Sharp Non-Asymptotic Analysis of the Penalized Challenger in β-EB-TCI for Bernoulli Bandits

Nam Nguyen[1,2] **Tuan Dam**[1*]
[1]Hanoi University of Science and Technology, Hanoi, Vietnam
[2]Center for AI Research, VinUniversity, Hanoi, Vietnam
26nam.n@vinuni.edu.vn, tuandq@soict.hust.edu.vn

Abstract

Top-two algorithms are simple and effective for fixed-confidence best-arm identification, but their sharp non-asymptotic behavior is still not well understood. We study this question for Bernoulli bandits through β-EB-TCI, the empirical-best top-two rule of Jourdan et al. [9], whose challenger minimizes a Bernoulli transportation cost plus a logarithmic penalty in its pull count. We first analyze the stabilized regime, in which the empirical leader has become the true best arm and its sampling fraction stays close to β. In this regime, every suboptimal arm is sampled linearly often, and, with high probability, the stopping time is at most $T^\star_\beta(\mu)\log(1/\delta)$ plus stabilization and lower-order concentration terms, with the exact fixed-β leading coefficient. The main remaining difficulty is therefore to control when the regime is reached. For the original rule, we obtain a non-asymptotic high-probability stopping bound on every Bernoulli instance with a unique best arm. Under a finite-mean sufficient-exploration condition, the bound yields the sharp expected sample complexity; combined with the exploration result of Jourdan et al. [9], this covers pairwise-distinct means under their tie-breaking convention. A mild forced-exploration guard, which uses only $O(\sqrt{Kt})$ pulls up to time t, gives a self-contained sharp expected bound for any number of arms with a unique best arm. We also identify a limitation of proof strategies that try to handle equal suboptimal means through a single index-comparison argument.

1 Introduction

Stochastic bandits model sequential decision-making under uncertainty. At each round, a learner chooses an arm, observes a reward, and uses the accumulated data to decide what to sample next [15]. We study fixed-confidence best-arm identification for a Bernoulli bandit with a unique best arm [3, 14]. A strategy stops at a random time τ_δ and recommends an arm, which must be the best one with probability at least $1-\delta$. The goal is to minimize the expected number of samples while preserving both termination and the target error probability. We use *finite-confidence* for guarantees that hold at each fixed δ, rather than only in the asymptotic limit $\delta \to 0$.

A central benchmark is the instance-dependent characteristic time $T^\star(\mu)$. Every δ-correct strategy needs at least about $T^\star(\mu)\log(1/\delta)$ samples in expectation as $\delta \to 0$, and $1/T^\star(\mu)$ is the largest rate at which an allocation of samples can accumulate evidence against the hardest competitor [14, 4]. Track-and-Stop attains this benchmark by tracking an estimated optimal allocation, which requires repeatedly solving a global optimization problem.

Top-two algorithms provide a simpler alternative [17, 18, 9]. They maintain a current leader and compare it with one challenger at a time; in the randomized rules considered here, the leader is

*Corresponding author.

40th Conference on Neural Information Processing Systems (NeurIPS 2026).

sampled with probability $\beta \in (0,1)$ and the challenger otherwise. This structure is easy to implement and has strong asymptotic guarantees. Once the leader remains the best arm, its sampling fraction converges to β, so the natural benchmark becomes the fixed-β characteristic time $T^\star_\beta(\mu) \geq T^\star(\mu)$: the best achievable when the remaining fraction $1-\beta$ is split optimally among the competitors [17, 9] (Section 3.3).

In particular, Jourdan et al. [9] prove asymptotic β-optimality, $\limsup_{\delta\to 0} \mathbb{E}_\mu[\tau_\delta]/\log(1/\delta) \leq T^\star_\beta(\mu)$, for several top-two rules on bounded rewards, including Bernoulli β-EB-TCI with uniform tie-breaking when all means are distinct. Non-asymptotic results with the exact leading coefficient are more limited: the limit does not control the remainder at a fixed δ, and high-probability stopping bounds alone do not control the expectation [2]. The closest existing analysis is the Gaussian Top-Two UCB (TTUCB) result of Jourdan and Degenne [8], which treats a different top-two rule and a different reward model. This leaves the following question. *Does Bernoulli β-EB-TCI stop after $T^\star_\beta(\mu)\log(1/\delta)$ samples up to explicit lower-order terms at each fixed δ?*

We analyze the exact Bernoulli β-EB-TCI rule, where EB-TCI stands for Empirical-Best leader with Transportation-Cost-Improved challenger. Under the unguarded rule, the empirical-best arm is sampled with probability β at each round. Otherwise, the algorithm samples the challenger against which the Bernoulli generalized likelihood-ratio (GLR) evidence is weakest, after adding the logarithm of its pull count. This term favors less-sampled challengers when the evidence values are comparable. It is part of the algorithm, not an auxiliary proof term.

