal{N}(\theta_i, \sigma^2)$ where σ^2 is the (known) observation noise variance¹,

$$\begin{aligned} \theta_i &\sim \mathcal{N}(\mu_{0,i}, \sigma_{0,i}^2) && \forall i \in [K] \\ Y_t \mid \theta, A_t &\sim \mathcal{N}(\theta_{A_t}, \sigma^2) && \forall t \in [n]. \end{aligned} \quad (2)$$

Under (2), the posterior distribution of θ_i given H_n is a Gaussian $\mathcal{N}(\hat{\mu}_{n,i}, \hat{\sigma}_{n,i}^2)$ ([Bishop, 2006](#)), where

$$\hat{\sigma}_{n,i}^{-2} = \frac{1}{\sigma_{0,i}^2} + \frac{n_i}{\sigma^2}, \quad \hat{\mu}_{n,i} = \hat{\sigma}_{n,i}^2 \left(\frac{\mu_{0,i}}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma^2} \right), \quad (3)$$

where $\mathcal{T}_i = \{t \in [n], A_t = i\}$ is the set of rounds when arm i is chosen, $n_i = |\mathcal{T}_i|$ is the number of times arm i is chosen and $B_{n,i} = \sum_{t \in \mathcal{T}_i} Y_t$ is the sum of rewards of arm i . Here the *mean posterior reward* is $\mathbb{E}[r(i; \theta) \mid H_n] = \hat{\mu}_{n,i}$.

2.2 Linear Bandit

One major drawback of model (2) is that it is not able to model situations where arms are dependent, thus leading to suboptimal exploration.

In linear bandits ([Abbasi-Yadkori et al., 2011](#)), arms share a common low-dimensional representation through parameter $\theta \in \mathbb{R}^d$. We denote $\mathcal{X} = \{x_1, \dots, x_K\}$ the set of arms where each $x_i \in \mathbb{R}^d$. We focus on the Gaussian case where θ is sampled from a Gaussian distribution with mean $\mu_0 \in \mathbb{R}^d$ and covariance matrix $\Sigma_0 \in \mathbb{R}^{d \times d}$. Given θ , the reward distribution of arm i is Gaussian with mean $r(i; \theta) = x_i^\top \theta$ and variance σ^2 ,

$$\begin{aligned} \theta &\sim \mathcal{N}(\mu_0, \Sigma_0) \\ Y_t \mid \theta, A_t &\sim \mathcal{N}(\theta^\top x_{A_t}, \sigma^2) && \forall t \in [n]. \end{aligned} \quad (4)$$

Similarly to (2), this model offers closed-form formulas, where the posterior of θ given the history H_n containing n_i samples from arm i is a Gaussian $\mathcal{N}(\hat{\mu}_n, \hat{\Sigma}_n)$:

$$\hat{\Sigma}_n^{-1} = \Sigma_0^{-1} + \sigma^{-2} \sum_{i \in [K]} n_i x_i x_i^\top, \quad \hat{\mu}_n = \hat{\Sigma}_n (\Sigma_0^{-1} \mu_0 + \sigma^{-2} B_n), \quad (5)$$

where $B_n = \sum_{t \in [n]} Y_t x_{A_t}$. The mean posterior reward of arm i is given by $\mathbb{E}[r(i; \theta) \mid H_n] = \hat{\mu}_n^\top x_i$. Note that the MAB (2) can be recovered from (4) when $x_i = e_i \in \mathbb{R}^K$ and $\Sigma_0 = \text{diag}(\sigma_{0,i}^2)_{i \in [K]}$.

2.3 Hierarchical Bandit

Another practical model that captures arm correlations is the hierarchical (or *mixed-effect*) model ([Bishop, 2006](#); [Wainwright et al., 2008](#); [Hong et al., 2022b](#); [Aouali et al., 2023](#)), defined in the Gaussian case as

$$\begin{aligned} \mu &\sim \mathcal{N}(\nu, \Sigma) \\ \theta_i &\sim \mathcal{N}(b_i^\top \mu, \sigma_{0,i}^2) && \forall i \in [K] \\ Y_t \mid \mu, \theta, A_t &\sim \mathcal{N}(\theta_{A_t}, \sigma^2) && \forall t \in [n]. \end{aligned} \quad (6)$$

¹Arm-dependent observation noise variances could be used but we choose to keep the notation simple.

This generative model reads as follows. First, $\mu = (\mu_\ell)_{\ell \in [L]}$ is an unknown latent vector composed of L effects μ_ℓ and it is sampled from a multivariate Gaussian with mean $\nu \in \mathbb{R}^L$ and covariance $\Sigma \in \mathbb{R}^{L \times L}$. Then, given μ , the mean rewards θ_i are independently sampled as $\theta_i \sim \mathcal{N}(b_i^\top \mu, \sigma_{0,i}^2)$, where $(b_i)_{i \in [K]}$ represent *known mixing weights*. In particular, $b_i^\top \mu$ creates a linear mixture of the L effects, with $b_{i,\ell}$ indicating a known score that quantifies the association between arm i and the effect ℓ . Concrete examples of $b_{i,\ell}$ are provided in Appendix A.1. Note that arm correlations arise because θ_i are derived from the same effect parameter μ . Finally, given μ and θ , the reward distribution of arm i is similar to the MAB (2) and writes $\mathcal{N}(\theta_i, \sigma^2)$.

With abuse of notation, the *effect posterior* distribution $Q_n(\mu) = \mathbb{P}(\mu | H_n)$ induces a *conditional arm posterior* distribution for each arm i , $P_{n,i}(\theta_i | \mu) = \mathbb{P}(\theta_i | H_n, \mu)$. Then, the *marginal arm posterior* density can be computed by marginalizing over Q_n such as $\mathbb{P}(\theta_i | H_n) = \mathbb{E}_{\mu \sim Q_n} [P_{n,i}(\theta_i | \mu)]$. Therefore, despite the hierarchical structure, these distributions can be derived in closed-form². First, the effect posterior writes $Q_n(\mu) = \mathcal{N}(\check{\mu}_n, \check{\Sigma}_n)$ with

$$\check{\Sigma}_n^{-1} = \Sigma^{-1} + \sum_{i \in [K]} \frac{n_i}{n_i \sigma_{0,i}^2 + \sigma^2} b_i b_i^\top, \quad \check{\mu}_n = \check{\Sigma}_n \left(\Sigma^{-1} \nu + \sum_{i \in [K]} \frac{B_{n,i}}{n_i \sigma_{0,i}^2 + \sigma^2} b_i \right). \quad (7)$$

Then, given $\mu \sim Q_n$, the conditional arm posteriors are $P_{n,i}(\theta_i | \mu) = \mathcal{N}(\tilde{\mu}_{n,i}, \tilde{\sigma}_{n,i}^2)$, where

$$\tilde{\sigma}_{n,i}^{-2} = \frac{1}{\sigma_{0,i}^2} + \frac{n_i}{\sigma^2}, \quad \tilde{\mu}_{n,i} = \tilde{\sigma}_{n,i}^2 \left(\frac{\mu^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma^2} \right). \quad (8)$$

Finally, combining (7) and (8) leads to the marginal arm posterior $\mathbb{P}(\theta_i | H_n) = \mathcal{N}(\hat{\mu}_{n,i}, \hat{\sigma}_{n,i}^2)$ where

$$\hat{\sigma}_{n,i}^2 = \tilde{\sigma}_{n,i}^2 + \frac{\tilde{\sigma}_{n,i}^4}{\sigma_{0,i}^4} b_i^\top \check{\Sigma}_n b_i, \quad \hat{\mu}_{n,i} = \tilde{\sigma}_{n,i}^2 \left(\frac{\check{\mu}_n^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma^2} \right), \quad (9)$$

and we have that $\mathbb{E}[r(i; \theta) | H_n] = \hat{\mu}_{n,i}$.

Link to linear bandit. (6) is a special case of a linear bandit (4), as can be seen by realizing that $\theta_i = b_i^\top \mu + \eta_i$ where $\eta_i \sim \mathcal{N}(0, \sigma_{0,i}^2)$ and the η_i are independent of μ . Hence, (6) can be rewritten by replacing ν with $\bar{\nu} \in \mathbb{R}^{L+K}$ defined as $\bar{\nu}^\top = (\nu^\top, 0, \dots, 0)$ and Σ with a block-diagonal matrix $\bar{\Sigma} = \langle \Sigma, \sigma_{0,1}^2, \dots, \sigma_{0,K}^2 \rangle$ with new actions $\bar{b}_i \in \mathbb{R}^{L+K}$ defined as $\bar{b}_i^\top = (b_i^\top, e_i^\top)$, leading to a linear bandit

$$\bar{\mu} \sim \mathcal{N}(\bar{\nu}, \bar{\Sigma}) \\ Y_t | \bar{\mu}, A_t \sim \mathcal{N}(\bar{b}_{A_t}^\top \bar{\mu}, \sigma^2) \quad \forall t \in [n]. \quad (10)$$

Note that reinterpreting hierarchical bandits in this way does not lead to practical benefits. In contrast, adhering to the initial formulation in (6) and the subsequent derivations in (9) is more computationally efficient. Indeed, this is one of the motivations behind the concept of hierarchical bandits.

3 Algorithm and Error Bounds

PI-BAI takes as input a budget n and an arbitrary vector of *allocation weights* $\omega = (\omega_i)_{i \in [K]} \in \Delta_K$. Then, it collects $n_i = \lfloor \omega_i n \rfloor$ samples for each arm i , and finally returns the arm with the highest posterior mean, defined as $J_n = \operatorname{argmax}_{i \in [K]} \mathbb{E}[r(i; \theta) | H_n]$. PI-BAI is described in Algorithm 1 where we make its dependence on the allocation weights ω explicit and call it PI-BAI(ω).

Our algorithm consists in coupling PI-BAI to a well-chosen allocation vector ω that depends on the prior. We will discuss the allocation strategies further below, but we first give theoretical guarantees that hold for any choice of fixed ω . This idea has the benefit of its versatility, as it naturally generalizes to structured bandit settings such as the linear or hierarchical problems defined above. The structure is a direct prior information and is taken into account in the computation of the allocation weights as well as in the posterior updates. Similarly and despite additional technicality, our novel proof scheme (see Section 4) is preserved across all settings and allows us to state our main theorems below.

²Full derivations are in Appendix C.

Algorithm 1 Prior-Informed BAI: PI-BAI(ω)

Input: Budget n , allocation weights $\omega \in \Delta_K$.

for $i = 1, \dots, K$ **do**

 Get $n_i = \lfloor \omega_i n \rfloor$ samples of arm i
 Compute mean posterior reward $\mathbb{E}[r(i; \theta) \mid H_n]$

Set $J_n = \operatorname{argmax}_{i \in [K]} \mathbb{E}[r(i; \theta) \mid H_n]$.

3.1 PoE Bounds for Multi-Armed Bandits

Theorem 3.1 presents an upper bound on the expected PoE of PI-BAI(ω) for MAB (2). The bound depends on the prior and allocation weights³ $\omega \in \Delta_K^+$.

Theorem 3.1 (Upper bound for multi-armed bandit). *For all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the MAB problem (2) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i, j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{2(\sigma_{0,i}^2 + \sigma_{0,j}^2)}}}{\sqrt{1 + n\phi_{i,j}}},$$

where $\phi_{i,j} = \Omega(1)$, and it depends on the prior parameters and allocation weights. In particular, $\mathcal{P}_n = \mathcal{O}(1/\sqrt{n})$. Full expressions and proofs are given in Appendix C.2.

The $\mathcal{O}(1/\sqrt{n})$ bound contrasts with frequentist results where PoE is typically $\mathcal{O}(e^{-n/f(\theta)})$, where $f(\theta)$ is a complexity measure that depends on the *fixed* bandit instance θ (Audibert et al., 2010; Carpentier and Locatelli, 2016). However, this $\mathcal{O}(1/\sqrt{n})$ rate is not surprising for the *expected* PoE. For instance, when $K = 2$, Atsidakou et al. (2022) states the existence of a prior distribution for which the expected PoE is lower bounded by $\mathcal{O}(1/\sqrt{n})$, and our result asymptotically matches this lower bound. To give more intuition, notice that averaging the frequentist PoE bound $\mathcal{O}(e^{-n/f(\theta)})$ under a Gaussian prior $\theta_i \sim \mathcal{N}(\mu_{0,i}, \sigma_{0,i}^2)$ leads to a $\mathcal{O}(1/\sqrt{n})$ rate for the expected PoE. Indeed, this idea was employed by Atsidakou et al. (2022) to achieve their $\mathcal{O}(1/\sqrt{n})$ rate. However, while we achieve similar asymptotic rates, our proof differs significantly as we do not average the frequentist bound. Beyond asymptotic behavior, our bound also captures the structure of the prior. In particular, if the prior is informative, either with small prior variances $\sigma_{0,i}^2 \rightarrow 0$ or large prior gaps $|\mu_{0,i} - \mu_{0,j}| \rightarrow \infty$, then $\mathcal{P}_n \rightarrow 0$ for any fixed allocation weights $\omega \in \Delta_K^+$. After interpreting our bound, we now compare it to that of another Bayesian algorithm, **BayesElim** (Atsidakou et al., 2022), in the simpler setting where their results are valid, that is, when the prior variances are homogeneous, $\sigma_{0,i}^2 = \sigma_0^2$. Their bound reads

$$\mathcal{P}_n^{\text{BayesElim}} \leq \sum_{\substack{i, j \in [K] \\ i \neq j}} \log_2(K) \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{4\sigma_0^2}}}{\sqrt{1 + \frac{n\sigma_0^2}{K \log_2(K) \sigma^2}}}.$$

Theorem 3.1 in this setting even with the simplest allocations $\omega_i = \frac{1}{K}$ simplifies to the similar expression

$$\mathcal{P}_n^{\text{PI-BAI}} \leq \sum_{\substack{i, j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{4\sigma_0^2}}}{\sqrt{1 + \frac{n\sigma_0^2}{K\sigma^2}}}.$$

By omitting elimination phases, we gain roughly a factor $\log_2^{3/2}(K)$ over the bound of Atsidakou et al. (2022) (highlighted in blue). This difference makes our bound smaller even in their setting with homogeneous prior variances and choosing uniform allocation weights for PI-BAI.

³For technical reasons we only allow positive allocation weights $\omega \in \Delta_K^+$.

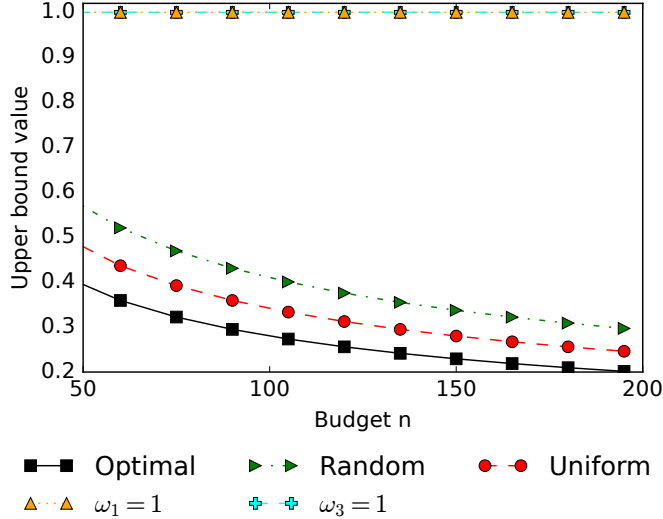


Figure 1: Upper bound value of our method instantiated with various weights compared to the upper bound of BayesElim (Atsidakou et al., 2022) in two settings.

We numerically compare these two bounds on the 3-armed bandit example described in Section 1, where one arm is a priori suboptimal, and one of the other two is a priori optimal, $\mu_0 = (1, 1.9, 2)$. We consider two scenarios: one with homogeneous prior variances ($\sigma_{0,i} = 0.3$ for all $i \in [3]$) and another with heterogeneous prior variances, $(\sigma_{0,1}, \sigma_{0,2}, \sigma_{0,3}) = (0.1, 0.5, 0.5)$. Since BayesElim’s bound does not handle heterogeneous prior variances, we use an average prior variance $\sigma_0^2 = \frac{1}{K} \sum_{i \in [K]} \sigma_{0,i}^2$ for comparison. As predicted, Figure 1 shows that the value of the upper bound is much lower for PI-BAI for optimized, uniform and random allocation weights. We also plot the upper bound of PI-BAI instantiated with a heuristic weight that favors higher prior means and higher prior variance arms, that is, $\omega_i^h = \frac{\mu_{0,i} \sigma_{0,i}}{\sum_{k \in [K]} \mu_{0,k} \sigma_{0,k}}$.

3.2 PoE Bounds for Structured Bandits

Importantly, our analysis extends to the linear and hierarchical bandits in (4) and (6). Theorems 3.2 and 3.3 provide upper bounds on the PoE of PI-BAI(ω) in these settings. To the best of our knowledge, these are the first *prior-dependent* bounds for fixed-budget Bayesian BAI in these settings.

Theorem 3.2 (Upper bound for linear bandit). *Assume that $x_i \neq x_j$ for any $i \neq j$, and that there exists $S > 0$ such that $\|x\|_2^2 \leq S$ for any $x \in \mathcal{X}$. Then, for all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the linear bandit problem (4) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{c_{i,j}}{\|x_i - x_j\|_{\Sigma_n}^2}}},$$

where $c_{i,j} = \Omega(1)$, $\lambda_1(\hat{\Sigma}_n) = \mathcal{O}(1/n)$, and they depend on both prior parameters and allocation weights. In particular, $\mathcal{P}_n = \mathcal{O}(1/\sqrt{n})$ and we recover Theorem 3.1 when $x_i = e_i \in \mathbb{R}^K$ and $\Sigma_0 = \text{diag}(\sigma_{0,i}^2)_{i \in [K]}$. Full expressions and proofs are given in Appendix C.3.