A finite-confidence analysis of this rule must address three coupled questions. First, how closely do the empirical GLR statistics follow their population counterparts under adaptive sampling? Second, once the true best arm is the persistent leader, does the penalized challenger rule sample every suboptimal arm often enough? Third, how long does it take until the leader stays correct and the best-arm sampling fraction stays near β? Since the error guarantee of Jourdan et al. [9] concerns runs that stop, termination must also be proved. We answer the first two questions in the *stabilized regime*, using a time-uniform comparison of empirical and population GLR statistics (Theorem 4.2) and a discrete benchmark (Proposition 4.1) showing that the penalty adds a common $\log t$ to the continuous rate $t/T^\star_\alpha(\mu)$, up to a bounded remainder. We then bound the time needed to reach this regime.

For Bernoulli means $\mu \in (0,1)^K$ with a unique best arm and fixed β, our expected-time bound reads informally, for every $\delta \in (0, 1/2]$,

$$\mathbb{E}_\mu[\tau_\delta] \leq T^\star_\beta(\mu)\log(1/\delta) + O\Big(\sqrt{\log(1/\delta)\,\log\log(e/\delta)}\Big),$$

under finite-mean sufficient exploration for the unguarded rule: every arm's count is at least a positive constant times $\sqrt{t/K}$ after a random time with finite mean. An exploration guard, which pulls a least-sampled arm whenever some count falls below $\sqrt{t/K}$ and uses $O(\sqrt{Kt})$ pulls by time t, gives the same bound for every such instance.

Contributions.

1. **Stopping and count growth after stabilization.** High-probability stopping bounds have the exact fixed-β leading coefficient, as do expected bounds when stabilization admits a summable upper-quantile bound (Theorem 5.2); each suboptimal count grows linearly after a transient (Proposition 5.3). Throughout, *sharp* refers only to this exact leading coefficient; we do not optimize lower-order constants or prove a matching non-asymptotic lower bound.
2. **Whole-run guarantees.** We prove a high-probability stopping bound for every unique-best instance (Corollary 6.4) and the expected bounds above (Theorems 6.5 and 6.8). Combined with Jourdan et al. [9], the unguarded expected bound covers pairwise-distinct means under their tie-breaking convention, which does not include our deterministic default (Corollary 6.6).
3. **A limitation of one-step comparisons.** With equal-mean suboptimal arms and logarithmically many pulls of a higher-mean arm, a population index comparison required to hold for all sufficiently large leader counts requires a number of pulls of an equal-mean arm that is polynomial in the inverse confidence level (Proposition 6.9).

Section 3 defines the problem and algorithm; Sections 4–6 develop the tools, stopping bounds, and entry guarantees; Section 7 reports experiments.

2 Related Work

We review fixed-confidence best-arm identification and the concentration tools used in our analysis.

Fixed-confidence identification and allocation tracking. Fixed-confidence best-arm identification has progressed from elimination methods to information-theoretic, instance-dependent analyses [16, 3, 1, 11]. The information-theoretic lower bounds of Kaufmann et al. [14] and the characteristic-time formulation and Track-and-Stop rule of Garivier and Kaufmann [4] turn these bounds into an adaptive rule that tracks a globally optimal allocation. Methods using lower and upper confidence bounds (LUCB), including KL-based refinements, form another family [10, 12, 7]. We fix the best-arm sampling fraction at β and use local leader–challenger updates rather than tracking the globally optimal allocation.

Top-two methods. Top-two algorithms repeatedly compare a current leader with a selected challenger. They were introduced in Bayesian form by Russo [17], studied in fixed-confidence settings by Shang et al. [18], and developed into a frequentist bounded-reward framework by Jourdan et al. [9]. In that framework, β-EB-TCI uses the empirical leader. We specialize the rule to Bernoulli rewards and analyze the challenger index with its exact logarithmic count penalty, rather than an unpenalized rule, a Gaussian approximation, or a different top-two construction. The closest existing non-asymptotic analysis is the Gaussian TTUCB result of Jourdan and Degenne [8].

Time-uniform concentration and stopping-time tails. Our concentration arguments use mixture martingales and time-uniform confidence bounds, or confidence sequences [13, 5, 6]. We convert armwise KL confidence sequences into time-uniform control of pairwise Bernoulli GLR statistics on the effective-sample-size scale of the two arms. Correctness and high-probability stopping bounds alone do not imply a sharp expected stopping-time bound [2]. We use the GLR stopping-rule error guarantee of Jourdan et al. [9] and prove termination and finite-time challenger-count bounds here. Appendix A records the role of each result used from prior work.

3 Problem Setting and Algorithm

We define the bandit model, evidence statistics, allocation benchmarks, and algorithm.