Similarly to our results in multi-armed bandits, $\mathcal{P}_n = \mathcal{O}(1/\sqrt{n})$ since $\lambda_1(\hat{\Sigma}_n) = \mathcal{O}(1/n)$. Also, this bound captures the benefit of using informative priors. Indeed, $\mathcal{P}_n \rightarrow 0$ when the prior variances are small, *i.e.* $\Sigma_0 \rightarrow 0_{L \times L}$, or when the prior gaps are large, $|\mu_0^\top x_i - \mu_0^\top x_j| \rightarrow \infty$.

Theorem 3.3 (Upper bound for hierarchical bandit). *For all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the hierarchical bandit problem (6) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} e^{-\frac{(\nu^\top b_i - \nu^\top b_j)^2}{2(\|b_i - b_j\|_\Sigma^2 + \sigma_{0,i}^2 + \sigma_{0,j}^2)}} \frac{1}{\sqrt{1 + \frac{c_{i,j}}{\hat{\sigma}_{n,i} + \hat{\sigma}_{n,j}}}},$$

where $c_{ij} = \Omega(1)$ and $\hat{\sigma}_{n,i}^2 = \Omega(1/n)$ and they all depend on both the prior parameters and allocation weights. In particular, $\mathcal{P}_n = \mathcal{O}(1/\sqrt{n})$. Full expressions and proofs are given in Appendix C.4.

The term $\|b_i - b_j\|_\Sigma^2 + \sigma_{0,i}^2 + \sigma_{0,j}^2$ accounts for the prior uncertainty of both arms and effects. If the effects are deterministic ($\Sigma \rightarrow 0_{L \times L}$) then our bound recovers the upper bound of MAB with prior mean $\mu_{0,i} = \nu^\top b_i$. On the other hand, if the arms are deterministic given the effects ($\sigma_{0,i}^2 \rightarrow 0$), the bound only depends on the effect covariance. Finally, if the prior is informative by its gap ($|\nu^\top b_i - \nu^\top b_j| \rightarrow \infty$) or by its variance ($\Sigma \rightarrow 0_{L \times L}$ and $\sigma_{0,i}^2 \rightarrow 0$), then $\mathcal{P}_n \rightarrow 0$.

3.3 Allocation Strategies

Instantiating our algorithm requires choosing the potentially prior-dependent allocation weights. Though our bound holds for any such choice, different principles can be used to find empirically satisfying allocations.

Allocation by optimization. Our upper bounds on the PoE in Theorems 3.1 to 3.3 are of the form

$$\mathcal{P}_n \leq C_{\text{prior}}(\omega, n),$$

where $C_{\text{prior}}(\omega, n) = \mathcal{O}(1/\sqrt{n})$ depends on the prior and allocation weights ω . Since it is valid for any $\omega \in \Delta_K^+$, we define the *optimized allocation weights* as

$$\omega^{\text{opt}} = \underset{w \in \Delta_K^+}{\text{argmin}} C_{\text{prior}}(w, n). \quad (11)$$

We denote this variant PI-BAI(ω^{opt}). To fix ideas and give intuition, we give the explicit solution for MAB with $K = 2$. By Theorem 3.1,

$$\mathcal{P}_n \leq \frac{e^{-\frac{(\mu_{0,1} - \mu_{0,2})^2}{2(\sigma_{0,1}^2 + \sigma_{0,2}^2)}} + e^{-\frac{(\mu_{0,1} - \mu_{0,2})^2}{2(\sigma_{0,1}^2 + \sigma_{0,2}^2)}}}{\sqrt{1 + n\phi_{1,2}(\omega_1, \omega_2)}},$$

which can be optimized to obtain

$$\omega_1^{\text{opt}} = \Pi_{[0,1]} \left(\frac{1}{2} - \frac{(\sigma_{0,2}^2 - \sigma_{0,1}^2)\sigma^2}{2n\sigma_{0,1}^2\sigma_{0,2}^2} \right), \quad \omega_2^{\text{opt}} = 1 - \omega_1^{\text{opt}}, \quad (12)$$

where $\Pi_{[a,b]}(\cdot)$ is the projection on $[a, b]$. This expression gives much insight into the allocation strategy. First, in the case of equal prior confidence $\sigma_{0,1}^2 = \sigma_{0,2}^2$, allocating the same amount of samples for each arm is optimal. On the other hand, if for example $\sigma_{0,1}^2 \ll \sigma_{0,2}^2$ for small budget n , we would have $\omega_2^{\text{opt}} \gg \omega_1^{\text{opt}}$, and hence most of the budget would be allocated to the arm with high prior variance (low initial confidence). This discussion is only valid for small budgets: as $n \rightarrow +\infty$, the optimal choice is to divide the budget equally between both arms, since the prior relevance vanishes asymptotically. On the other hand, a saturation phenomenon happens for ‘too small’ budgets: if $n < 2\sigma^2 \left| \frac{\sigma_{0,2}^2}{\sigma_{0,1}^2} - \frac{\sigma_{0,1}^2}{\sigma_{0,2}^2} \right|$, the weight of the arm with larger prior variance is equal to 1 due to the projection. Note that (12) does not depend on the prior gap Δ_0 , which is coherent since identifying the optimal arm is strictly equivalent to identifying the worst arm when dealing with 2 arms. However, this is not necessarily the case when K is larger, as discussed in Appendix D.8.

Since the objective function in (11) is non-convex for $K > 2$, we use numerical optimization to solve it (e.g. L-BFGS-B (Virtanen et al., 2020)). Thankfully, this optimization is done just once before interacting

with the environment. These optimized weights remain non-guaranteed to be good, because (11) is only optimal with respect to the bound we derived. We found it useful in practice to mix them with the heuristic weight ω^h defined in Section 3.1. This allocation is motivated by having a small probability of error when plugged in the bound (Figure 1), and our theoretical guarantees (Theorem 3.1) still hold because the weights are a function of the prior. Hence, we define the new optimized weight as $\alpha\omega_i^{\text{opt}} + (1-\alpha)\frac{\mu_{0,i}\sigma_{0,i}}{\sum_{k \in [K]} \mu_{0,k}\sigma_{0,k}}$. For simplicity, we also refer this as ω^{opt} . We tested the value of the parameter α in various settings, and found that it is generally around 0.5 (Appendix D.7).

Allocation by optimal design. In the linear bandit setting, we generalize ideas from optimal experimental design (Lattimore and Szepesvári, 2020, Chapter 21) to *Bayesian* MAB, linear and hierarchical bandits. To the best of our knowledge, our work is the first to exploit this idea for such settings. Finding an (approximate) Bayesian G-optimal design (see *e.g.* López-Fidalgo (2023, Chapter 4) for a review) is equivalent to maximizing the log-determinant of the regularized information matrix defined as

$$\omega^{\text{G-opt}} = \underset{\omega \in \Delta_K^+}{\operatorname{argmax}} \log \det \left(n\sigma^{-2} \sum_{k \in [K]} \omega_k x_k x_k^\top + \Sigma_0^{-1} \right).$$

This leads to budget allocations that minimize the worst-case posterior variance in all directions. We quantify the effects of using such an allocation in the upper bound:

Corollary 3.4 (Upper bound of PI-BAI($\omega^{\text{G-opt}}$)). *Assume that $x_i \neq x_j$ for any $i \neq j$, and that there exists $S > 0$ such that $\|x\|_2^2 \leq S$ for any $x \in \mathcal{X}$. The expected PoE of PI-BAI($\omega^{\text{G-opt}}$) in the linear bandit problem (4) with diagonal covariance matrix Σ_0 is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{n}{2d\sigma^2} c_{i,j}}},$$

where $c_{i,j} = \Omega(1)$. The full expressions and proofs are given in Appendix C.3. The result also holds for the MAB setting when $x_i = e_i \in \mathbb{R}^K$ and $\Sigma_0 = \operatorname{diag}(\sigma_{0,i}^2)_{i \in [K]}$, and for the hierarchical setting with the equivalent model (10) when Σ is diagonal.

G-optimal design can be applied for MAB, by using $x_i = e_i$ for any $i \in [K]$ and $\Sigma_0 = \operatorname{diag}(\sigma_{0,i}^2)_{i \in [K]}$, and also for hierarchical bandits by using its connection to the linear model (10).

Note that both allocation weights ω^{opt} and $\omega^{\text{G-opt}}$ are prior-dependent but not instance-dependent. In particular, they enjoy the theoretical guarantees we derived in Section 4.

Allocation by warm-up. Finally, we present an adaptive allocation rule, for which our theoretical results do not apply directly, but performs well in practice. Here we use a *warm-up policy* π_w to interact with the bandit environment for n_w rounds (the warm-up phase), and then choose some allocation weights ω_{π_w} based on the prior and the data collected through the interaction. The warm-up policy can be any decision-making policy that (preferably) takes as input the prior distribution. Motivated by its good performance in BAI (Lee et al., 2023), we choose π_w to be Thompson sampling, and select the allocation weights to be proportional to the number of pulls to each arm during the warm-up phase:

$$\omega_i^{\text{TS}} = \frac{\sum_{t \in [n_w]} \mathbb{1}\{A_t = i\}}{n_w}, \quad \forall i \in [K]. \quad (13)$$

To illustrate the differences between optimized weights (PI-BAI(ω^{opt})) and learned weights with Thompson sampling as a warm-up policy (PI-BAI(ω^{TS})), we return to our motivating example in Section 3.1, where $K = 3$, $\mu_0 = (1, 1.9, 2)$ and $\sigma_{0,i} = 0.3$ for all $i \in [3]$. We set the budget as $n = 100$. We repeat 10^4 times the following experiment: we sample a bandit instance from the prior and run Thompson sampling for $n_w = 20$ rounds, then construct the allocation weights ω^{TS} . Computing the weights ω^{opt} by numerical optimization of (11) is done once at the beginning of these experiments.




Figure 2: Allocations weights ω_i^{opt} and ω_i^{TS} .

Figure 2 shows an empirical comparison of the weights on 2 problem instances and on average over 10^4 runs. We see that, in this example, both allocation strategies assign high weights to arms 2 and 3, while allocating a small weight to arm 1. This is because, based on the prior information, arm 1 is highly unlikely to be the optimal arm. Then, the primary objective revolves around selecting the optimal arm among arms 2 and 3. Also, while the allocation weights ω_i^{TS} vary with each bandit instance, their average values in all instances are similar to those of ω_i^{opt} . Thus $\text{PI-BAI}(\omega^{\text{TS}})$ is more adaptive than $\text{PI-BAI}(\omega^{\text{opt}})$, while both have similar average behavior.

4 General Proof Scheme

We outline the key technical insights to derive our Bayesian proofs. The idea is general and can be applied to all our settings. Specific proofs for these settings are in Appendix C.

From Frequentist to Bayesian proof. To analyze their algorithm in the MAB setting, [Atsidakou et al. \(2022\)](#) rely on the strong restriction that $\sigma_{0,i} = \sigma_0$ for all arms $i \in [K]$ and tune their allocations as a function of the noise variance σ^2 such that in the Gaussian setting, the posterior variances $\hat{\sigma}_{n,i}^2$ in (3) are equal for all arms $i \in [K]$. This assumption is needed to directly leverage results from [Karnin et al. \(2013\)](#), thus allowing them to bound the frequentist PoE $\mathbb{P}(J_n \neq i_*(\theta) \mid \theta) \leq B(\theta)$ for a fixed instance θ . Then, the expected PoE, $\mathcal{P}_n = \mathbb{E}[\mathbb{P}(J_n \neq i_*(\theta) \mid \theta)]$, is bounded by computing $\mathbb{E}_{\theta \sim P_0}[B(\theta)]$. We believe it is not possible to extend such technique for general choices of allocations n_i and prior variances $\sigma_{0,i}^2$. Thus, we pursue an alternative approach, establishing the result in a fully Bayesian fashion. We start with a key observation.

Key reformulation of the expected PoE. We observe that the expected PoE can be reformulated as follows

$$\mathcal{P}_n = \mathbb{E}[\mathbb{P}(J_n \neq i_*(\theta) \mid \theta)] = \mathbb{E}[\mathbb{P}(J_n \neq i_*(\theta) \mid H_n)].$$

This swap of measures means that to bound \mathcal{P}_n , we no longer bound the probability of

$$J_n = \underset{i \in [K]}{\operatorname{argmax}} \mathbb{E}[r(i; \theta) \mid H_n] \neq \underset{i \in [K]}{\operatorname{argmax}} r(i; \theta) = i_*(\theta)$$

for any fixed θ . Instead, we only need to bound that probability when θ is drawn from the posterior distribution. Precisely, we bound the probability that the arm i maximizing the posterior mean $\mathbb{E}[r(i; \theta) \mid H_n]$ differs from the arm i maximizing the posterior sample $r(i; \theta) \mid H_n$. This is achieved by first noticing that \mathcal{P}_n can be rewritten as

$$\mathcal{P}_n = \sum_{\substack{i, j \in [K] \\ i \neq j}} \mathbb{E}[\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbf{1}\{J_n = j\}].$$

The rest of the proof consists of bounding the above conditional probabilities for distincts i and j , and this depends on the setting (MAB, linear or hierarchical). Roughly speaking, this is achieved as follows. $\mathbb{P}(i_*(\theta) = i \mid H_n, J_n = j)$ is the probability that arm i maximizes the posterior sample $r(\cdot; \theta)$, given that

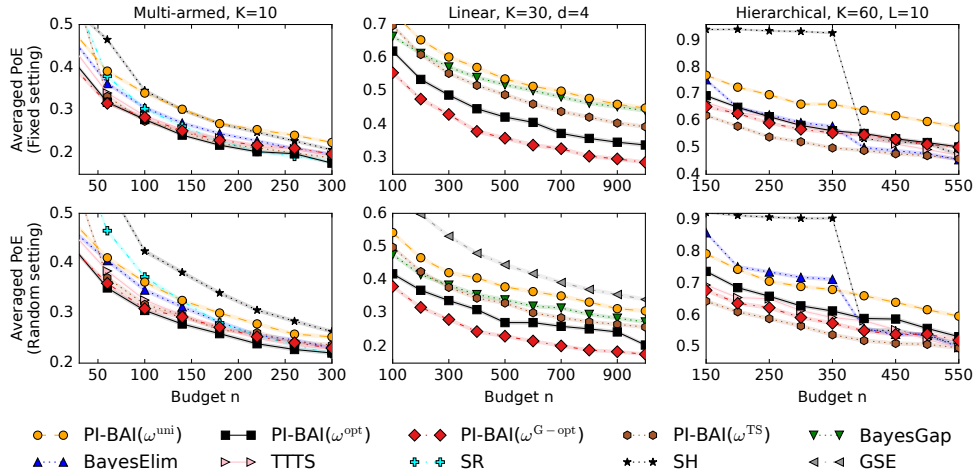


Figure 3: Averaged PoE with varying budgets for the Fixed and Random settings. We compare PI-BAI instantiated with different allocations strategies to the baselines BayesElim, TTS, SR, SH in multi-armed and hierarchical bandits, and to BayesGap and GSE in linear bandits.

arm j maximizes the posterior mean $\mathbb{E}[r(\cdot; \theta) | H_n]$. We show that this probability decays exponentially with the squared difference $(\mathbb{E}[r(i; \theta) | H_n] - \mathbb{E}[r(j; \theta) | H_n])^2$. Taking the expectation of this term under the history H_n gives the desired $\mathcal{O}(1/\sqrt{n})$ rate. All technical details and full explanations can be found in Appendix C. This proof introduces a novel perspective for Bayesian BAI, distinguished by its tighter prior-dependent bounds on the expected PoE and ease of extension to more complex settings like linear and hierarchical bandits. However, its application to adaptive algorithms could be challenging, particularly due to the complexity of taking expectations under the history H_n in that case. Also, extending this proof to non-Gaussian distributions is an interesting direction for future work. Note that PI-BAI is applicable beyond the Gaussian case, as done in Appendices A.2 and D.2 featuring an approximate approach to logistic bandits (Chapelle and Li, 2012).

5 Experiments

We conduct several experiments to evaluate the performance of PI-BAI. In all experiments, we set the observation noise to $\sigma = 1$ and run algorithms 10^4 times and display the (narrow) standard error. The code is available in the supplementary material, and additional experiments and details are presented in Appendix D.1. We consider four variants of PI-BAI(ω) varying in allocation weights: PI-BAI(ω^{uni}) (uniform weights), PI-BAI(ω^{opt}) (optimized weights with mixing), PI-BAI($\omega^{\text{G-opt}}$) (G-optimal design) and PI-BAI(ω^{TS}) (warmed-up with Thompson Sampling). The question of tuning the warm-up length n_w is discussed in Appendix D.7 and we set $n_w = K$.

Multi-armed bandit. We consider two settings, *Fixed* and *Random*. For both, we set $K = 10$ and $\sigma_{0,i}$ evenly spaced between 0.1 and 0.5. In the Fixed setting, $\mu_{0,i}$ is evenly spaced between 0 and 1 whereas in the Random one, $\mu_{0,i} \sim \mathcal{U}([0, 1])$. We compare PI-BAI variants to state-of-the-art Bayesian methods, top-two Thompson sampling (TTS) (Russo, 2016; Jourdan et al., 2022) and BayesElim (Atsidakou et al., 2022), as well as to frequentist elimination algorithms: successive rejects (SR) (Audibert et al., 2010) and sequential halving (SH) (Karnin et al., 2013). Note that TTS does not come with theoretical guarantees in the fixed-budget setting, but we include it given its good empirical performance.