3.1 Problem Formulation

Consider $K \geq 2$ Bernoulli arms, indexed by $[K] := \{1, \dots, K\}$, with means $\mu = (\mu_1, \dots, \mu_K) \in (0,1)^K$. We assume a unique best arm, $i^\star = \arg\max_{i \in [K]} \mu_i$. The algorithm first pulls each arm once. After $t \geq K$ pulls, $N_{t,i}$ and $\widehat{\mu}_{t,i}$ denote arm i's pull count and empirical mean; I_{t+1} denotes the arm sampled at round $t+1$. Unless stated otherwise, all finite-set $\arg\max$ and $\arg\min$ operations use fixed deterministic tie-breaking (smallest-index in experiments), with the exception stated in Corollary 6.6.

A strategy samples adaptively, stops at a stopping time τ_δ for its observation and randomization history, and recommends $\widehat{i}_{\tau_\delta}$. Write $\mathbb{P}_\mu$ for the law of this history on instance μ. For $\delta \in (0,1)$, the strategy is δ-correct if, for every instance μ with a unique best arm,

$$\mathbb{P}_\mu(\tau_\delta < \infty) = 1 \quad \text{and} \quad \mathbb{P}_\mu(\widehat{i}_{\tau_\delta} \neq i^\star) \leq \delta.$$

The goal is to minimize the expected number of samples subject to these conditions. *Finite-confidence* guarantees hold at each fixed δ, not only as $\delta \to 0$.

3.2 Evidence Statistics

The Bernoulli divergence, measuring the cost of explaining mean p by q, is defined for $p \in [0,1]$ and $q \in (0,1)$ by

$$d(p,q) := p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q},$$

with $0 \log 0 = 0$ and $1 \log 1 = 0$. Here $\mathrm{Ber}(p)$ denotes the Bernoulli law with mean p. For interior arguments, $d(p,q) = \mathrm{KL}(\mathrm{Ber}(p), \mathrm{Ber}(q))$, the Kullback–Leibler divergence between these laws. All logarithms are natural.

Evidence that distinct arms $i, j \in [K]$ have different means is measured by the symmetric two-sample GLR statistic $V_t(i,j) := \inf_{x\in(0,1)} \{N_{t,i} d(\widehat{\mu}_{t,i}, x) + N_{t,j} d(\widehat{\mu}_{t,j}, x)\}$. The infimum is unchanged over $[0,1]$ with the standard Bernoulli-KL boundary convention. Replacing empirical means by true means isolates allocation information in the population counterpart $V_t^{\mu}(i,j) := \inf_{x\in(0,1)} \{N_{t,i} d(\mu_i, x) + N_{t,j} d(\mu_j, x)\}$.

The one-sided transportation cost of Jourdan et al. [9] counts evidence that i beats j only under that empirical ordering:

$$W_t(i,j) := \begin{cases} V_t(i,j), & \widehat{\mu}_{t,i} \geq \widehat{\mu}_{t,j}, \\ 0, & \widehat{\mu}_{t,i} < \widehat{\mu}_{t,j}. \end{cases}$$

If i is the empirical leader and $j \neq i$, then $W_t(i,j) = V_t(i,j)$. The symmetric GLR is minimized at the sample-size-weighted pooled mean, with the boundary convention above: $x_t(i,j) = \frac{N_{t,i}\widehat{\mu}_{t,i} + N_{t,j}\widehat{\mu}_{t,j}}{N_{t,i}+N_{t,j}}$. For $V_t^{\mu}(i,j)$, the pooled mean uses true rather than empirical means. Appendix B.1 derives these identities.

3.3 Characteristic Times

Prior work gives the leading lower-bound term $T^\star(\mu)\log(1/\delta)$ for the expected number of samples used by δ-correct strategies as $\delta \to 0$ [14, 4]. The fixed-β benchmark of Jourdan et al. [9] prescribes the best arm's sampling fraction.

Allocate fraction α of the samples to the best arm and split $1-\alpha$ among competitors. If arm j receives fraction ω, its population GLR grows at rate $I_j^\alpha(\omega)$ per sample, defined below. With $K = 2$ and arm 1 best, all remaining samples go to arm 2, giving $1/T_\alpha^\star(\mu) = I_2^\alpha(1-\alpha)$.

For $\alpha \in (0,1)$, $j \in [K] \setminus \{i^\star\}$, and scalar $\omega > 0$, define the pairwise information rate by

$$I_j^\alpha(\omega) := \inf_{x\in(0,1)} \{\alpha d(\mu_{i^\star}, x) + \omega d(\mu_j, x)\}.$$

For the allocation vector $w = (w_j)_{j\neq i^\star}$, the feasible set and fixed-fraction characteristic time are

$$\Delta_{1-\alpha} := \left\{ w \in (0,\infty)^{K-1} : \sum_{j\neq i^\star} w_j = 1-\alpha \right\},$$
$$T_\alpha^\star(\mu)^{-1} := \max_{w\in\Delta_{1-\alpha}} \min_{j\neq i^\star} I_j^\alpha(w_j).$$

The global characteristic time is $T^\star(\mu) := \min_{\alpha\in(0,1)} T_\alpha^\star(\mu)$.

For more than two arms, the smallest pairwise rate limits progress. The best split equalizes these rates (Appendix C.2), giving the leading sample requirement $T_\beta^\star(\mu)\log(1/\delta)$ at best-arm fraction β. Since $T_\beta^\star(\mu) \geq T^\star(\mu)$, fixing β can increase this requirement. We keep β fixed throughout; online estimation or adaptive choice of β is outside our scope.

3.4 The β-EB-TCI Algorithm

We now specify the leader–challenger comparisons and stopping rule. After t pulls, the empirical leader is $B_{t+1} \in \arg\max_{i\in[K]} \widehat{\mu}_{t,i}$. Set $\widehat{i}_t := B_{t+1}$ as the recommendation. On an unguarded round before stopping, the challenger minimizes evidence plus a logarithmic count penalty, favoring less-sampled arms when evidence is comparable: $C_{t+1} \in \arg\min_{j\in[K]\setminus\{B_{t+1}\}} \{W_t(B_{t+1}, j) + \log N_{t,j}\}$. On an unguarded round, pull the leader with probability $\beta \in (0,1)$ and the challenger with probability $1-\beta$; after stabilization, β controls the true best arm's sampling fraction.

The optional *exploration guard* instead pulls a least-sampled arm when $\min_{i\in[K]} N_{t,i} < \sqrt{t/K}$. The leader is still defined; the challenger is used only on unguarded rounds. Write G_t^{tot} for the number of guarded pulls by time t; the *unguarded algorithm* omits the guard. The GLR stopping threshold of Jourdan et al. [9] is

$$c(t,\delta) = \log(1/\delta) + 2\log(1 + t/2) + 2 + \log(K-1). \tag{3.1}$$

The stopping time and terminal recommendation are

$$\tau_\delta := \inf\left\{t \geq K : \min_{j \neq B_{t+1}} W_t(B_{t+1}, j) > c(t,\delta)\right\}, \qquad \widehat{i}_{\tau_\delta} := B_{\tau_\delta+1}. \tag{3.2}$$

For bounded rewards and any adaptive sampling rule, Jourdan et al. [9] give

$$\mathbb{P}_\mu\left(\tau_\delta < \infty, \widehat{i}_{\tau_\delta} \neq i^\star\right) \leq \delta.$$

This bounds the probability of a wrong recommendation on runs that stop. The almost-sure termination results in Sections 5–6 then give δ-correctness in the covered regimes (Appendix B). Algorithm 1 summarizes the rule.

Algorithm 1 β-EB-TCI with GLR stopping and optional exploration guard.

Pull each arm once.
for $t = K, K+1, \ldots$ **do**
  $B_{t+1} \in \arg\max_{i \in [K]} \widehat{\mu}_{t,i}$ *fixed deterministic tie-breaking*
  $\widehat{i}_t \leftarrow B_{t+1}$
  if $\min_{j \neq B_{t+1}} W_t(B_{t+1}, j) > c(t,\delta)$ **then**
    stop and recommend $\widehat{i}_t$
  end if
  if the guard is enabled and $\min_{i \in [K]} N_{t,i} < \sqrt{t/K}$ **then**
    $I_{t+1} \in \arg\min_{i \in [K]} N_{t,i}$ *guarded pull*
  else
    $C_{t+1} \in \arg\min_{j \in [K] \setminus \{B_{t+1}\}} \{W_t(B_{t+1}, j) + \log N_{t,j}\}$
    Draw $U_{t+1} \sim \mathrm{Ber}(\beta)$ independently of the past.
    $I_{t+1} \leftarrow B_{t+1}$ if $U_{t+1} = 1$, and $I_{t+1} \leftarrow C_{t+1}$ otherwise.
  end if
end for

4 Population Allocation Benchmark and Uniform GLR Control

We develop a population allocation benchmark and a uniform empirical GLR comparison to control challenger evidence once the empirical leader is correct. Throughout, bound constants $C_{\ldots}$ are positive and finite unless stated otherwise; their subscripts and arguments specify dependence, with K included in μ. Values may differ between results.

4.1 A Discrete Penalized Allocation Benchmark

The algorithm's integer counts and logarithmic count penalty require a discrete counterpart to $T^\star_\alpha(\mu)$ from Section 3.3. Let integers $n \geq 1$ and $s \geq K-1$ be the best-arm pull count and total suboptimal-arm pull count. For each suboptimal arm j and integer $m \geq 1$, define $\Phi_j(n,m) := \inf_{x \in (0,1)} \{nd(\mu_{i^\star}, x) + md(\mu_j, x)\} + \log m$. For count vectors $m = (m_j)_{j \neq i^\star}$, define the discrete penalized allocation benchmark

$$\Gamma^{\mathrm{disc}}_{n,s}(\mu) := \max_{m \in \{1,2,\ldots\}^{K-1} : \sum_{j \neq i^\star} m_j = s} \min_{j \neq i^\star} \Phi_j(n, m_j).$$

This benchmark gives the largest minimum population index at the prescribed counts, including β-EB-TCI's logarithmic count penalty. It depends on true means unavailable to the learner.