Linear bandit. We let $d = 4$ and $K = 30$. Then we construct the arm set \mathcal{X} by sampling arms x_i from a multivariate Gaussian distribution with mean 0_d and covariance I_d . In the Fixed setting, the prior mean is flat $\mu_0 = (1, \dots, 1)$ whereas in the Random setting, the prior means are sampled uniformly from $[0, 1]$ as

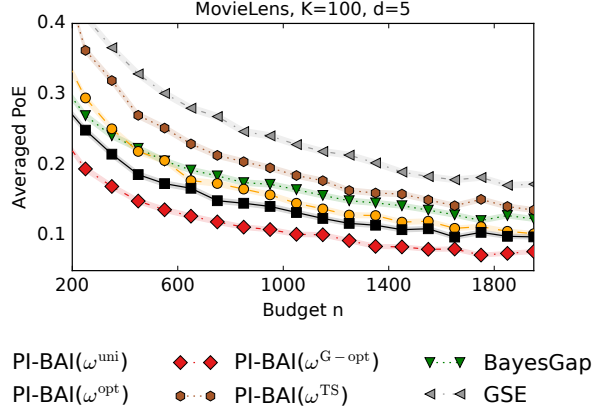


Figure 4: Averaged PoE on MovieLens.

$\mu_{0,i} \sim \mathcal{U}([0, 1])$. For both settings, $\Sigma_0 = \text{diag}(\sigma_{0,i}^2)_{i \in [K]}$ where the $\sigma_{0,i}$'s are evenly spaced between 0.1 and 0.5. We compare our methods with two algorithms that were designed for linear bandits; **BayesGap** (Hoffman et al., 2014) and **GSE** (Azizi et al., 2021), the current (tractable⁴) state-of-the-art that leverages G-optimal design to perform successive elimination. Other methods that leverage the same elimination idea and have lower performances on these settings (Alieva et al., 2021; Yang and Tan, 2022) are not tested.

Hierarchical bandit. Here, each mixing weight b_i is chosen uniformly between 0 and 1, and then normalized to form a probability vector. In the Fixed setting, The ν_i 's are evenly spaced in $[-1, 1]$, whereas each $\nu_i \sim \mathcal{U}([-1, 1])$ in the Random setting. In both settings, Σ is diagonal with entries evenly spaced in $[0.1^2, 0.5^2]$, and Σ_0 is also diagonal, where the $\sigma_{0,i}$'s are evenly spaced between 0.1 and 0.5. The prior distribution for all Bayesian algorithms (except **PI-BAI**) is simply obtained by marginalizing out the effects. This allows obtaining $\mu_{0,i}$ and $\sigma_{0,i}$, even for algorithms that are not suitable for hierarchical priors. Though there is no explicit baseline for this setting, we implement **TS** based on **meTS** (Aouali et al., 2023) and we compare with frequentist and Bayesian elimination strategies agnostic to the structure. Despite the connection to linear bandits, we do not include such baselines as they do not perform well due to their inefficient representation of the structure.

Results on simulated data (Figure 3). Overall, despite setting-dependent variations, **PI-BAI($\omega^{\text{G-opt}}$)** is the best-performing version, closely followed by ω^{opt} . In the hierarchical experiments **PI-BAI(ω^{TS})** surpasses all baselines, highlighting the effectiveness of this method in capturing the underlying problem structure. These observations reaffirm that leveraging prior information is a powerful and practical tool to scale the applicability of BAI to cases with a large number of arms in limited data regimes.

MovieLens data experiment (Figure 4). The MovieLens (Lam and Herlocker, 2016) dataset is a large sparse matrix of ratings from 6040 users (rows) on movies (columns, we subsampled $K = 100$). We first perform a low-rank matrix factorization to obtain $d = 5$ dimensional vectors representing users (context) and movies (actions) as well as estimated corresponding values. We then simulate an online interaction setting: at each round $t \in [n]$, a user vector is picked uniformly at random and a movie is chosen by the policy, leading to a reward $Y_t \sim \mathcal{N}(x_t^\top \theta_t, \sigma^2)$. This semi-synthetic problem allows us to assess **PI-BAI**'s robustness to prior misspecification since the bandit instances are no longer sampled from the prior. More details in Appendix D.1.

⁴Katz-Samuels et al. (2020) has an algorithm with tighter bounds but it is not tractable.

6 Conclusion

We revisit the Bayesian fixed-budget BAI for PoE minimization (Atsidakou et al., 2022) and propose a simple yet efficient algorithm for MAB, linear and hierarchical bandits. Our new proof technique reveals a flexible and smaller upper bound on the expected PoE. In particular, this allows us to derive the first prior-dependent Bayesian PoE upper bounds for linear and hierarchical bandits. We believe that our work sheds a new light on the adaptativity-vs-generality trade-off in BAI algorithm design while opening several avenues for future research. Our work relies on the assumption that the algorithm has access to the true instance-generating distribution, *i.e.* the prior. Though this assumption is very common in the Bayesian bandits literature, it is unrealistic in many scenarios. An interesting question for future work is to explore methods to learn the generative distribution and the related consequences of prior misspecification on the expected PoE of PI-BAI (Kveton et al., 2021; Simchowit et al., 2021; Nguyen and Vernade, 2023).

Broader Impact Statement

In this work we have developed and analyzed generic algorithms for certain optimization problems. Employing our methods may lead to savings in computation and energy. Since our problem setting and our algorithms are generic, the broader (social) impact is unforeseeable and depends on the area where the methods are applied.

References

- Yasin Abbasi-Yadkori, David Pal, and Csaba Szepesvari. Improved algorithms for linear stochastic bandits. In *Advances in Neural Information Processing Systems 24*, pages 2312–2320, 2011.
- Ayya Alieva, Ashok Cutkosky, and Abhimanyu Das. Robust pure exploration in linear bandits with limited budget. In *International Conference on Machine Learning*, pages 187–195. PMLR, 2021.
- Imad Aouali. Linear diffusion models meet contextual bandits with large action spaces. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*, 2023.
- Imad Aouali, Branislav Kveton, and Sumeet Katariya. Mixed-effect thompson sampling. In *International Conference on Artificial Intelligence and Statistics*, pages 2087–2115. PMLR, 2023.
- Alexia Atsidakou, Sumeet Katariya, Sujay Sanghavi, and Branislav Kveton. Bayesian fixed-budget best-arm identification. *arXiv preprint arXiv:2211.08572*, 2022.
- Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In *COLT*, pages 41–53, 2010.
- Javad Azizi, Branislav Kveton, Mohammad Ghavamzadeh, and Sumeet Katariya. Meta-learning for simple regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6709–6717, 2023.
- Mohammad Javad Azizi, Branislav Kveton, and Mohammad Ghavamzadeh. Fixed-budget best-arm identification in structured bandits. *arXiv preprint arXiv:2106.04763*, 2021.
- Soumya Basu, Branislav Kveton, Manzil Zaheer, and Csaba Szepesvari. No regrets for learning the prior in bandits. In *Advances in Neural Information Processing Systems 34*, 2021.
- Christopher M Bishop. *Pattern Recognition and Machine Learning*, volume 4 of *Information science and statistics*. Springer, 2006.
- Sébastien Bubeck, Rémi Munos, and Gilles Stoltz. Pure exploration in multi-armed bandits problems. In *Algorithmic Learning Theory: 20th International Conference, ALT 2009, Porto, Portugal, October 3-5, 2009. Proceedings 20*, pages 23–37. Springer, 2009.

- Alexandra Carpentier and Andrea Locatelli. Tight (lower) bounds for the fixed budget best arm identification bandit problem. In *Conference on Learning Theory*, pages 590–604. PMLR, 2016.
- Olivier Chapelle and Lihong Li. An empirical evaluation of Thompson sampling. In *Advances in Neural Information Processing Systems 24*, pages 2249–2257, 2012.
- Rémy Degenne. On the existence of a complexity in fixed budget bandit identification. *arXiv preprint arXiv:2303.09468*, 2023.
- Eyal Even-Dar, Shie Mannor, Yishay Mansour, and Sridhar Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(6), 2006.
- Sarah Filippi, Olivier Cappé, Aurelien Garivier, and Csaba Szepesvari. Parametric bandits: The generalized linear case. In *Advances in Neural Information Processing Systems 23*, pages 586–594, 2010.
- Matthew Hoffman, Bobak Shahriari, and Nando Freitas. On correlation and budget constraints in model-based bandit optimization with application to automatic machine learning. In *Artificial Intelligence and Statistics*, pages 365–374. PMLR, 2014.
- Joey Hong, Branislav Kveton, Manzil Zaheer, Yinlam Chow, Amr Ahmed, and Craig Boutilier. Latent bandits revisited. In *Advances in Neural Information Processing Systems 33*, 2020.
- Joey Hong, Branislav Kveton, Sumeet Katariya, Manzil Zaheer, and Mohammad Ghavamzadeh. Deep hierarchy in bandits. *arXiv preprint arXiv:2202.01454*, 2022a.
- Joey Hong, Branislav Kveton, Manzil Zaheer, and Mohammad Ghavamzadeh. Hierarchical Bayesian bandits. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, 2022b.
- Joey Hong, Branislav Kveton, Manzil Zaheer, Sumeet Katariya, and Mohammad Ghavamzadeh. Multi-task off-policy learning from bandit feedback. In *International Conference on Machine Learning*, pages 13157–13173. PMLR, 2023.
- Marc Jourdan, Rémy Degenne, Dorian Baudry, Rianne de Heide, and Emilie Kaufmann. Top two algorithms revisited. *Advances in Neural Information Processing Systems*, 35:26791–26803, 2022.
- Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *International conference on machine learning*, pages 1238–1246. PMLR, 2013.
- Julian Katz-Samuels, Lalit Jain, Kevin G Jamieson, et al. An empirical process approach to the union bound: Practical algorithms for combinatorial and linear bandits. *Advances in Neural Information Processing Systems*, 33:10371–10382, 2020.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17:1–42, 2016.
- Junpei Komiyama, Kaito Ariu, Masahiro Kato, and Chao Qin. Rate-optimal bayesian simple regret in best arm identification. *Mathematics of Operations Research*, 2023.
- Branislav Kveton, Manzil Zaheer, Csaba Szepesvari, Lihong Li, Mohammad Ghavamzadeh, and Craig Boutilier. Randomized exploration in generalized linear bandits. In *International Conference on Artificial Intelligence and Statistics*, pages 2066–2076. PMLR, 2020.
- Branislav Kveton, Mikhail Konobeev, Manzil Zaheer, Chih-Wei Hsu, Martin Mladenov, Craig Boutilier, and Csaba Szepesvari. Meta-Thompson sampling. In *Proceedings of the 38th International Conference on Machine Learning*, 2021.
- Shyong Lam and Jon Herlocker. MovieLens Dataset. <http://grouplens.org/datasets/movielens/>, 2016.
- Tor Lattimore and Csaba Szepesvári. *Bandit algorithms*. Cambridge University Press, 2020.

- Jongyeong Lee, Junya Honda, and Masashi Sugiyama. Thompson exploration with best challenger rule in best arm identification. *arXiv preprint arXiv:2310.00539*, 2023.
- Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *The journal of machine learning research*, 18(1): 6765–6816, 2017.
- Jesús López-Fidalgo. *Optimal Experimental Design: A Concise Introduction for Researchers*, volume 226. Springer Nature, 2023.
- Nicolas Nguyen and Claire Vernade. Lifelong best-arm identification with misspecified priors. In *Sixteenth European Workshop on Reinforcement Learning*, 2023.
- Amit Peleg, Naama Pearl, and Ron Meir. Metalearning linear bandits by prior update. In *International Conference on Artificial Intelligence and Statistics*, pages 2885–2926. PMLR, 2022.
- My Phan, Yasin Abbasi Yadkori, and Justin Domke. Thompson sampling and approximate inference. *Advances in Neural Information Processing Systems*, 32, 2019.
- Chao Qin. Open problem: Optimal best arm identification with fixed-budget. In *Conference on Learning Theory*, pages 5650–5654. PMLR, 2022.
- Daniel Russo. Simple bayesian algorithms for best arm identification. In *Conference on Learning Theory*, pages 1417–1418. PMLR, 2016.
- Max Simchowitz, Christopher Tosh, Akshay Krishnamurthy, Daniel Hsu, Thodoris Lykouris, Miro Dudik, and Robert Schapire. Bayesian decision-making under misspecified priors with applications to meta-learning. In *Advances in Neural Information Processing Systems 34*, 2021.
- Inigo Urteaga and Chris Wiggins. Variational inference for the multi-armed contextual bandit. In *Proceedings of the 21st International Conference on Artificial Intelligence and Statistics*, pages 698–706, 2018.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- Po-An Wang, Kaito Ariu, and Alexandre Proutiere. On uniformly optimal algorithms for best arm identification in two-armed bandits with fixed budget. *arXiv preprint arXiv:2308.12000*, 2023.
- R. Wolke and H. Schwetlick. Iteratively reweighted least squares: Algorithms, convergence analysis, and numerical comparisons. *SIAM Journal on Scientific and Statistical Computing*, 9(5):907–921, 1988.
- Junwen Yang and Vincent Tan. Minimax optimal fixed-budget best arm identification in linear bandits. *Advances in Neural Information Processing Systems*, 35:12253–12266, 2022.

Organization of the Appendix

The supplementary material is organized as follows. In Appendix A, we provide additional general additional remarks. In Appendix B, we mention additional existing works relevant to our work. In Appendix C, we give complete statements and proofs of our theoretical results. In Appendix D, we provide additional numerical experiments.

A Additional Discussions

A.1 Motivating examples for hierarchical bandits

In this section, we discuss motivating examples for using hierarchical models in pure exploration settings.

Hyper-parameter tuning. Here the goal is to find the best configuration for a neural network using n epochs (Li et al., 2017). A configuration i is represented by a scalar $\theta_i \in \mathbb{R}$ which quantifies the expected performance of a neural network with such configuration. Again, it is intuitive to learn all θ_i individually. Roughly speaking, this means running each configuration for $\lfloor \frac{n}{K} \rfloor$ epochs and selecting the one with the highest performance. This is statistically inefficient since the number of configurations can be combinatorially large. Fortunately, we can leverage the fact that configurations often share the values of many hyper-parameters. Precisely, a configuration is a combination of multiple hyper-parameters, each set to a specific value. Then we represent each hyper-parameter $\ell \in [L]$ by a scalar parameter $\mu_\ell \in \mathbb{R}$, and the configuration parameter θ_i is a mixture of its hyper-parameters μ_ℓ weighted by their values. That is, $\theta_i = \sum_{\ell \in [L]} b_{i,\ell} \mu_\ell + \epsilon_i$, where $b_{i,\ell}$ is the value of hyper-parameter ℓ in configuration i and ϵ_i is a random noise to incorporate uncertainty due to model misspecification.

Drug design. In clinical trials, K drugs are administrated to n subjects, with the goal of finding the optimal drug design. Each drug is parameterized by its expected efficiency $\theta_i \in \mathbb{R}$. As in the previous example, it is intuitive to learn each θ_i individually by assigning a drug to $\lceil \frac{n}{K} \rceil$ subjects. However, this is inefficient when K is large. We leverage the idea that drugs often share the same components: each drug parameter θ_i is a combination of component parameters μ_l , each accounting for a specific dosage. More precisely, the parameter of drug i can be modeled as $\theta_i = \sum_{l \in [L]} b_{i,l} \mu_l + \epsilon_i$ where ϵ_i accounts for uncertainty due to model misspecification. Similarly to the hyper-parameter tuning example, this models correlations between drugs and it can be leveraged for more efficient use of the whole budget of n epochs.

A.2 Beyond Gaussian distributions

The standard linear model (4) can be generalized beyond linear mean rewards. The *Generalized Linear Bandit* (GLB) model with prior P_0 writes (Filippi et al., 2010; Kveton et al., 2020)

$$\begin{aligned} \theta &\sim P_0 \\ Y_t | \theta, A_t &\sim P(\cdot; \theta, A_t) \quad \forall t \in [n], \end{aligned} \tag{14}$$

where the reward distribution $P(\cdot; \theta, A_t)$ belongs to some exponential family with mean reward $r(A_t; \theta) = \phi(\theta^\top x_{A_t})$. ϕ is called the *link function*. The *log-likelihood* of such reward distribution can be written

$$\mathcal{L}_n(\theta) = \sum_{t=1}^n \log P(Y_t; \theta, A_t) = \sum_{t=1}^n Y_t \theta^\top x_{A_t} - A(\theta^\top x_{A_t}) + h(Y_t),$$

where A is a *log-partition* function and h another function. Importantly, (14) encompasses the *logistic bandit* model with the particular link function $\phi(z) = \frac{1}{1+e^{-z}}$.

The main challenge of (14) is that closed-form posterior generally does not exist. One method is to approximate the posterior distribution of θ given H_n with Laplace approximation: $\mathbb{P}(\theta | H_n)$ is approximated

with a multivariate Gaussian distribution with mean $\hat{\theta}_{\text{MAP}}$ and covariance $\hat{\Sigma}_{\text{Lap}}$, where $\hat{\theta}_{\text{MAP}}$ is the *maximum a posteriori*,

$$\begin{aligned}\hat{\theta}_{\text{MAP}} &= \underset{\theta}{\operatorname{argmax}} \mathcal{L}_n(\theta) P_0(\theta) \\ \hat{\Sigma}_{\text{Lap}}^{-1} &= \sum_{t=1}^n \dot{\phi}(\hat{\theta}_{\text{MAP}}^\top x_{A_t}) x_{A_t} x_{A_t}^\top,\end{aligned}$$

where ϕ is assumed continuously differentiable and increasing. Note that $\hat{\theta}_{\text{MAP}}$ can be computed efficiently by iteratively reweighted least squares (Wolke and Schwetlick, 1988).

Logistic Bandit. In the particular case where the reward distribution is Bernoulli, the model writes

$$\begin{aligned}\theta &\sim \mathcal{N}(\mu_0, \Sigma_0) \\ Y_t \mid \theta, A_t &\sim \mathcal{B}(\phi(\theta^\top x_{A_t})) \quad \forall t \in [n],\end{aligned}\tag{15}$$

where ϕ is the logistic function. Then the mean posterior reward can be approximated as

$$\mathbb{E}_{\theta \sim \mathcal{N}(\hat{\theta}_{\text{MAP}}, \hat{\Sigma}_{\text{Lap}})} [\phi(\theta^\top x_i)] \approx \frac{\phi(\hat{\theta}_{\text{MAP}}^\top x_i)}{\sqrt{1 + \frac{\pi}{8} \|x_i\|_{\hat{\Sigma}_{\text{Lap}}}}},$$

and the decision after n rounds is $J_n = \operatorname{argmax}_{i \in [K]} \frac{\phi(\hat{\theta}_{\text{MAP}}^\top x_i)}{\sqrt{1 + \frac{\pi}{8} \|x_i\|_{\hat{\Sigma}_{\text{Lap}}}}$.

Proving an upper bound on the expected PoE of this algorithm is challenging. Particularly, upper bounding the expectation with respect to H_n is hard because one needs to show that $\hat{\theta}_{\text{MAP}}$ concentrates in norm towards its expectation $\mathbb{E}_{H_n} [\hat{\theta}_{\text{MAP}}]$. We leave this study for future work. However, we provide numerical experiments for this setting in Appendix D.2.

A.3 Additional Remarks on Hierarchical Models

The two-level prior we consider has a shared latent parameter $\mu = (\mu_\ell)_{\ell \in [L]} \in \mathbb{R}^L$ representing L effects impacting each of the K arm means:

$$\begin{aligned}\mu &\sim Q \\ \theta_i &\sim P_{0,i}(\cdot; \mu) \quad \forall i \in [K] \\ Y_t \mid \mu, \theta, A_t &\sim P(\cdot; \theta_{A_t}) \quad \forall t \in [n],\end{aligned}$$

where Q is the latent prior on $\mu \in \mathbb{R}^L$.

In the Gaussian setting (6), the maximum likelihood estimate of the reward mean of action i , $B_{n,i}/n_i$, contributes to (7) proportionally to its precision $n_i/(n_i\sigma_{0,i}^2 + \sigma^2)$ and is weighted by its mixing weight b_i . (8) is a standard Gaussian posterior, and (9) takes into account the information of the conditional posterior. Finally, (9) takes into account the arm correlation through its dependence on $\hat{\sigma}_{n,i}$ and $\hat{\mu}_{n,i}$. While the properties of conjugate priors are useful for inference, other models could be considered with approximate inference techniques (Urteaga and Wiggins, 2018; Phan et al., 2019).

Link with linear bandit (cont.). The slightly unusual characteristic of (10) is that the prior distribution has correlated components. This can be addressed by the whitening trick (Bishop, 2006), defining $\tilde{\mu} = \Sigma^{-1/2}\bar{\mu}$ and $\tilde{b}_i = \Sigma^{1/2}\bar{b}_i$, giving

$$\begin{aligned}\tilde{\mu} &\sim \mathcal{N}(\bar{\Sigma}^{-1/2}\bar{\nu}, I_{L+K}) \\ Y_t \mid \tilde{\mu}, A_t &\sim \mathcal{N}(\tilde{b}_{A_t}^\top \tilde{\mu}, \sigma^2) \quad \forall t \in [n],\end{aligned}\tag{16}$$

where I_{L+K} is the $(L+K)$ -dimensional identity matrix. Then, (16) corresponds to a linear bandit model with K arms and $d = K + L$ features. However, this model comes with some limitations. First, when computing

posteriors under (16), the time and space complexities are $\mathcal{O}((K+L)^3)$ and $\mathcal{O}((K+L)^2)$ respectively, compared to the $\mathcal{O}(K+L^3)$ and $\mathcal{O}(K+L^2)$ for our model (9). The feature dimension $d = K+L$ can be reduced to $d = K$ through the following QR decomposition: $\tilde{B} = (\tilde{b}_1, \dots, \tilde{b}_K) \in \mathbb{R}^{(K+L) \times K}$ can be expressed as $\tilde{B} = VR$, where $V \in \mathbb{R}^{(K+L) \times K}$ is an orthogonal matrix and $R \in \mathbb{R}^{K \times K}$. This leads to the following model $\check{\mu} \sim \mathcal{N}(V^\top \bar{\Sigma}^{-1/2} \bar{\nu}, V^\top V)$ and $Y_t | \check{\mu}, A_t \sim \mathcal{N}(R_{A_t}^\top \check{\mu}, \sigma^2)$, yet the feature dimension d remains at the order of K , and computational efficiency is not improved with respect to K .

From hierarchical bandit to MAB. Marginilizing the hyper-prior in (6) leads to a MAB model,

$$\begin{aligned} \theta_i &\sim \mathcal{N}(b_i^\top \nu, \sigma_{0,i}^2 + b_i^\top \Sigma b_i) & \forall i \in [K] & (17) \\ Y_t | \mu, \theta, A_t &\sim \mathcal{N}(\theta_{A_t}, \sigma^2) & \forall t \in [n]. & (18) \end{aligned}$$

In this marginalized model, the agent does not know μ and he doesn't want to model it. Therefore, only θ is learned. The marginalized prior variance $\sigma_{0,i}^2 + b_i^\top \Sigma b_i$ accounts for the uncertainty of the not-modeled effects.

B Extended Related Work

In this section, we provide additional references relevant to our work.

Bayesian bandits in structured environments. Bayesian bandits algorithms under hierarchical models have been heavily studied (Hong et al., 2020; Kveton et al., 2021; Basu et al., 2021; Hong et al., 2022a,b; Peleg et al., 2022; Aouali et al., 2023; Aouali, 2023). All the aforementioned papers propose methods to explore efficiently in the structured environment to minimize the (Bayesian) regret. The hierarchical model we use is derived from Aouali et al. (2023). Beyond regret minimization, Bayesian structured models also found success in simple regret minimization (Azizi et al., 2023) and off-policy learning in bandits (Hong et al., 2023).

Bayesian simple regret minimization. Azizi et al. (2023) considers the problem of simple regret minimization in a Bayesian hierarchical setting. Their algorithm is based on Thompson sampling, and choose an arm at the last round by sampling according to the number of pulls. This leads to a $\mathcal{O}(1/\sqrt{n})$ rate on the Bayesian simple regret. Recently, Komiyama et al. (2023) derived a method for Bayesian simple regret minimization that asymptotically matches their proposed lower bound scaling in $\mathcal{O}(1/n)$. Their result does not contradict our analysis because as mentioned in their work, the difference between the simple regret and PoE matters in terms of rate when considering a Bayesian objective, unlike in the frequentist case. Moreover, their method is designed for Bernoulli rewards, and their algorithm does not use the prior distribution.

C Proofs

In this section, we give complete proof of our theoretical results. In Appendix C.1, we give proofs for the Bayesian posterior derivations and we provide technical results. In Appendix C.2, we prove Theorem 3.1. The proofs for linear bandits (Theorem 3.2 and Corollary 3.4) are given in Appendix C.3. In Appendix C.4, we provide proofs and technical remarks for hierarchical bandits (Theorem 3.3).

C.1 Technical Proofs and Posteriors Derivations

Bayesian computations. We focus on the hierarchical Gaussian case (6) and detail the computations of posterior distributions. We first recall the model,

$$\begin{aligned} \mu &\sim \mathcal{N}(\nu, \Sigma) \\ \theta_i &\sim \mathcal{N}(b_i^\top \mu, \sigma_{0,i}^2) & \forall i \in [K] \\ Y_t | \mu, \theta, A_t &\sim \mathcal{N}(\theta_{A_t}, \sigma^2) & \forall t \in [n], \end{aligned}$$

where we recall that $B_{n,i} = \sum_{t \in \mathcal{T}_i} Y_t$ and $\mathcal{T}_i = \{t \in [n], A_t = i\}$.

Lemma C.1 (Gaussian posterior update). *For any $\rho \in \mathbb{R}, \mu \in \mathbb{R}^L, b \in \mathbb{R}^L$ and $\sigma, \sigma_0 > 0, m \in \mathbb{N}$, we have*

$$\int_{\rho} \prod_{t \in [m]} \mathcal{N}(Y_t; \rho, \sigma^2) \mathcal{N}(\rho; b^\top \mu, \sigma_0^2) d\rho \propto \mathcal{N}(\mu; \mu_m, \Sigma_m),$$

with

$$\Sigma_m = \frac{m}{m\sigma_0^2 + \sigma^2} bb^\top; \quad \mu_m = \Sigma_m^{-1} \frac{\sum_{t \in [m]} Y_t}{m\sigma_0^2 + \sigma^2} b.$$

Proof of Lemma C.1. By keeping only terms that depend on μ ,

$$\begin{aligned} f(\mu) &= \int_{\rho} \prod_{t \in [m]} \mathcal{N}(Y_t; \rho, \sigma^2) \mathcal{N}(\rho; b^\top \mu, \sigma_0^2) d\rho \\ &\propto \int_{\rho} \exp \left\{ -\frac{1}{2\sigma^2} \sum_{t \in [m]} (Y_t - \rho)^2 - \frac{1}{2\sigma_0^2} (\rho - b^\top \mu)^2 \right\} d\rho \\ &\propto \int_{\rho} \exp \left\{ -\frac{1}{2} \rho^2 \left(\frac{1}{\sigma_0^2} + \frac{m}{\sigma^2} \right) - 2\rho \left(\frac{\sum_{t \in [m]} Y_t}{\sigma^2} + \frac{b^\top \mu}{\sigma_0^2} \right) \right\} d\rho \exp \left\{ -\frac{1}{2\sigma_0^2} \mu^\top bb^\top \mu \right\} \\ &\propto \exp \left\{ \frac{1}{2} \left(\frac{1}{\sigma^2} \sum_{t \in [m]} Y_t + \frac{b^\top \mu}{\sigma_0^2} \right)^2 \frac{\sigma_0^2 \sigma^2}{\sigma^2 + m\sigma_0^2} - \frac{1}{2\sigma_0^2} \mu^\top bb^\top \mu \right\} \\ &\propto \exp \left\{ \frac{\sum_{t \in [m]} Y_t b^\top \mu}{\sigma^2 + m\sigma_0^2} + \frac{\sigma^2}{2(\sigma^2 + m\sigma_0^2)} \mu^\top bb^\top \mu - \frac{1}{\sigma_0^2} bb^\top \mu \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\mu^\top \frac{m}{\sigma^2 + m\sigma_0^2} bb^\top \mu - 2\mu^\top \frac{\sum_{t \in [m]} Y_t}{\sigma^2 + m\sigma_0^2} b \right) \right\} \\ &\propto \mathcal{N}(\mu; \mu_m, \Sigma_m). \end{aligned}$$

□

Lemma C.2 (Joint effect posterior). *For any $n \in [N]$, the joint effect posterior is a multivariate Gaussian $Q_n(\mu) = \mathcal{N}(\check{\mu}, \check{\Sigma}_n)$, where*

$$\check{\Sigma}_n^{-1} = \Sigma^{-1} + \sum_{i \in [K]} \frac{n_i}{n_i \sigma_{0,i}^2 + \sigma^2} b_i b_i^\top, \quad \check{\mu}_n = \check{\Sigma}_n \left(\Sigma^{-1} \nu + \sum_{i \in [K]} \frac{B_{n,i}}{n_i \sigma_{0,i}^2 + \sigma^2} b_i \right).$$

Proof of Lemma C.2. The joint effect posterior can be written as

$$\begin{aligned} Q_n(\mu) &\propto \int_{\theta} \mathcal{L}_{\theta}(Y_{A_1}, \dots, Y_{A_n}) P_0(\theta | \mu) d\theta Q(\mu) \\ &= \prod_{i \in [K]} \int_{\theta_i} \prod_{t \in \mathcal{T}_i} \mathcal{N}(Y_t; \theta_i, \sigma^2) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) d\theta_i \mathcal{N}(\mu; \nu, \Sigma). \end{aligned}$$

Applying Lemma C.1 gives

$$\int_{\theta_i} \prod_{t \in \mathcal{T}_i} \mathcal{N}(Y_t; \theta_i, \sigma^2) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) d\theta_i \propto \mathcal{N}(\mu; \bar{\mu}_{n,i}, \bar{\Sigma}_{n,i}),$$

with

$$\bar{\Sigma}_{n,i}^{-1} = \frac{n_i}{\sigma_{0,i}^2 n_i + \sigma^2} b_i b_i^\top, \quad \bar{\mu}_{n,i} = \bar{\Sigma}_{n,i} b_i \frac{n_i}{\sigma_{0,i}^2 n_i + \sigma^2} \frac{B_{n,i}}{n_i}.$$

Therefore, the joint effect posterior is a product of Gaussian distributions,

$$Q_n(\mu) \propto \prod_{i \in [K]} \mathcal{N}(\mu; \check{\mu}_{n,i}, \check{\Sigma}_{n,i}) \mathcal{N}(\mu; \nu, \Sigma) \propto \mathcal{N}(\mu; \check{\mu}_n, \check{\Sigma}_n),$$

where

$$\begin{aligned} \check{\Sigma}_n &= \Sigma^{-1} + \sum_{i \in [K]} \check{\Sigma}_{n,i} = \Sigma^{-1} + \sum_{i \in [K]} \frac{n_i}{\sigma_{0,i}^2 n_i + \sigma^2} b_i b_i^\top \\ \check{\mu}_n &= \check{\Sigma}_n^{-1} \left(\Sigma^{-1} \nu + \sum_{i \in [K]} \check{\Sigma}_{n,i}^{-1} \check{\mu}_{n,i} \right) = \check{\Sigma}_n^{-1} \left(\Sigma^{-1} \nu + \sum_{i \in [K]} \frac{B_{n,i}}{n_i \sigma_{0,i}^2 + \sigma^2} b_i \right). \end{aligned}$$

□

Lemma C.3 (Conditional arm posteriors). *For any $n \in [n]$ and any arm $i \in [K]$, the conditional posterior distribution of arm i is a Gaussian distribution $P_{n,i}(\theta_i | \mu) = \mathcal{N}(\check{\mu}_{n,i}, \check{\sigma}_{n,i}^2)$, where*

$$\check{\sigma}_{n,i}^{-2} = \frac{1}{\sigma_{0,i}^2} + \frac{n_i}{\sigma^2}, \quad \check{\mu}_{n,i} = \check{\sigma}_{n,i}^2 \left(\frac{\mu^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma^2} \right).$$

Proof of Lemma C.3. The conditional posterior of arm i can be written as

$$\begin{aligned} P_{n,i}(\theta_i | \mu) &\propto \mathcal{L}_{\theta_i}(Y_{A_1}, \dots, Y_{A_n}) P_{0,i}(\theta_i | \mu) \\ &\propto \prod_{t \in \mathcal{T}_i} \mathcal{N}(Y_t; \theta_i, \sigma^2) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) \\ &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{t \in \mathcal{T}_i} (Y_t - \theta_i)^2 - \frac{1}{2\sigma_{0,i}^2} (\theta_i - b_i^\top \mu)^2 \right\} \\ &\propto \exp \left\{ -\frac{1}{2\sigma^2} \sum_{t \in \mathcal{T}_i} (-2Y_t \theta_i + \theta_i^2) - \frac{1}{2\sigma_{0,i}^2} (\theta_i^2 - 2\theta_i b_i^\top \mu) \right\} \\ &\propto \exp \left\{ -\frac{1}{2} \left(\theta_i^2 \left(\frac{n_i}{\sigma^2} + \frac{1}{\sigma_{0,i}^2} \right) - 2\theta_i \left(\frac{1}{\sigma^2} \sum_{t \in \mathcal{T}_i} Y_t + \frac{1}{\sigma_{0,i}^2} b_i^\top \mu \right) \right) \right\} \\ &\propto \mathcal{N}(\theta_i; \check{\mu}_{n,i}, \check{\sigma}_{n,i}^2) \end{aligned}$$

□

Lemma C.4 (Marginal arm posterior). *For any $n \in [n]$ and any arm $i \in [K]$, the marginal posterior distribution of arm i is a Gaussian distribution $\mathbb{P}(\theta_i | H_n) = \mathcal{N}(\hat{\mu}_{n,i}, \hat{\sigma}_{n,i}^2)$, where*

$$\hat{\sigma}_{n,i}^2 = \check{\sigma}_{n,i}^2 + \frac{\check{\sigma}_{n,i}^4}{\sigma_{0,i}^4} b_i^\top \check{\Sigma}_n b_i, \quad \hat{\mu}_{n,i} = \check{\sigma}_{n,i}^2 \left(\frac{\check{\mu}_{n,i}^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma_i^2} \right).$$

Proof of Lemma C.4. The marginal distribution of arm i can be written as

$$\begin{aligned} \int_{\mu} P_{n,i}(\theta_i | \mu) Q_n(\mu) d\mu &= \int_{\mu} \mathcal{N}(\theta_i; \check{\mu}_{n,i}, \check{\sigma}_{n,i}^2) \mathcal{N}(\mu; \check{\mu}_n, \check{\Sigma}_n) d\mu \\ &\propto \int_{\mu} \mathcal{N} \left(\theta_i; \check{\sigma}_{n,i}^2 \left(\frac{\mu^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma^2} \right), \check{\sigma}_{n,i}^2 \right) \mathcal{N}(\mu; \check{\mu}_n, \check{\Sigma}_n) d\mu. \end{aligned}$$

The line above is a convolution of Gaussian measures, and can be written as (Bishop, 2006),

$$\begin{aligned} \int_{\mu} P_{n,i}(\theta_i | \mu) Q_n(\mu) d\mu &\propto \mathcal{N} \left(\theta_i; \check{\sigma}_{n,i}^2 \left(\frac{\check{\mu}_n^\top b_i}{\sigma_{0,i}^2} + \frac{B_{n,i}}{\sigma_i^2} \right), \check{\sigma}_{n,i}^2 + \frac{\check{\sigma}_{n,i}^2}{\sigma_{0,i}^2} b_i^\top \check{\Sigma}_n \frac{\check{\sigma}_{n,i}^2}{\sigma_{0,i}^2} b_i \right) d\mu \\ &= \mathcal{N}(\theta_i; \hat{\mu}_{n,i}, \hat{\sigma}_{n,i}^2). \end{aligned}$$

□

Lemma C.5 (Technical lemma). *Let $a > 0$ and $X \sim \mathcal{N}(\mu, \sigma^2)$. Then $\mathbb{E}_X \left[e^{-\frac{X^2}{2a^2}} \right] = \frac{1}{\sqrt{1 + \frac{\sigma^2}{a^2}}} e^{-\frac{\mu^2}{2(a^2 + \sigma^2)}}$.*

C.2 Proofs for MAB

From now, we consider that $\lfloor \omega_k n \rfloor = \omega_k n \in \mathbb{N}$ for sake of simplicity.