Proposition 4.1 (Approximation of the discrete penalized benchmark)**.** *Fix a compact interval $U \Subset (0,1)$. There exists a constant $C_\mu(U) < \infty$ such that for all integers $n \geq 1$ and $s \geq K-1$, if $t = n+s$, $\alpha = n/t$, and $\alpha \in U$, then*

$$\left|\Gamma^{\mathrm{disc}}_{n,s}(\mu) - \left(\frac{t}{T^\star_\alpha(\mu)} + \log t\right)\right| \leq C_\mu(U).$$

Writing $w_j = m_j/t$ separates the continuous contribution $t/T^\star_\alpha(\mu)$ from the common penalty $\log t$. Rounding and arm-dependent penalty terms have bounded cost uniformly for $\alpha \in U$ (Appendix C.3).

4.2 Uniform Comparison of Empirical and Population Pairwise GLR Statistics

We next control the error from replacing true means by adaptive empirical means. For $t \geq K$ and $\eta \in (0,1)$, define the time-uniform confidence term $H_t(\eta) := \log(K/\eta) + \log\log(et)$ and, for distinct arms $i, j \in [K]$, the two-arm effective sample size $N_{\mathrm{eff},t}(i,j) := \frac{N_{t,i}N_{t,j}}{N_{t,i}+N_{t,j}}$. It satisfies $\frac{1}{2}\min\{N_{t,i}, N_{t,j}\} \leq N_{\mathrm{eff},t}(i,j) \leq \min\{N_{t,i}, N_{t,j}\}$ and sets the fluctuation scale in the next theorem.

Theorem 4.2 (Time-uniform pairwise Bernoulli GLR approximation)**.** *There exists an instance-dependent constant $C_\mu < \infty$ such that for every $\eta \in (0,1)$ there is an event $\mathcal{G}_\eta$ with $\mathbb{P}_\mu(\mathcal{G}_\eta) \geq 1-\eta$ such that, on $\mathcal{G}_\eta$, simultaneously for all $t \geq K$ and all pairs $i, j \in [K]$ with $i \neq j$,*

$$|V_t(i,j) - V_t^\mu(i,j)| \leq C_\mu \left(\sqrt{N_{\mathrm{eff},t}(i,j) H_t(\eta)} + H_t(\eta) \right).$$

For distinct means, $V_t^\mu(i,j)$ grows linearly in $N_{\mathrm{eff},t}(i,j)$, so the error is lower order when $H_t(\eta) = o(N_{\mathrm{eff},t}(i,j))$. For equal means, $V_t^\mu(i,j) = 0$ and the theorem bounds absolute fluctuations (Appendix D.3).

5 Sharp Guarantees in the Stabilized Regime

Starting from stabilization, we lower-bound evidence against every challenger (Proposition 5.1), invert the stopping threshold (Theorem 5.2), and deduce linear count growth (Proposition 5.3). Section 6 then bounds the time needed to reach this regime. For analysis, continue sampling beyond τ_δ, which remains the first threshold crossing. For integers $t \geq K$, let $\alpha_t := N_{t,i^\star}/t$. Fix a compact interval $U \Subset (0,1)$ with β in its interior. A *U-stabilization time* is a random time $t_U \in \{K, K+1, \dots\} \cup \{\infty\}$ such that $B_{t+1} = i^\star$ and $\alpha_t \in U, \qquad \forall t \geq t_U$. Throughout this section, *stabilization* refers to this joint condition.

For an almost surely finite nonnegative random time T, an *upper-quantile bound* is a function $u_T : (0,1] \to [0,\infty)$ satisfying

$$\mathbb{P}_\mu\left(T \leq u_T(\eta)\right) \geq 1-\eta \qquad \text{for every } \eta \in (0,1].$$

We call the bound *summable* if $\sum_{k\geq 0} 2^{-k} u_T(2^{-k}) < \infty$, which implies $\mathbb{E}_\mu[T] < \infty$. Constants in expected bounds may also depend on the assumed upper-quantile bound or sufficient-exploration parameters.

For $t \geq t_U$ and $j \neq i^\star$, define the population penalized index $Q_t(j) := \Phi_j(N_{t,i^\star}, N_{t,j}) = V_t^\mu(i^\star, j) + \log N_{t,j}$. Since $B_{t+1} = i^\star$, $\widehat{\mu}_{t,i^\star} \geq \widehat{\mu}_{t,j}$ for every $j \neq i^\star$ and $W_t(i^\star, j) = V_t(i^\star, j)$. For $\eta \in (0,1)$, use the looser confidence term $\widetilde{H}_t(\eta) := \log(K/\eta) + \log(et)$.