Theorem C.6 (Complete statement of Theorem 3.1). *For all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the MAB problem (2) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{2(\sigma_{0,i}^2 + \sigma_{0,j}^2)}}}{\sqrt{1 + n \frac{\sigma_{0,i}^4 \omega_i \left(\frac{\sigma^2}{n} + \omega_j \sigma_{0,j}^2 \right) + \sigma_{0,j}^4 \omega_j \left(\frac{\sigma^2}{n} + \omega_i \sigma_{0,i}^2 \right)}{\sigma^2 \sigma_{0,i}^2 \left(\frac{\sigma^2}{n} + \omega_j \sigma_{0,j}^2 \right) + \sigma^2 \sigma_{0,j}^2 \left(\frac{\sigma^2}{n} + \omega_i \sigma_{0,i}^2 \right)}}} := \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{2(\sigma_{0,i}^2 + \sigma_{0,j}^2)}}}{\sqrt{1 + n \phi_{i,j}}},$$

Remark C.7. When $\sigma_{0,i}^2 = \sigma_{0,j}^2$, $\lim_{n \rightarrow +\infty} \phi_{i,j} = \lim_{n \rightarrow +\infty} \frac{\sigma_{0,i}^4 \omega_i \left(\frac{\sigma^2}{n} + \omega_j \sigma_{0,j}^2 \right) + \sigma_{0,j}^4 \omega_j \left(\frac{\sigma^2}{n} + \omega_i \sigma_{0,i}^2 \right)}{\sigma^2 \sigma_{0,i}^2 \left(\frac{\sigma^2}{n} + \omega_j \sigma_{0,j}^2 \right) + \sigma^2 \sigma_{0,j}^2 \left(\frac{\sigma^2}{n} + \omega_i \sigma_{0,i}^2 \right)} = \frac{2\sigma_{0,i}^2 \omega_i \omega_j}{\sigma^2 \omega_i + \omega_j} = \mathcal{O}(1)$.

Proof of Theorem 3.1. We first write \mathcal{P}_n as a double sum over all possible distinct arms,

$$\begin{aligned} \mathbb{E} [\mathbb{P}(J_n \neq i_*(\theta) \mid H_n)] &= \mathbb{E} [\mathbb{1}\{J_n \neq i_*(\theta)\}] \\ &= \mathbb{E} \left[\mathbb{E} \left[\sum_{i=1}^K \sum_{j=1}^K \mathbb{1}\{i \neq j\} \mathbb{1}\{i_*(\theta) = i\} \mathbb{1}\{J_n = j\} \mid H_n \right] \right] \\ &= \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{E} [\mathbb{1}\{i_*(\theta) = i\} \mathbb{1}\{J_n = j\} \mid H_n]] \\ &= \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \cap J_n = j \mid H_n)] \end{aligned}$$

Since $J_n : H_n \rightarrow [K]$, $\mathbb{P}(J_n = j \mid H_n) = \mathbb{1}\{J_n = j\}$. Considering both events $\{J_n = j\}$ or $\{J_n \neq j\}$ under H_n ,

$$\begin{aligned} \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \cap J_n = j \mid H_n)] &= \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \cap J_n = j \mid H_n) (\mathbb{1}\{J_n = j\} + \mathbb{1}\{J_n \neq j\})] \\ &= \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mid J_n = j] \mathbb{P}(J_n = j) \cdot \\ &= \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbb{1}\{J_n = j\}]. \end{aligned}$$

Overall,

$$\mathcal{P}_n = \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbb{1}\{J_n = j\}].$$

By definition of $i_*(\theta)$ in the MAB setting and applying Hoeffding inequality for sub-Gaussian random variables,

$$\begin{aligned} \mathbb{P}\left(\operatorname{argmax}_{k \in [K]} \theta_k = i \mid H_n, J_n = j\right) &\leq \mathbb{P}(\theta_i \geq \theta_j \mid H_n, J_n = j) \\ &= \mathbb{P}((\theta_i - \theta_j) - (\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) \geq -(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) \mid H_n, J_n = j) \\ &\leq \exp\left(-\frac{(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)}\right). \end{aligned} \quad (19)$$

Therefore,

$$\mathbb{E}[\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbb{1}\{J_n = j\}] \leq \mathbb{E}\left[\exp\left(-\frac{(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)}\right)\right] \quad (20)$$

We now want to compute this above expectation with respect to H_n .

First, we remark that because the scheduling of arms (A_1, \dots, A_n) is deterministic, the law of $H_n = (A_1, Y_{A_1}, \dots, A_n, Y_{A_n})$ is the law of $(Y_{A_1}, \dots, Y_{A_n})$. Denoting π_{H_n} the marginal distribution of H_n ,

$$\pi_{H_n}(H_n) = \pi_{H_n}(Y_{A_1}, \dots, Y_{A_n}) = \int_{\theta} \mathcal{L}_{\theta}(Y_{A_1}, \dots, Y_{A_n}) P_0(\theta) d\theta,$$

where $\mathcal{L}_{\theta}(Y_{A_1}, \dots, Y_{A_n})$ denotes the likelihood of $(Y_{A_1}, \dots, Y_{A_n})$ given parameter θ and $P_0(\theta) = \prod_{i \in [K]} P_{0,i}(\theta_i)$ since each mean reward θ_i is drawn independently from $P_{0,i}$ in the MAB setting. Since rewards given parameter θ are independent and identically distributed,

$$\begin{aligned} \pi_{H_n}(H_n) &= \int_{\theta} \prod_{i \in [K]} \mathcal{L}_{\theta_i}((Y_t)_{t \in \mathcal{T}_i}) P_{0,i}(\theta_i) d\theta_i \\ &= \int_{\theta} \prod_{i \in [K]} \mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; \theta_i \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n}) \mathcal{N}(\theta_i; \mu_{0,i}, \sigma_{0,i}^2) d\theta_i, \end{aligned} \quad (21)$$

where $\mathbf{1}_q$ denotes the vector of size q whose all components are 1s.

(21) is a convolution of Gaussians and can be computed easily (Bishop, 2006),

$$\mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; \theta_i \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n}) \mathcal{N}(\theta_i; \mu_{0,i}, \sigma_{0,i}^2) = \mathcal{N}((Y_t)_{t \in \mathcal{A}_i}; \mu_{0,i} \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n} + \sigma_{0,i}^2 \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^{\top}).$$

The above covariance matrix exhibits $\sigma^2 + \sigma_{0,i}^2$ on the diagonal and $\sigma_{0,i}^2$ out of diagonal.

We are now ready to compute some useful statistics : for any $i \in [K]$,

$$\mathbb{E}[\hat{\mu}_{n,i}] = \mathbb{E}\left[\frac{\sigma^2}{\sigma^2 + \sigma_{0,i}^2 \omega_i n} \mu_{0,i} + \frac{\sigma_{0,i}^2}{\sigma^2 + \sigma_{0,i}^2 \omega_i n} \sum_{t \in \mathcal{T}_i} Y_t\right] = \mu_{0,i} \quad (22)$$

$$\mathbb{V}(\hat{\mu}_{n,i}) = \frac{\sigma_{0,i}^4}{(\sigma^2 + \sigma_{0,i}^2 \omega_i n)^2} \mathbb{V}\left(\sum_{t \in \mathcal{T}_i} Y_t\right) = \frac{\sigma_{0,i}^4}{(\sigma^2 + \sigma_{0,i}^2 \omega_i n)} \omega_i n \quad (23)$$

$$\mathbb{E}\left[\frac{1}{\omega_i n} \sum_{t \in \mathcal{T}_i} Y_t\right] = \mu_{0,i} \quad (24)$$

$$\mathbb{V}\left(\frac{1}{\omega_i n} \sum_{t \in \mathcal{T}_i} Y_t\right) = \frac{1}{\omega_i^2 n^2} (\omega_i n (\sigma^2 + \sigma_{0,i}^2) + (\omega_i^2 n^2 - \omega_i n) \sigma_{0,i}^2) = \frac{\sigma^2}{\omega_i n} + \sigma_{0,i}^2 \quad (25)$$

Applying Lemma C.5 on (20) and simplifying terms gives

$$\mathbb{E}[\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbb{1}\{J_n = j\}] \leq \mathbb{E}\left[\exp\left(-\frac{(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)}\right)\right] = \frac{e^{-\frac{(\mu_{0,i} - \mu_{0,j})^2}{2(\sigma_{0,i}^2 + \sigma_{0,j}^2)}}}{\sqrt{1 + n\phi_{i,j}}}.$$

C.3 Proofs for Linear Bandits

Theorem C.8 (Complete statement of Theorem 3.2). *Assume that $x_i \neq x_j$ for any $i \neq j$, and that there exists $S > 0$ such that $\|x\|_2^2 \leq S$ for any $x \in \mathcal{X}$. Then, for all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the linear bandit problem (4) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{1}{\sqrt{1 + \frac{c_{i,j}}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2}}} e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\hat{\Sigma}_0}^2}},$$

where:

$$c_{i,j} = \|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2, \quad \text{Cov}(\hat{\mu}_n) = \frac{1}{\sigma^4} \hat{\Sigma}_n \left(\underbrace{\sum_{i \in [K]} (\omega_i n (\sigma^2 + x_i^\top \Sigma_0 x_i)) x_i x_i^\top}_{\text{variance terms}} + \underbrace{\sum_{i \in [K]} \sum_{j \in [K] \setminus \{i\}} x_i^\top \Sigma_0 x_j \omega_j \omega_j n^2 x_i x_j^\top}_{\text{covariance between arms}} \right) \hat{\Sigma}_n.$$

Proof of Theorem 3.2. The proof for the linear model follows the same steps as the MAB model by rewriting \mathcal{P}_n as

$$\mathcal{P}_n = \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbf{1}\{J_n = j\}].$$

By definition of $i_*(\theta)$ and J_n in the linear bandit setting,

$$\begin{aligned} \mathbb{P}(i_*(\theta) = i \mid H_n, J_n = j) &= \mathbb{P}(\forall k \in [K], \theta^\top x_i \geq \theta^\top x_k \mid H_n, J_n = j) \\ &\leq \mathbb{P}(\theta^\top x_i \geq \theta^\top x_j \mid J_n = j, H_n) \\ &\leq \exp\left(-\frac{\|\hat{\mu}_n\|_{(x_i - x_j)(x_i - x_j)^\top}^2}{2\|x_i - x_j\|_{\hat{\Sigma}_n}^2}\right), \end{aligned}$$

where the last inequality follows from Hoeffding inequality for sub-Gaussian random variables. Taking the expectation with respect to H_n ,

$$\mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid H_n, J_n = j) \mathbf{1}\{J_n = j\}] \leq \mathbb{E} \left[\exp\left(-\frac{\|\hat{\mu}_n\|_{(x_i - x_j)(x_i - x_j)^\top}^2}{2\|x_i - x_j\|_{\hat{\Sigma}_n}^2}\right) \right]. \quad (26)$$

Then we remark that the expectation of $\hat{\mu}_n$ with respect to H_n is

$$\mathbb{E} [\hat{\mu}_n] = \mathbb{E} \left[\hat{\Sigma}_n \left(\Sigma_0^{-1} \mu_0 + \frac{1}{\sigma^2} \sum_{t \in [n]} Y_t x_{A_t} \right) \right] = \hat{\Sigma}_n \left(\Sigma_0^{-1} \mu_0 + \frac{1}{\sigma^2} \mathbb{E} \left[\sum_{t \in [n]} Y_t x_{A_t} \right] \right),$$

since the scheduling (A_1, \dots, A_n) is known beforehand. Now,

$$\mathbb{E} \left[\sum_{t \in [n]} Y_t x_{A_t} \right] = \sum_{t \in [n]} \mathbb{E} [Y_t] x_{A_t} = \sum_{t \in [n]} \mu_0^\top x_{A_t} x_{A_t},$$

where $\mathbb{E} [Y_t]$ was obtained by marginalizing the likelihood over the prior distribution as in (21).

Rearranging the terms permits to conclude that $\mathbb{E} [\hat{\mu}_n] = \mu_0$. Then we can compute the expectation in (26) by applying Lemma C.5, Sylvester identity, and some simplifications:

$$\mathbb{E} \left[\exp\left(-\frac{\|\hat{\mu}_n\|_{(x_i - x_j)(x_i - x_j)^\top}^2}{2\|x_i - x_j\|_{\hat{\Sigma}_n}^2}\right) \right] = \frac{1}{\sqrt{1 + \frac{\|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2}}} e^{-\frac{1}{2} \|\mu_0\|_{\Lambda_{ij}}^2}. \quad (27)$$

where from Lemma C.5,

$$\begin{aligned}
\Lambda_{i,j} &= \text{Cov}(\hat{\mu}_n)^{-1} - \text{Cov}(\hat{\mu}_n)^{-1} \left(\text{Cov}(\hat{\mu}_n)^{-1} + \frac{(A_i - A_j)(A_i - A_j)}{\|A_i - A_j\|_{\hat{\Sigma}_n}^2} \right)^{-1} \text{Cov}(\hat{\mu}_n)^{-1} \\
&= \text{Cov}(\hat{\mu}_n)^{-1} - \text{Cov}(\hat{\mu}_n)^{-1} \left(I_d + \text{Cov}(\hat{\mu}_n) \frac{(x_i - x_j)(x_i - x_j)}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2} \right)^{-1} \\
&= \frac{(x_i - x_j)(x_i - x_j)^\top}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2 + \|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2}.
\end{aligned}$$

The last equality follows from an application of Sherman-Morrison identity. Applying the law of total expectation,

$$\text{Cov}(\theta) = \mathbb{E}[\text{Cov}(\theta | H_n)] + \text{Cov}(\mathbb{E}[\theta | H_n]) = \hat{\Sigma}_n + \text{Cov}(\hat{\mu}_n).$$

Therefore,

$$\|x_i - x_j\|_{\hat{\Sigma}_n}^2 + \|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2 = \|x_i - x_j\|_{\hat{\Sigma}_n + \text{Cov}(\hat{\mu}_n)}^2 = \|x_i - x_j\|_{\Sigma_0}^2.$$

Plugging these into (27), we obtain

$$\mathbb{E}[\mathbb{P}(i_*(\theta) = i | H_n, J_n = j)] \leq \frac{e^{-\frac{\|\mu_0\|_{(x_i - x_j)(x_i - x_j)^\top}^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{\|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2}}} = \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{\|x_i - x_j\|_{\text{Cov}(\hat{\mu}_n)}^2}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2}}}.$$

Computation of $\text{Cov}(\hat{\mu}_n)$. By definition of Gaussian posteriors in linear bandit in (5),

$$\text{Cov}(\hat{\mu}_n) = \text{Cov}\left(\hat{\Sigma}_n \left(\Sigma_0^{-1} \mu_0 + \frac{1}{\sigma^2} B_n\right)\right) = \hat{\Sigma}_n \text{Cov}\left(\Sigma_0^{-1} \mu_0 + \frac{1}{\sigma^2} B_n\right) \hat{\Sigma}_n = \frac{1}{\sigma^4} \hat{\Sigma}_n \text{Cov}(B_n) \hat{\Sigma}_n,$$

and

$$\begin{aligned}
\text{Cov}(B_n) &= \sum_{t \in [n]} \mathbb{V}(Y_t x_{A_t}) + \sum_{t \in [n]} \sum_{t' \in [n], t \neq t'} \text{Cov}(Y_t x_{A_t}, Y_{t'} x_{A_{t'}}) \\
&= \sum_{t \in [n]} \mathbb{V}(Y_t) x_{A_t} x_{A_t}^\top + \sum_{t \in [n]} \sum_{t' \in [n], t \neq t'} \text{Cov}(Y_t, Y_{t'}) x_{A_t} x_{A_{t'}}^\top \\
&= \sum_{k \in [K]} \omega_k n \mathbb{V}(Y_{x_k}) x_k x_k^\top + \sum_{\substack{i, j \in [K] \\ i \neq j}} \omega_i \omega_j n^2 \text{Cov}(Y_{x_i}, Y_{x_j}) x_i x_j^\top \\
&= \sum_{k \in [K]} \omega_k n (\sigma^2 + x_k^\top \Sigma_0 x_k) x_k x_k^\top + \sum_{\substack{i, j \in [K] \\ i \neq j}} \omega_i \omega_j n^2 (x_i^\top \Sigma_0 x_j) x_i x_j^\top.
\end{aligned}$$

Proof of Corollary 3.4. We first prove a useful lemma that holds for Bayesian G-optimal design.

Lemma C.9. *Let \mathcal{X} a finite set such that $|\mathcal{X}| = K$, $\xi : \mathcal{X} \rightarrow [0, 1]$ a distribution on \mathcal{X} so that $\sum_{x \in \mathcal{X}} \xi(x) = 1$, $V_n(\xi) = \sum_{x \in \mathcal{X}} \xi(x) x x^\top + \frac{\sigma^2}{n} \Sigma_0^{-1}$, $\Sigma_0 \in \mathbb{R}^{d \times d}$ a diagonal matrix, and $f(\xi) = \log \det(V_n(\xi))$. If $\xi^* = \text{argmin}_{\xi \in \Delta_K} f(\xi)$, then $\max_{x \in \mathcal{X}} \|x\|_{V_n(\xi^*)}^2 \leq d$.*

Proof of Lemma C.9. By concavity of $\xi \mapsto f(\xi)$, we have for any ξ that

$$0 \geq \langle \nabla f(\xi^*), \xi - \xi^* \rangle = \sum_{x \in \mathcal{X}} \xi(x) [\nabla f(\xi^*)]_x - \langle \xi^*, \nabla f(\xi^*) \rangle,$$

and since this holds for any pdf ξ , choosing $\xi = \delta_{x'}$ for an arbitrary action x' yields

$$[\nabla f(\xi^*)]_{x'} \leq \langle \xi^*, \nabla f(\xi^*) \rangle \quad \text{for any } x' \in \mathcal{X}.$$

Since r.h.s. does not depend on x' ,

$$\max_{x \in \mathcal{X}} [\nabla f(\xi^*)]_x \leq \langle \xi^*, \nabla f(\xi^*) \rangle. \quad (28)$$

By the property of the gradient of log-determinant, $[\nabla f(\xi^*)]_x = \|x\|_{V_n(\xi^*)}^{-1}$. Therefore, for any ξ ,

$$\begin{aligned} \langle \xi, \nabla f(\xi) \rangle &= \sum_{x \in \mathcal{X}} \xi(x) \|x\|_{V^{-1}(\xi)}^2 \\ &= \sum_{x \in \mathcal{X}} \xi(x) x^\top V_n(\xi)^{-1} x \\ &= \text{Tr} \left(\sum_{x \in \mathcal{X}} \xi(x) x^\top V_n(\xi)^{-1} x \right) \\ &= \text{Tr} \left(\sum_{x \in \mathcal{X}} \xi(x) x x^\top \left(\sum_{x' \in \mathcal{X}} \xi(x') x' x'^\top + \frac{\sigma^2}{n} \Sigma_0^{-1} \right)^{-1} \right) \\ &= \text{Tr} \left(A \left(A + \frac{\sigma^2}{n} \Sigma_0^{-1} \right)^{-1} \right) \quad \text{where } E = \sum_{x \in \mathcal{X}} \xi(x) x x^\top \\ &= \text{Tr} \left(\left(I_d + \frac{\sigma^2}{n} \Sigma_0^{-1} E^{-1} \right)^{-1} \right) \\ &= \text{Tr} \left(I_d - \frac{\sigma^2}{n} \Sigma_0^{-1} \left(I_d + \frac{\sigma^2}{n} E^{-1} \Sigma_0^{-1} \right)^{-1} E^{-1} \right) \quad (\text{Woodbury}) \\ &= \text{Tr}(I_d) - \frac{\sigma^2}{n} \text{Tr} \left(\left(E \Sigma_0 + \frac{\sigma^2}{n} I_d \right)^{-1} \right) \\ &\leq \text{Tr}(I_d) = d. \end{aligned}$$

All putting together in (28) with $\xi^* = \text{argmin}_{\xi \in \Delta_K} f(\xi)$ implies $\max_{x \in \mathcal{X}} \|x\|_{V(\xi^*)}^{-1} \leq d$. \square

A direct implication of Lemma C.9 is that $\max_{x \in \mathcal{X}} \|x\|_{\hat{\Sigma}_n} \leq \frac{d\sigma^2}{n}$. Therefore,

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{c_{i,j}}{\|x_i - x_j\|_{\hat{\Sigma}_n}^2}}} \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + \frac{c_{i,j}}{2 \max_{x \in \mathcal{X}} \|x\|_{\hat{\Sigma}_n}^2}}} \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{e^{-\frac{(\mu_0^\top x_i - \mu_0^\top x_j)^2}{2\|x_i - x_j\|_{\Sigma_0}^2}}}{\sqrt{1 + n \frac{c_{i,j}}{2d\sigma^2}}}.$$

C.4 Proofs for Hierarchical Bandits

We begin by stating the complete proof.

Theorem C.10 (Complete statement of Theorem 3.3). *For all $\omega \in \Delta_K^+$, the expected PoE of PI-BAI(ω) under the hierarchical bandit problem (6) is upper bounded as*

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{1}{\sqrt{1 + \frac{c_{i,j}}{\hat{\sigma}_{n,i} + \hat{\sigma}_{n,j}}}} e^{-\frac{(\nu^\top b_i - \nu^\top b_j)^2}{2(\|b_i - b_j\|_{\Sigma}^2 + \sigma_{0,i}^2 + \sigma_{0,j}^2)}},$$

where

$$\begin{aligned}
c_{i,j} &= \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) = \frac{\sigma_{0,i}^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \mathbb{V}(B_{n,i}) + \frac{\sigma_{0,j}^4}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \mathbb{V}(B_{n,j}) \\
&+ \frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_i)^2}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) + \frac{\sigma^4}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_j)^2}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) \\
&+ \frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \sum_{k \in [K]} \sum_{k' \in [K] \setminus \{k\}} \frac{(b_k^\top \check{\Sigma}_n b_i) \cdot (b_{k'}^\top \check{\Sigma}_n b_i)}{(\sigma_{0,k}^2 \omega_k n + \sigma^2) \cdot (\sigma_{0,k'}^2 \omega_{k'} n + \sigma^2)} \text{cov}(B_{n,k}, B_{n,k'}) \\
&+ \frac{\sigma^4}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \sum_{k \in [K]} \sum_{k' \in [K] \setminus \{k\}} \frac{(b_k^\top \check{\Sigma}_n b_j) \cdot (b_{k'}^\top \check{\Sigma}_n b_j)}{(\sigma_{0,k}^2 \omega_k n + \sigma^2) \cdot (\sigma_{0,k'}^2 \omega_{k'} n + \sigma^2)} \text{cov}(B_{n,k}, B_{n,k'}) \\
&- \frac{2\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \left[\sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_i)(b_k^\top \check{\Sigma}_n b_j)}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) + \sum_{k \in [K]} \sum_{k' \in [K] \setminus \{k\}} \frac{(b_k^\top \check{\Sigma}_n b_i) \cdot (b_{k'}^\top \check{\Sigma}_n b_j)}{(\sigma_{0,k}^2 \omega_k n + \sigma^2) \cdot (\sigma_{0,k'}^2 \omega_{k'} n + \sigma^2)} \text{cov}(B_{n,k}, B_{n,k'}) \right] \\
&- \frac{2\sigma_{0,i}^2 \sigma_{0,j}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \text{cov}(B_{n,i}, B_{n,j}) \\
&+ \frac{2\sigma^2 \sigma_{0,i}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \left[\sum_{k \in [K] \setminus \{i\}} \frac{b_k^\top \check{\Sigma}_n b_i}{(\sigma_{0,k}^2 \omega_k n + \sigma^2)} \text{cov}(B_{n,k}, B_{n,i}) + \frac{b_i^\top \check{\Sigma}_n b_i}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)} \mathbb{V}(B_{n,i}) \right] \\
&+ \frac{2\sigma^2 \sigma_{0,j}^2}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \left[\sum_{k \in [K] \setminus \{j\}} \frac{b_k^\top \check{\Sigma}_n b_j}{(\sigma_{0,k}^2 \omega_k n + \sigma^2)} \text{cov}(B_{n,k}, B_{n,j}) + \frac{b_j^\top \check{\Sigma}_n b_j}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \mathbb{V}(B_{n,j}) \right] \\
&- \frac{2\sigma^2 \sigma_{0,j}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \left[\sum_{k \in [K] \setminus \{j\}} \frac{b_k^\top \check{\Sigma}_n b_i}{\sigma^2 + \omega_k n \sigma_{0,k}^2} \text{cov}(B_{n,k}, B_{n,i}) + \frac{b_j^\top \check{\Sigma}_n b_i}{\sigma^2 + \omega_j n \sigma_{0,j}^2} \mathbb{V}(B_{n,j}) \right] \\
&- \frac{2\sigma^2 \sigma_{0,i}^2}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)(\sigma_{0,i}^2 \omega_i n + \sigma^2)} \left[\sum_{k \in [K] \setminus \{i\}} \frac{b_k^\top \check{\Sigma}_n b_j}{\sigma^2 + \omega_k n \sigma_{0,k}^2} \text{cov}(B_{n,k}, B_{n,j}) + \frac{b_i^\top \check{\Sigma}_n b_j}{\sigma^2 + \omega_i n \sigma_{0,i}^2} \mathbb{V}(B_{n,i}) \right],
\end{aligned}$$

where we defined $\check{\Sigma}_n$ from (7),

$$\check{\Sigma}_n^{-1} = \Sigma^{-1} + \sum_{k \in [K]} \frac{\omega_k n}{\sigma^2 + \omega_k n \sigma_{0,k}^2} b_k b_k^\top \quad ; \quad \mathbb{V}(B_{n,k}) = \omega_k n \sigma^2 + \omega_k^2 n^2 (\sigma_{0,k}^2 + b_k^\top \Sigma b_k) \quad ; \quad \text{cov}(B_{n,k}, B_{n,k'}) = \omega_k \omega_{k'} n^2 b_k^\top \Sigma b_{k'}.$$

Proof of Theorem C.10. This proof follows the same idea of the proof of Theorem 3.1. We first write \mathcal{P}_n as

$$\mathcal{P}_n = \sum_{\substack{i,j \in [K] \\ i \neq j}} \mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbf{1}\{J_n = j\}].$$

Following (19), by applying Hoeffding inequality for sub-Gaussian random variables,

$$\begin{aligned}
\mathbb{P} \left(\underset{k \in [K]}{\text{argmax}} \theta_k = i \mid H_n, J_n = j \right) &\leq \mathbb{P}(\theta_i \geq \theta_j \mid H_n, J_n = j) \\
&= \mathbb{P}((\theta_i - \theta_j) - (\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) \geq -(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) \mid H_n, J_n = j) \\
&\leq \exp \left(-\frac{(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)} \right),
\end{aligned}$$

where $\hat{\mu}_{n,i}$ and $\hat{\sigma}_{n,i}^2$ are given by (9). Taking the expectation with respect to the history H_n ,

$$\mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n) \mathbf{1}\{J_n = j\}] \leq \mathbb{E} \left[e^{-\frac{(\hat{\mu}_{n,j} - \hat{\mu}_{n,i})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)}} \right].$$

Since $\hat{\mu}_{n,i} - \hat{\mu}_{n,j} \sim \mathcal{N}(\mathbb{E}[\hat{\mu}_{n,i}] - \mathbb{E}[\hat{\mu}_{n,j}], \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}))$, applying Lemma C.5 gives

$$\mathbb{E} \left[e^{-\frac{(\hat{\mu}_{n,j} - \hat{\mu}_{n,i})^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)}} \right] = \frac{1}{\sqrt{1 + \frac{\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})}{\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2}}} \exp \left(-\frac{(\mathbb{E}[\hat{\mu}_{n,i}] - \mathbb{E}[\hat{\mu}_{n,j}])^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2)} \frac{1}{1 + \frac{\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})}{\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2}} \right)$$

Therefore,

$$\mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n)] \leq \frac{1}{\sqrt{1 + \frac{\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})}{\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2}}} e^{-\frac{(\mathbb{E}[\hat{\mu}_{n,i}] - \mathbb{E}[\hat{\mu}_{n,j}])^2}{2(\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2 + \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}))}}.$$

Now we want to simplify $\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2 + \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})$. on one hand, by the law of total variance,

$$\mathbb{V}(\theta_i - \theta_j) = \mathbb{E} [\mathbb{V}(\theta_i - \theta_j \mid H_n)] + \mathbb{V}(\mathbb{E}[\theta_i - \theta_j \mid H_n]) = \hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2 + \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}),$$

On the other hand,

$$\begin{aligned} \mathbb{V}(\theta_i - \theta_j) &= \mathbb{E} [\mathbb{V}(\theta_i - \theta_j \mid \mu)] + \mathbb{V}(\mathbb{E}[\theta_i - \theta_j \mid \mu]) = \sigma_{0,i}^2 + \sigma_{0,j}^2 + \mathbb{V}((b_i - b_j)^\top \mu) \\ &= \sigma_{0,i}^2 + \sigma_{0,j}^2 + \|b_i - b_j\|_\Sigma^2. \end{aligned}$$

Combining these two last equations gives $\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2 + \mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j}) = \sigma_{0,i}^2 + \sigma_{0,j}^2 + \|b_i - b_j\|_\Sigma^2$.

Therefore,

$$\mathbb{E} [\mathbb{P}(i_*(\theta) = i \mid J_n = j, H_n)] \leq \frac{1}{\sqrt{1 + \frac{\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})}{\hat{\sigma}_{n,i}^2 + \hat{\sigma}_{n,j}^2}}} e^{-\frac{(\mathbb{E}[\hat{\mu}_{n,i}] - \mathbb{E}[\hat{\mu}_{n,j}])^2}{2(\sigma_{0,i}^2 + \sigma_{0,j}^2 + \|b_i - b_j\|_\Sigma^2)}}. \quad (29)$$

Computing $\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})$.

The rest of the proof consists to compute $\mathbb{E}[\hat{\mu}_{n,i}]$ and $\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})$ for (i, j) . Denoting Q the latent prior distribution $\mu \sim Q$ and π_{H_n} the law of H_n ,

$$\begin{aligned} \pi_{H_n}(H_n) &= \pi_{H_n}(Y_{A_1}, \dots, Y_{A_n}) \\ &= \iint_{(\theta, \mu)} \mathcal{L}_\theta(Y_{A_1}, \dots, Y_{A_n}) P_0(\theta \mid \mu) Q(\mu) d\theta d\mu \\ &= \iint_{(\theta, \mu)} \prod_{i \in [K]} \mathcal{L}_{\theta_i}((Y_t)_{t \in \mathcal{T}_i}) P_{0,i}(\theta_i \mid \mu) Q(\mu) d\theta_i d\mu \\ &= \int_\mu \left[\prod_{i \in [K]} \int_{\theta_i} \mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; \theta_i \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n}) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) d\theta_i \right] Q(\mu) d\mu. \end{aligned}$$

From properties of Gaussian convolutions (Bishop, 2006),

$$\int_{\theta_i} \mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; \theta_i \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n}) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) d\theta_i = \mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; (b_i^\top \mu) \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n} + \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^\top \sigma_{0,i}^2).$$

Therefore,

$$\begin{aligned} &\prod_{i \in [K]} \int_{\theta_i} \mathcal{N}((Y_t)_{t \in \mathcal{T}_i}; \theta_i \mathbf{1}_{\omega_i n}, \sigma^2 I_{\omega_i n}) \mathcal{N}(\theta_i; b_i^\top \mu, \sigma_{0,i}^2) d\theta_i \\ &= \mathcal{N}(H_n; \sum_{i \in [K]} e_i(\mathbb{R}^K) \otimes \left(\sum_{t \in [\omega_i n]} e_t(\mathbb{R}^{\omega_i n}) \otimes b_i^\top \right) \otimes \mu, I_K \otimes (\sigma^2 I_{\omega_i n} + \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^\top \sigma_{0,i}^2)), \end{aligned}$$

where we define explicitly $e_i (\mathbb{R}^K)$ as the i^{th} base vector of \mathbb{R}^K .