Proposition 5.1 (Uniform lower bound on evidence against all challengers)**.** *Fix a compact interval $U \Subset (0,1)$ containing β in its interior. There exist constants $C_\mu(U), \overline{C}_\mu(U) < \infty$ such that for every $\eta \in (0,1)$ there is an event $\mathcal{E}_{\mathrm{env}}(\eta)$ with $\mathbb{P}_\mu(\mathcal{E}_{\mathrm{env}}(\eta)) \geq 1-2\eta$ on which the following holds simultaneously for every U-stabilization time t_U and every $s \geq t_U$:*

$$\min_{j \neq i^\star} Q_s(j) \geq \Gamma^{\mathrm{disc}}_{N_{s,i^\star},\, s-N_{s,i^\star}}(\mu) - C_\mu(U)\left(t_U + \sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta)\right) - \overline{C}_\mu(U).$$

Consequently, on the same event,

$$\min_{j \neq i^\star} W_s(i^\star, j) \geq \frac{s}{T^\star_{\alpha_s}(\mu)} - C_\mu(U)\left(t_U + \sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta)\right) - \overline{C}_\mu(U).$$

The evidence against every challenger tracks the fixed-fraction information rate, up to stabilization and concentration errors. Appendix E.4 gives the proof.

Theorem 5.2 (Non-asymptotic stopping bound in the stabilized regime)**.** *Fix $\beta \in (0,1)$ and a compact interval $U \Subset (0,1)$ containing β in its interior. Let t_U be an almost surely finite U-stabilization time. Then there exist constants $C_{\mu,k}(U) < \infty$, $k \in \{0,1,2,3\}$, such that, for every*

$\delta \in (0, 1/2]$ *and every* $\eta \in (0, 1/4]$*, with probability at least* $1 - 3\eta$*,*

$$\tau_\delta \le (1 + C_{\mu,0}(U))t_U + T^\star_\beta(\mu)\log(1/\delta) + C_{\mu,1}(U)\sqrt{\log(1/\delta)\left(\log\frac{K}{\eta} + \log\log\frac{e}{\delta}\right)} + C_{\mu,2}(U)\left(\log\frac{K}{\eta} + \log\log\frac{e}{\delta}\right) + C_{\mu,3}(U).$$

Moreover, if t_U *admits a summable upper-quantile bound, then there exist constants* $C'_{\mu,k}(U) < \infty$*,* $k \in \{1, 2, 3\}$*, such that, for all* $\delta \in (0, 1/2]$*,*

$$\mathbb{E}_\mu[\tau_\delta] \le T^\star_\beta(\mu)\log(1/\delta) + C'_{\mu,1}(U)\sqrt{\log(1/\delta)\log\log(e/\delta)} + C'_{\mu,2}(U)\log\log(e/\delta) + C'_{\mu,3}(U).$$

The bound has the exact fixed-β leading coefficient, with corrections for stabilization and concentration errors and bounded remainders (Appendix F.5). A high-probability bound on t_U yields a high-probability stopping bound; a summable upper-quantile bound yields an expected stopping-time bound.

Proposition 5.3 (Linear growth of every suboptimal-arm pull count in the stabilized regime)**.** *Fix a compact interval* $U \Subset (0, 1)$ *containing* β *in its interior and define* $\lambda_{\mu,U} := \inf_{\alpha\in U} T^\star_\alpha(\mu)^{-1} > 0$*. For every* $\eta \in (0, 1)$*, on* $\mathcal{E}_{\mathrm{env}}(\eta)$ *from Proposition 5.1, simultaneously for every* U*-stabilization time* t_U*, every* $s \ge t_U$*, and every* $j \ne i^\star$*,*

$$N_{s,j} \ge \frac{\lambda_{\mu,U}s - C_\mu(U)\left(t_U + \sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta)\right) - \overline{C}_\mu(U) - \log s}{d(\mu_j, \mu_{i^\star})}.$$

Each suboptimal-arm count grows linearly after a transient controlled by t_U and the concentration error; Appendix E.5 gives the proof.

6 Reaching the Stabilized Regime

We bound entry into the regime of Section 5 using armwise confidence bounds and sufficient sampling. Here t_U denotes the earliest U-stabilization time, or ∞ if none exists. To quantify empirical separation, let c_0 be the positive finite universal constant from Lemma D.1 in Appendix D.1. For $\eta \in (0, 1)$, positive integers n, and $j \in [K] \setminus \{i^\star\}$, define

$$r(n, \eta) := \sqrt{\frac{c_0}{2n}\left(\log\frac{8K^2}{\eta} + \log\log(en)\right)},$$

$$n_j(\eta) := \min\left\{n \ge 1 : r(n, \eta) \le \frac{\mu_{i^\star} - \mu_j}{4}\right\}, \qquad n_\star(\eta) := \max_{j\ne i^\star} n_j(\eta).$$

The count $n_j(\eta)$ makes the confidence radius at most one quarter of the gap between arms $i^\star$ and j. The first time all count requirements hold is $t_{\mathrm{sep}}(\eta) := \inf\{t \ge K : N_{t,i^\star} \ge n_\star(\eta) \text{ and } N_{t,j} \ge n_j(\eta)\ \forall j \ne i^\star\}$.