Therefore,

$$\begin{aligned}\pi(H_n) &= \int_{\mu} \mathcal{N}(H_n; \sum_{i \in [K]} e_i (\mathbb{R}^K) \otimes \left(\sum_{t \in \omega_i n} e_t (\mathbb{R}^{\omega_i n}) \otimes b_i^\top \right) \otimes \mu, I_K \otimes (\sigma^2 I_{\omega_i n} + \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^\top \sigma_{0,i}^2)) \mathcal{N}(\mu; \nu, \Sigma) d\mu \\ &= \mathcal{N}(H_n; \hat{\mu}, \hat{\Sigma}),\end{aligned}$$

where $\hat{\mu} \in \mathbb{R}^n$, $\hat{\Sigma} \in \mathbb{R}^{n \times n}$ with

$$\begin{aligned}\hat{\mu} &= \sum_{i \in [K]} e_i (\mathbb{R}^K) \otimes \left(\sum_{t \in \omega_i n} e_t (\mathbb{R}^{\omega_i n}) \otimes b_i^\top \right) \otimes \nu \\ \hat{\Sigma} &= I_K \otimes (\sigma^2 I_{\omega_i n} + \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^\top \sigma_{0,i}^2) \\ &+ \left[\sum_{i \in [K]} e_i (\mathbb{R}^K) \otimes \left(\sum_{t \in [\omega_i n]} e_t (\mathbb{R}^{\omega_i n}) \otimes b_i^\top \right) \right] \Sigma \left[\sum_{i \in [K]} e_i (\mathbb{R}^K) \otimes \left(\sum_{t \in [\omega_i n]} e_t (\mathbb{R}^{\omega_i n}) \otimes b_i^\top \right) \right]^\top.\end{aligned}\quad (30)$$

The covariance matrix $\hat{\Sigma}$ seems complex but has a simple structure. The first term $I_K \otimes (\sigma^2 I_{\omega_i n} + \mathbf{1}_{\omega_i n} \mathbf{1}_{\omega_i n}^\top \sigma_{0,i}^2)$ is the same as in the standard model. The remaining term accounts for the correlation between distinct arms (i, j) , and this correlation is of the form $b_i^\top \Sigma b_j$.

Now we are ready to compute $\mathbb{E} [\hat{\mu}_{n,k}]$ for any arm $k \in [K]$: from (8) and (9),

$$\begin{aligned}\mathbb{E} [\hat{\mu}_{n,k}] &= \mathbb{E} \left[\tilde{\sigma}_{n,k}^2 \left(\frac{\check{\mu}_n^\top b_k}{\sigma_{0,k}^2} + \frac{B_{n,k}}{\sigma^2} \right) \right] \\ &= \mathbb{E} \left[\frac{\sigma^2 \sigma_{0,i}^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \left(\frac{\check{\mu}_n^\top b_k}{\sigma_{0,k}^2} + \frac{B_{n,k}}{\sigma^2} \right) \right] \\ &= \frac{\sigma^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \mathbb{E} [\check{\mu}_n^\top b_k] + \frac{\sigma_{0,k}^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \mathbb{E} [B_{n,k}].\end{aligned}$$

From (7),

$$\check{\mu}_n^\top b_k = \left(\nu^\top \Sigma^{-1} + \sum_{i \in [K]} \frac{B_{n,i}}{\sigma^2 + \omega_i n \sigma_{0,i}^2} b_i^\top \right) \check{\Sigma}_n b_k = \nu^\top \Sigma^{-1} \check{\Sigma}_n b_k + \sum_{i \in [K]} \frac{B_{n,i}}{\sigma^2 + \omega_i n \sigma_{0,i}^2} b_i^\top \check{\Sigma}_n b_k.$$

By linearity,

$$\mathbb{E} [\hat{\mu}_{n,k}] = \frac{\sigma^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \left(\nu^\top \Sigma^{-1} \check{\Sigma}_n b_k + \sum_{i \in [K]} \frac{\mathbb{E} [B_{n,i}]}{\sigma^2 + \omega_i n \sigma_{0,i}^2} b_i^\top \check{\Sigma}_n b_k \right) + \frac{\sigma_{0,k}^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \mathbb{E} [B_{n,k}].$$

From Equation (30), $\mathbb{E} [B_{n,i}] = \omega_i n \nu^\top b_i$. Therefore,

$$\mathbb{E} [\hat{\mu}_{n,k}] = \frac{\sigma^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \left(\nu^\top \Sigma^{-1} \check{\Sigma}_n b_k + \sum_{i \in [K]} \frac{\omega_i n \nu^\top b_i}{\sigma^2 + \omega_i n \sigma_{0,i}^2} b_i^\top \check{\Sigma}_n b_k \right) + \frac{\sigma_{0,k}^2}{\sigma_{0,k}^2 \omega_k n + \sigma^2} \omega_k n \nu^\top b_k = \nu^\top b_k \quad (31)$$

Now we are ready to compute $\mathbb{V}(\hat{\mu}_{n,i} - \hat{\mu}_{n,j})$ for any (i, j) . From (31),

$$\begin{aligned}
\hat{\mu}_{n,i} - \hat{\mu}_{n,j} &= \frac{\sigma^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} \check{\mu}_n^\top b_i + \frac{\sigma_{0,i}^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} B_{n,i} - \frac{\sigma^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} \check{\mu}_n^\top b_j - \frac{\sigma_{0,j}^2}{\sigma_{0,i}^2 \omega_j n + \sigma^2} B_{n,j} \\
&= \underbrace{\frac{\sigma^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} \nu^\top \Sigma^{-1} \check{\Sigma}_n b_i - \frac{\sigma^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} \nu^\top \Sigma^{-1} \check{\Sigma}_n b_j}_{\text{does not depend on observations}} \\
&\quad + \underbrace{\frac{\sigma^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} \cdot \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_i}_{(1)} + \underbrace{\frac{-\sigma^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} \cdot \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_j}_{(2)} \\
&\quad + \underbrace{\frac{\sigma_{0,j}^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} B_{n,i}}_{(3)} + \underbrace{\frac{-\sigma_{0,j}^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} B_{n,j}}_{(4)}.
\end{aligned}$$

Since (1), (2), (3) and (4) are correlated,

$$\mathbb{V}(\hat{\mu}_{n,i} + \hat{\mu}_{n,j}) = \mathbb{V}((1) + (2) + (3) + (4)) = \sum_{i=1}^4 \mathbb{V}((i)) + \sum_{i=1}^4 \sum_{j=1, j \neq i}^4 2\text{cov}((i), (j)). \quad (32)$$

We now compute each term of (32):

$$\begin{aligned}
\mathbb{V}((1)) &= \frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \mathbb{V} \left(\sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_i \right) \\
&= \frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \left(\sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_i)^2}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) + \sum_{(k,k'), k \neq k'} \text{cov} \left(\frac{B_{n,k}}{\sigma^2 + \omega_k n \sigma_{0,k}^2} b_k^\top \check{\Sigma}_n b_i, \frac{B_{n,k'}}{\sigma^2 + \omega_{k'} n \sigma_{0,k'}^2} b_{k'}^\top \check{\Sigma}_n b_i \right) \right) \\
&= \frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \left(\sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_i)^2}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) + \sum_{(k,k'), k \neq k'} \frac{(b_k^\top \check{\Sigma}_n b_i)(b_{k'}^\top \check{\Sigma}_n b_i)}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)(\sigma^2 + \omega_{k'} n \sigma_{0,k'}^2)} \text{cov}(B_{n,k}, B_{n,k'}) \right) \\
\mathbb{V}((2)) &= \frac{\sigma^4}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \left(\sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_j)^2}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)^2} \mathbb{V}(B_{n,k}) + \sum_{(k,k'), k \neq k'} \frac{(b_k^\top \check{\Sigma}_n b_j)(b_{k'}^\top \check{\Sigma}_n b_j)}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)(\sigma^2 + \omega_{k'} n \sigma_{0,k'}^2)} \text{cov}(B_{n,k}, B_{n,k'}) \right), \\
\mathbb{V}((3)) &= \frac{\sigma_{0,i}^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)^2} \mathbb{V}(B_{n,i}) \\
\mathbb{V}((4)) &= \frac{\sigma_{0,j}^4}{(\sigma_{0,j}^2 \omega_j n + \sigma^2)^2} \mathbb{V}(B_{n,j}),
\end{aligned}$$

$$\begin{aligned}
\text{cov}((1), (2)) &= -\text{cov} \left(\frac{\sigma^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} \cdot \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_i, \frac{\sigma^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} \cdot \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_j \right) \\
&= -\frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \text{cov} \left(\sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_i, \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_j \right) \\
&= -\frac{\sigma^4}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \left[\sum_{k \in [K]} \frac{(b_k^\top \check{\Sigma}_n b_i)(b_k^\top \check{\Sigma}_n b_j)}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)} \mathbb{V}(B_{n,k}) \right. \\
&\quad \left. + \sum_{(k,k'), k \neq k'} \frac{(b_k^\top \check{\Sigma}_n b_i)(b_{k'}^\top \check{\Sigma}_n b_j)}{(\sigma^2 + \omega_k n \sigma_{0,k}^2)(\sigma^2 + \omega_{k'} n \sigma_{0,k'}^2)} \text{cov}(B_{n,k}, B_{n,k'}) \right] \\
\text{cov}((3), (4)) &= -\text{cov} \left(\frac{\sigma_{0,i}^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} B_{n,i}, \frac{\sigma_{0,j}^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} B_{n,j} \right) \\
&= -\frac{\sigma_{0,i}^2 \sigma_{0,j}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \text{cov}(B_{n,i}, B_{n,j}) \\
\text{cov}((1), (4)) &= -\text{cov} \left(\frac{\sigma^2}{\sigma_{0,i}^2 \omega_i n + \sigma^2} \cdot \sum_{k \in [K]} \frac{B_{n,k}}{\omega_k n \sigma_{0,k}^2 + \sigma^2} b_k^\top \check{\Sigma}_n b_i, \frac{\sigma_{0,j}^2}{\sigma_{0,j}^2 \omega_j n + \sigma^2} B_{n,j} \right) \\
&= -\frac{\sigma^2 \sigma_{0,j}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \text{cov} \left(\sum_{k \in [K]} \frac{B_{n,k}}{\sigma^2 + \omega_k n \sigma_{0,k}^2} b_k^\top \check{\Sigma}_n b_i, B_{n,j} \right) \\
&= -\frac{\sigma^2 \sigma_{0,j}^2}{(\sigma_{0,i}^2 \omega_i n + \sigma^2)(\sigma_{0,j}^2 \omega_j n + \sigma^2)} \left[\sum_{k \in [K] \setminus \{j\}} \frac{b_k^\top \check{\Sigma}_n b_i}{\sigma_{0,k}^2 \omega_k n + \sigma^2} + \frac{b_j^\top \check{\Sigma}_n b_i}{\sigma_{0,j}^2 \omega_j n + \sigma^2} \mathbb{V}(B_{n,j}) \right].
\end{aligned}$$

The remaining terms are obtained by symmetry. Finally, for any (i, j) :

$$\begin{aligned}
\mathbb{V}(B_{n,i}) &= \mathbb{V} \left(\sum_{t \in \mathcal{A}_i} Y_t \right) = \omega_i n \sigma^2 + \omega_i^2 n^2 (\sigma_{0,i}^2 + b_i^\top \Sigma b_i) \\
\text{cov}(B_{n,i}, B_{n,j}) &= \text{cov} \left(\sum_{t \in \mathcal{A}_i} Y_t, \sum_{t \in \mathcal{A}_j} Y_t \right) = \sum_{t \in \mathcal{A}_i} \sum_{t \in \mathcal{A}_j} \text{cov}(Y_t, Y_t) = \omega_i \omega_j n^2 b_i^\top \Sigma b_j.
\end{aligned}$$

Remark C.11 (Computing the upper bound for hierarchical bandit with Theorem 3.2). The reader can wonder why transforming the hierarchical model into a linear model thanks to (10), and plug directly the transformed prior and actions to the linear upper bound (Theorem 3.2). While this is what we do to optimize numerically the bound, it is challenging to give explicit terms with this method. In fact, it would yield to the following upper bound,

$$\mathcal{P}_n \leq \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{1}{\sqrt{1 + \frac{\|\bar{b}_i - \bar{b}_j\|_{\text{Cov}_n}^2}{\|\bar{b}_i - \bar{b}_j\|_{\bar{\Sigma}_n}^2}}} e^{-\frac{(\bar{\nu}^\top \bar{b}_i - \bar{\nu}^\top \bar{b}_j)^2}{2(\|\bar{b}_i - \bar{b}_j\|_{\bar{\Sigma}_n}^2)}} = \sum_{\substack{i,j \in [K] \\ i \neq j}} \frac{1}{\sqrt{1 + \frac{\|\bar{b}_i - \bar{b}_j\|_{\text{Cov}_n}^2}{\|\bar{b}_i - \bar{b}_j\|_{\bar{\Sigma}_n}^2}}} e^{-\frac{(\nu^\top b_i - \nu^\top b_j)^2}{2(\|b_i - b_j\|_{\Sigma + \sigma_{0,i}^2 + \sigma_{0,j}^2}^2)}},$$

where

$$\bar{\Sigma}_n = \left(\bar{\Sigma} + \frac{1}{\sigma^2} \sum_{i \in [K]} \omega_i n \bar{b}_i \bar{b}_i^\top \right)^{-1}, \quad \text{Cov}_n = \frac{1}{\sigma^4} \bar{\Sigma}_n \left(\sum_{i \in [K]} \omega_i n (\sigma^2 + \bar{b}_i \bar{\Sigma} \bar{b}_i) \bar{b}_i \bar{b}_i^\top + \sum_{\substack{i,j \in [K] \\ i \neq j}} \bar{b}_i \bar{\Sigma} \bar{b}_j \omega_i \omega_j n^2 \bar{b}_i \bar{b}_j^\top \right) \bar{\Sigma}_n.$$

However, computing $\|\bar{b}_i - \bar{b}_j\|_{\text{Cov}_n}^2$ and $\|\bar{b}_i - \bar{b}_j\|_{\bar{\Sigma}_n}^2$ is computationally challenging because it requires first to compute $\bar{\Sigma}$ with block-matrix inversion, then to recover the marginal and posterior covariances $\bar{\sigma}_{n,i}^2$, $\check{\sigma}_{n,i}^2$ and $\hat{\sigma}_{n,i}^2$ from (7), (8) and (9).

D Additional Experiments

We provide additional numerical experiments on synthetic data.

- Appendix D.1 provides additional details for the MovieLens experiments.
- Appendix D.2 provides experiments for the logistic bandit model.
- Appendix D.3 provides justifications for the choice of baselines. In particular, we discuss the choice of the warm-up policy and the influence of adding elimination on top of our method.
- Appendix D.4 gives experiments when focusing as the simple regret as a metric.
- Appendix D.5 provides another type of confidence intervals on the experiments of Section 5.
- Appendix D.6 explains in which setting the hierarchical model benefits from model structure.
- Appendix D.7 tackles the problem of tuning the warm-up length n_w and the influence of the choice of π_w on the PoE.
- Appendix D.8 provides toy example when deriving ω^{opt} .

D.1 MovieLens Experiments

We provide more information on our MovieLens experiments in Figure 4. The MovieLens dataset contains ratings given by 6040 users to 3952 movies. We use a subset of $K = 100$ randomly picked movies for our experiments. The prior used for inference in PI-BAI and BayesGap is set to be Gaussian with mean μ_0 and Σ_0 . These parameters are estimated by taking the empirical mean and empirical covariance over the wall dataset. All results are averaged over 10^4 rounds.

D.2 Logistic Bandits

We consider two main settings as in Section 5. In the Fixed setting, μ_0 is flat, $\mu_0 = (1, \dots, 1)$ whereas in the Random setting, the prior means are uniformly sampled from $[0, 1]$. For both settings, $\Sigma_0 = \text{diag}(\sigma_{0,i}^2)_{i \in [K]}$ where the $\sigma_{0,i}$'s are evenly spaced between 0.1 and 0.5. We ran experiments for $K = 30$ arms and $d \in \{3, 4\}$. Figure 5 shows that the generalization of PI-BAI with G-optimal design allocations on has good performances beyond linear settings.




Figure 5: Average PoE with varying budgets for fixed and randomized settings in the GLB framework.

D.3 Choice of Baselines

A remark on TTTS. Top two sampling algorithms is a family of algorithms that is known to have good performances in BAI. In Section 5, we used TS-TCI with $\beta = 0.5$ from Jourdan et al. (2022) and denoted it as TTTS for sake of notation simplicity.

Choice of warm-up policy. We evaluate different warm-up policies, TS and two Top-Two algorithms, TSTCI and T3C from Jourdan et al. (2022). The experiments shown in Figure 6 are run in the same setting as in Section 5, with $K = 10$ arms in the MAB setting, and with $K = 60$ and $d = 4$ in the hierarchical setting. Figure 6 suggests to pick TS as a warm-up policy for the MAB setting and $meTS$ for the hierarchical setting.

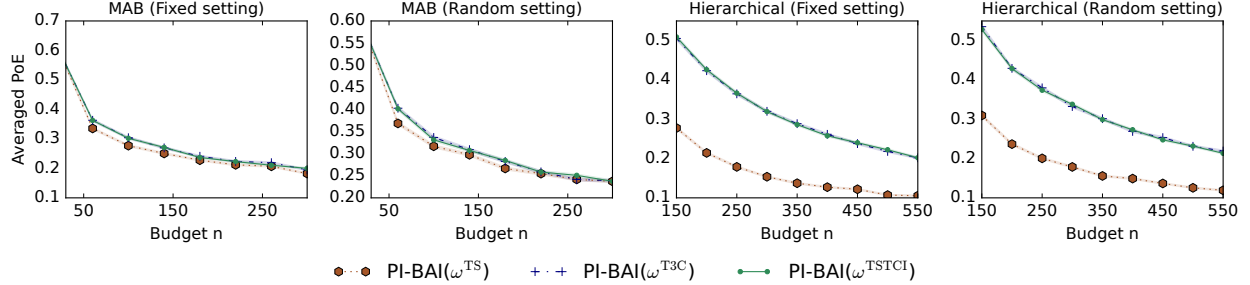


Figure 6: Average PoE of PI-BAI instantiated with different warm-up policies.

Influence of elimination. We empirically compare the influence of using elimination on top of our methods. The elimination procedure is the same as the one used in Atsidakou et al. (2022). There are $\lfloor \log_2(K) \rfloor$ rounds, and each lasts $\lfloor \frac{n}{R} \rfloor$ steps. At each round, we pull each remaining arm $i \lfloor \frac{\omega_i n}{R} \rfloor$ times. At the end of the round, half of arms are eliminated. These correspond to the arms that have the least posterior mean reward (so $\hat{\mu}_{n,i}$ in the MAB setting). The allocation ω is then normalized to allocate more budget to remaining arms. Note that we draft all observations at the end of each round, as it is the case in (Karnin et al., 2013; Azizi et al., 2021; Atsidakou et al., 2022). Figure 7 shows that using elimination does not give better performances, and hence we chose to not add these baselines in Section 5.

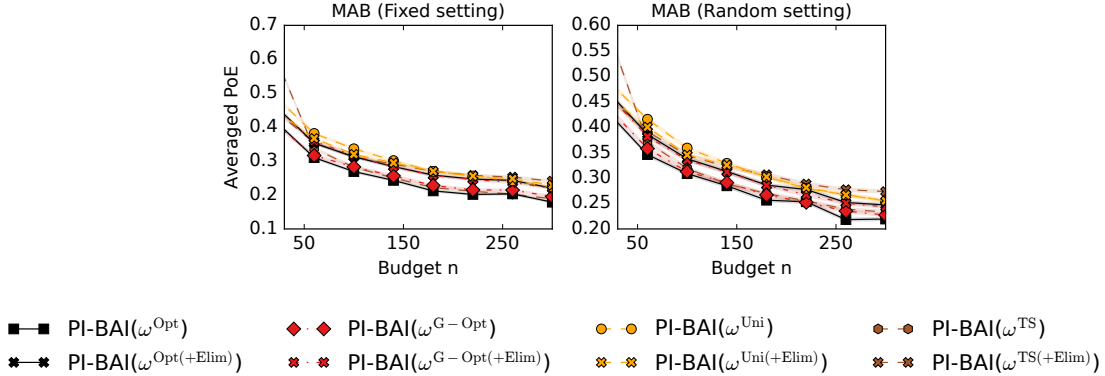


Figure 7: Average PoE of PI-BAI instantiated with different weights with or without elimination in the MAB setting.

D.4 Simple Regret

Figure 8 compares the performances of our methods based on the Bayesian simple regret $\mathbb{E} [\max_{i \in [K]} \theta - \theta_{J_n}]$, where the expectation is taken with respect to the prior distribution over instances θ . Overall, it shows that our methods also have low simple regret in these settings.

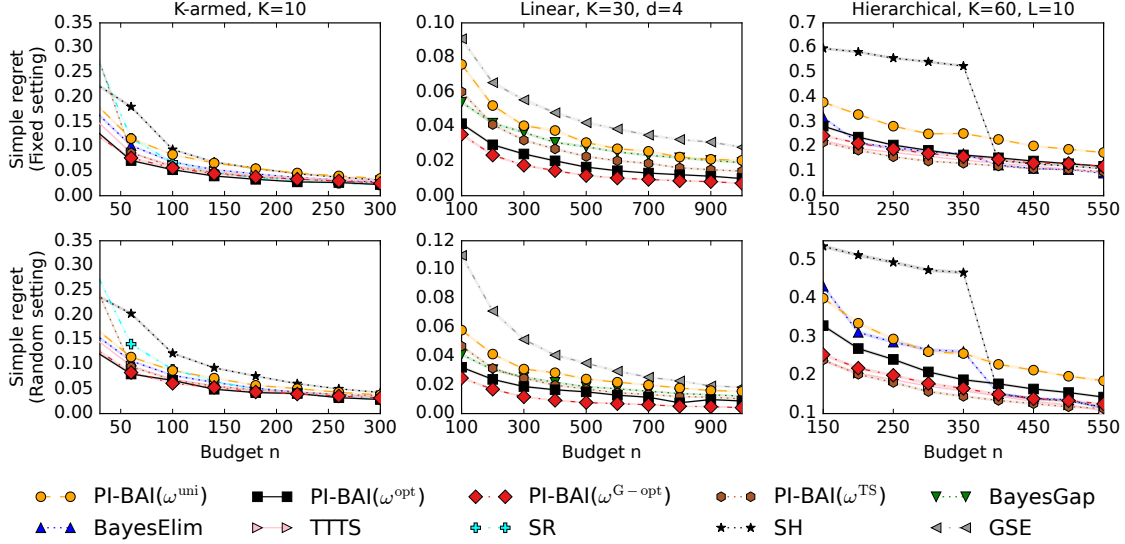


Figure 8: Average simple regret with varying budgets for fixed and randomized settings.

D.5 Confidence Intervals on Sampled Instances

We provide additional plots in the same settings of Section 5. The first row of Figure 9 shows the PoE of the methods averaged over 1000 different instances sampled from the prior distribution. For each instance, we repeat the experiments 100 times to get an estimate of the probability. We show one standard deviation around the averaged mean of PoE over instances. In the second row of the same figure, we plot the PoE of each method subtracted by the PoE of the method having the least PoE in each setting, that is, PI-BAI(ω^{opt}) in MAB, PI-BAI($\omega^{\text{G-opt}}$) in linear bandits and PI-BAI(ω^{TS}) in hierarchical bandits.

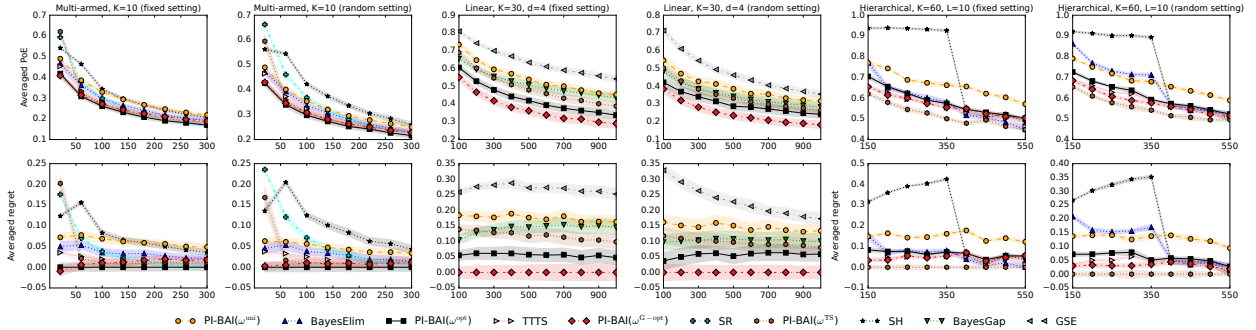


Figure 9: PoE of several methods (first row) and PoE of each method subtracted by the PoE of the most performing method in each setting (second row). For each instance, we repeat the experiments 100 times and we average the results over 1000 instances. The confidence intervals show one standard deviation.

D.6 Benefits of Hierarchical Models

To illustrate the benefits of using a hierarchical structure (6), we compare the posterior variances $\hat{\sigma}_{n,i}^2$ under a standard model and a hierarchical model. The standard model is obtained by marginalizing over the effects μ ,

$$\begin{aligned}
 \theta_i &\sim \mathcal{N}(b_i^\top \nu, \sigma_{0,i}^2 + b_i^\top \Sigma b_i) & \forall i \in [K] \\
 Y_t \mid \mu, \theta, A_t &\sim \mathcal{N}(\theta_{A_t}, \sigma^2) & \forall t \in [n].
 \end{aligned} \tag{33}$$

From (3), the corresponding posterior covariance of an arm $i \in [K]$ is

$$\hat{\sigma}_{n,i}^{-2} = \frac{1}{\sigma_{0,i}^2 + b_i^T \Sigma b_i} + \frac{\omega_i n}{\sigma^2}.$$

For the first setting, we uniformly draw a vector $u \in [0, 1]$ and set $\sigma_0 = 0.1u$ and $\Sigma = 2I_L$. For the second setting, we set $\sigma_0 = u$ and $\Sigma = 10^{-3}I_L$. In both settings, we consider $K = 50$ arms, and $L = 10$ effects. We draw uniformly $\nu, b_i \in [-1, 1]^L$ and the allocation vector is set to uniform allocation $\omega_i^{\text{uni}} = \frac{1}{K}$ for any $i \in [K]$. In Figure 10 we plot the average posterior covariance $\frac{1}{K} \sum_{i \in [K]} \sigma_{n,i}^2$ across all arms for both the (marginalized) standard model (33) and the hierarchical model (6). The goal of this experiment is to show for which setting the benefits of the hierarchy are pronounced. The results show that this difference is more pronounced when the initial uncertainty of the effects Σ is greater than the initial uncertainty of the mean rewards $(\sigma_{0,i}^2)_{i \in [K]}$.

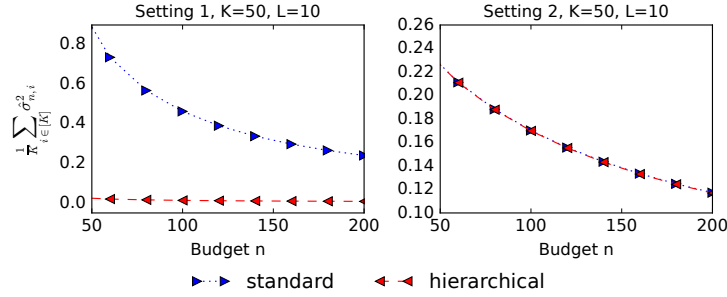


Figure 10: Average posterior covariance across all arms for standard and hierarchical model for two settings.

D.7 Hyperparameters

Warm-up length n_w . We try different values of warm-up length n_w for our warm-up policies. We emphasize that methods based on TTTS require $n_w > K$ because each arm has to be pulled at the beginning. Figure 11 suggests picking $n_w = 2K$ for the warm-up with T3C and TSTCI, and $n_w = K$ for the warm-up with TS.

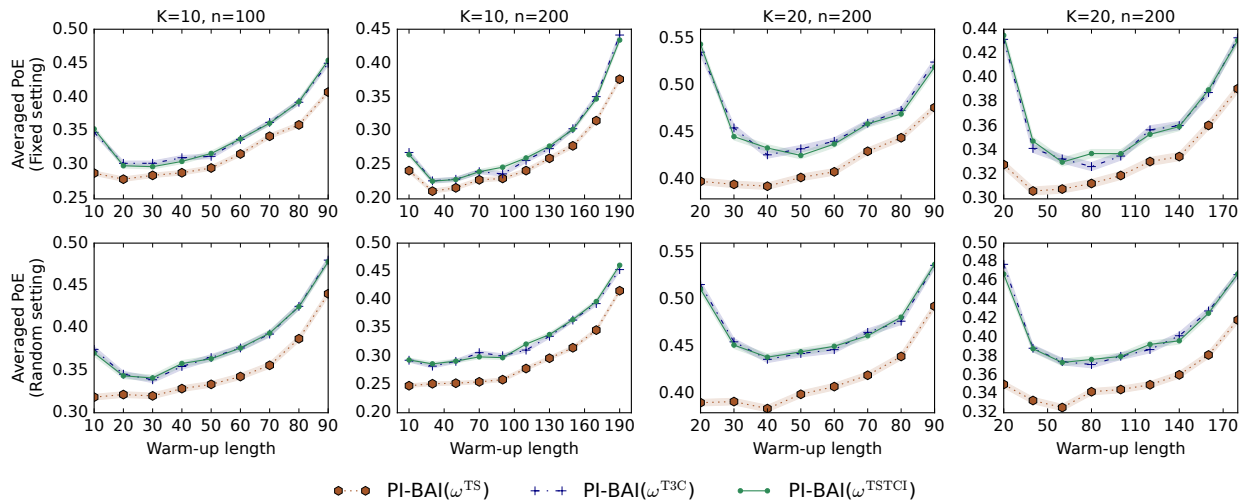


Figure 11: Average PoE of PI-BAI instantiated with different warm-up policies for different warm-up lengths n_w .

Mixture parameter α . We discuss the choice of the mixture parameter α . We recall that we use the heuristic $\alpha\omega_i^{\text{opt}} + (1-\alpha)\frac{\mu_{0,i}\sigma_{0,i}}{\sum_{k\in[K]}\mu_{0,k}\sigma_{0,k}}$ in our experiments. Figure 12 shows that for the fixed setting, adding the vector $\frac{\mu_{0,i}\sigma_{0,i}}{\sum_{k\in[K]}\mu_{0,k}\sigma_{0,k}}$ helps improve the performances. This is not necessarily the case in the random setting.




Figure 12: Average PoE of PI-BAI(ω^{opt}) for different mixture parameter α .

D.8 Toy Experiments for ω^{opt}

We provide additional experiments to evaluate the optimized weights ω^{opt} in different settings. In Figure 13, we set $K = 3$, $\mu_0 = (2, 1.9, 0)$, $(\sigma_{0,1}, \sigma_{0,2}, \sigma_{0,3}) = (10^{-2}, 0.5, 0.5)$. This corresponds to the motivating setting depicted in the introduction (Section 1). We provide a comprehensive illustration of the prior bandit instance $\mu_{0,i} \pm 2\sigma_{0,i}$ (left plot). Then we let the budget vary, $n \in [10, 200]$, and for each n we (numerically) optimize (11) to get ω_n^{opt} (middle plot). On the right plot, we let the prior mean of arm 2 vary, $\mu_{0,2} \in [0, 2]$, and get ω^{opt} as a function of $\mu_{0,2}$. In Figure 14, we do the same experiments for 6 arms with 2 good arms *a priori*, $\mu_0 = (2, 1.9, 1, 0.6, 0.3, 0)$ and σ_0 uniformly spaced in $[0.5, 0.1]$. As n increases, ω^{opt} suggests distributing roughly one third of the budget for each arm 1 and 2, and the remaining to the rest of the arms.

These results give insight on the behavior of ω^{opt} . In Figure 13, for small budgets, ω^{opt} suggests pulling a lot the second arm because of its wide prior confidence $\sigma_{0,2}^2 \gg \sigma_{0,1}^2$. It does not give much allocation to the last arm since it is statistically unlikely to become optimal. As the budget n grows, most of the allocations are almost equal to arm 1 and arm 2. Since $\sigma_{0,1}^2 \neq \sigma_{0,2}^2$, $\omega_1^{\text{opt}} \neq \omega_2^{\text{opt}}$ even for a large budget $n = 200$. The right plot of Figure 13 shows that ω^{opt} depends on prior gaps. Interestingly, this was not the case for the very special example $K = 2$ depicted in Section 3.3.

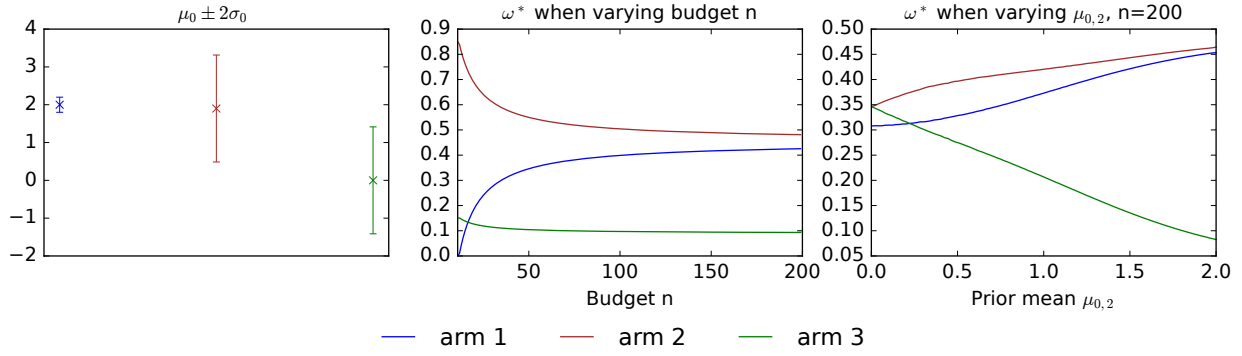


Figure 13: Illustration of the prior bandit instance (left), and optimized allocation ω^{opt} when varying budget n (middle) or one coordinate of μ_0 (right).

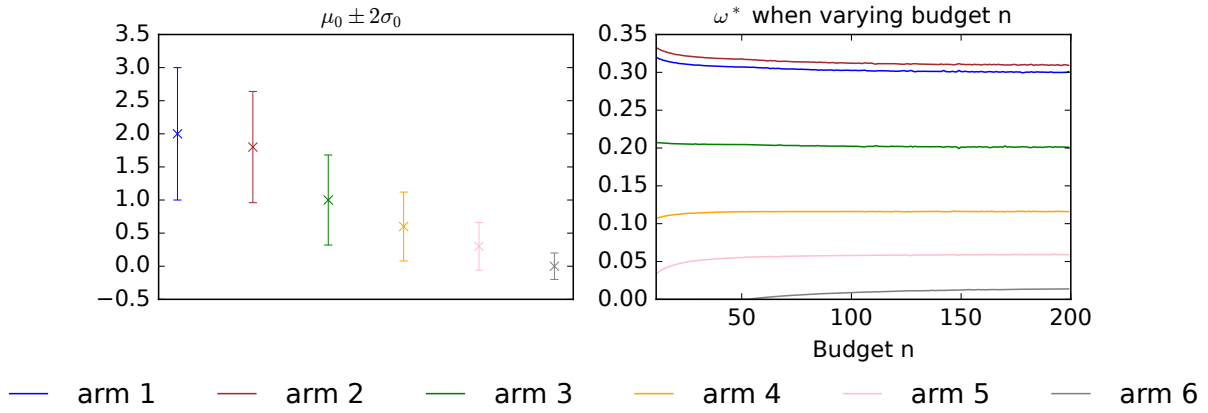


Figure 14: Illustration of the prior bandit instance (left), and optimized allocation ω^{opt} when varying budget n (right).