Let $\mathcal{K}^{\mathrm{sep}}_\eta$ denote the simultaneous armwise event

$$|\widehat{\mu}_{t,i} - \mu_i| \le r(N_{t,i}, \eta), \qquad \forall t \ge K,\ i \in [K].$$

Then $\mathbb{P}_\mu(\mathcal{K}^{\mathrm{sep}}_\eta) \ge 1 - \eta$ (Appendix D.1).

Corollary 6.1 (Time until all arm counts support empirical separation)**.** *Fix a compact interval* $U \Subset (0, 1)$ *containing* β *in its interior. There exists* $C_{\mu,\beta}(U) < \infty$ *such that for every* $\eta \in (0, 1/4]$*, on the event of Proposition 5.1,* $t_{\mathrm{sep}}(\eta) \le C_{\mu,\beta}(U)\left(t_U + \log(K/\eta) + \log\log(e/\eta) + 1\right)$.

Linear count growth controls separation by t_U and logarithmic confidence terms (Appendix E.5). We next show why these counts ensure permanent leader correctness, then bound their accumulation before stabilization.

Proposition 6.2 (Sufficient sampling makes the leader permanently correct)**.** *For every $\eta \in (0,1)$, on $\mathcal{K}_\eta^{\mathrm{sep}}$, the empirical leader is permanently correct after $t_{\mathrm{sep}}(\eta)$: $B_{t+1} = i^\star, \forall t \geq t_{\mathrm{sep}}(\eta)$.*

The confidence event excludes later incorrect leaders (Appendix G.1).

Proposition 6.3 (High-probability stabilization bound for the unguarded algorithm)**.** *The unguarded β-EB-TCI algorithm admits finite exponents $a_\mu > 0$, $b_\mu \geq 0$, and a constant $C_{\mu,\beta,1} < \infty$ such that for every $\eta \in (0, 1/4]$, $\mathbb{P}_\mu\left(t_{\mathrm{sep}}(\eta/6) \leq C_{\mu,\beta,1}\eta^{-a_\mu}(\log(e/\eta))^{b_\mu}\right) \geq 1 - \eta/2$. For every compact interval $U \Subset (0,1)$ containing β in its interior, there is a constant $C_{\mu,\beta,2}(U) < \infty$ such that $\mathbb{P}_\mu\left(t_U \leq C_{\mu,\beta,2}(U)\eta^{-a_\mu}(\log(e/\eta))^{b_\mu}\right) \geq 1 - \eta$.*

The earliest t_U is defined on each sample path without reference to η. The bounds for all η therefore concern the same random time. Letting $\eta \downarrow 0$ gives $t_U < \infty$ almost surely (Appendix G.2).

Corollary 6.4 (Non-asymptotic high-probability bound for the unguarded algorithm)**.** *Fix a compact interval $U \Subset (0,1)$ containing β in its interior. Under the standing unique-best-arm assumption, the unguarded algorithm admits constants $C_{\mu,\beta,k}(U) < \infty$, $k \in \{1,2,3\}$, such that for every $\delta \in (0, 1/2]$ and $\eta \in (0, 1/4]$,*

$$\mathbb{P}_\mu\Bigg(\tau_\delta \leq T_\beta^\star(\mu)\log(1/\delta) + C_{\mu,\beta,1}(U)\sqrt{\log(1/\delta)\left(\log(K/\eta) + \log\log(e/\delta)\right)} + C_{\mu,\beta,2}(U)\left(\log(K/\eta) + \log\log(e/\delta)\right) + C_{\mu,\beta,3}(U)\eta^{-a_\mu}(\log(e/\eta))^{b_\mu}\Bigg) \geq 1 - 4\eta.$$

Appendix G.2 derives the corollary from Theorem 5.2. The polynomial-in-$1/\eta$ bound need not be summable, so it gives no expected bound by itself. The next theorem supplies a condition that does.

Theorem 6.5 (Expected stopping time under finite-mean sufficient exploration)**.** *Consider the unguarded algorithm and fix a compact interval $U \Subset (0,1)$ containing β in its interior. Assume* finite-mean sufficient exploration*: there exist a positive finite constant c_{se} and an almost surely finite random time T_{se} with $\mathbb{E}_\mu[T_{\mathrm{se}}] < \infty$ such that*

$$\min_{i \in [K]} N_{t,i} \geq c_{\mathrm{se}}\sqrt{t/K} \qquad \forall t \geq T_{\mathrm{se}}.$$

Then there exists a U-stabilization time t_U that admits a summable upper-quantile bound. Consequently, there are constants $C_{\mu,k}(U) < \infty$, $k \in \{1,2,3\}$, such that, for all $\delta \in (0, 1/2]$,

$$\mathbb{E}_\mu[\tau_\delta] \leq T_\beta^\star(\mu)\log(1/\delta) + C_{\mu,1}(U)\sqrt{\log(1/\delta)\log\log(e/\delta)} + C_{\mu,2}(U)\log\log(e/\delta) + C_{\mu,3}(U).$$

Finite-mean sufficient exploration makes stabilization integrable (Appendix G.3).

Corollary 6.6 (Expected bound for pairwise-distinct means)**.** *Assume pairwise-distinct arm means and let the unguarded Bernoulli β-EB-TCI rule break ties as in Jourdan et al. [9] (e.g., uniformly at random). Then there exist instance-dependent constants $C_{\mu,\beta,1}, C_{\mu,\beta,2}, C_{\mu,\beta,3} < \infty$ such that, for all $\delta \in (0, 1/2]$,*

$$\mathbb{E}_\mu[\tau_\delta] \leq T_\beta^\star(\mu)\log(1/\delta) + C_{\mu,\beta,1}\sqrt{\log(1/\delta)\log\log(e/\delta)} + C_{\mu,\beta,2}\log\log(e/\delta) + C_{\mu,\beta,3}.$$

The cited exploration result supplies the premise of Theorem 6.5 (Appendix G.4).

A simple exploration guard. Pulling a least-sampled arm when its count is below $\sqrt{t/K}$ controls expected stabilization time under a unique best arm.

Proposition 6.7 (Deterministic exploration rate of the guarded algorithm)**.** *Assume the exploration guard in Algorithm 1 is enabled. Then for all $t \geq K$, $\min_{i \in [K]} N_{t,i} \geq \left\lfloor \sqrt{t/K} \right\rfloor - 1$. Consequently, for every $\eta \in (0,1)$, $t_{\mathrm{sep}}(\eta) \leq K\left(n_\star(\eta) + 1\right)^2$. For every compact interval $U \Subset (0,1)$ containing β in its interior, there exists a constant $C_{\mu,\beta}(U) < \infty$ such that the earliest U-stabilization time t_U satisfies, for every $\eta \in (0, 1/4]$,*

$$\mathbb{P}_\mu\left(t_U \leq C_{\mu,\beta}(U)(\log(e/\eta))^2\right) \geq 1 - \eta.$$

This bound is summable (Appendix G.5).

Guarded rounds. Guarded rounds bypass the $\mathrm{Ber}(\beta)$ leader–challenger coin, changing the best-arm sampling fraction by at most G_t^{tot}/t. Proposition F.1 in Appendix F.1 gives concentration; Lemma E.3 in Appendix E.2 gives $G_t^{\mathrm{tot}} \le \sqrt{Kt} + K$.

Theorem 6.8 (Expected stopping time of the guarded algorithm)**.** *For any Bernoulli instance μ with a unique best arm, Algorithm 1 with the exploration guard enabled admits instance-dependent constants $C_{\mu,\beta,1}, C_{\mu,\beta,2}, C_{\mu,\beta,3} < \infty$ such that, for all $\delta \in (0, 1/2]$,*

$$\mathbb{E}_\mu[\tau_\delta] \le T^\star_\beta(\mu)\log(1/\delta) + C_{\mu,\beta,1}\sqrt{\log(1/\delta)\log\log(e/\delta)} + C_{\mu,\beta,2}\log\log(e/\delta) + C_{\mu,\beta,3}.$$

The guard recovers the fixed-β leading term (Appendix G.5).

The unguarded expected bound. For the unguarded algorithm under the default convention, finite-mean sufficient exploration and expected stabilization time remain unresolved when suboptimal means coincide. We examine one proof route.

Fix a provisional leader $b \in [K] \setminus \{i^\star\}$ and a candidate challenger $u \in [K] \setminus \{b\}$. For positive integer counts n, m, define the population penalized index with the leader count held at n by $\Phi_u^{(b)}(n, m) := \inf_{x\in(0,1)} \{nd(\mu_b, x) + md(\mu_u, x)\} + \log m$.

Proposition 6.9 (Limitation of one-step penalized-index comparisons)**.** *Fix distinct arms $b, j, u \in [K]$ such that $\mu_b = \mu_j = \lambda < \mu_u$. Let $C > 0$ and, for $\eta \in (0, e^{-1}]$, define $m_u(\eta) := \lceil C\log(e/\eta)\rceil$. Suppose $M : (0, e^{-1}] \to \{1, 2, \dots\}$ is such that for every sufficiently small η and all sufficiently large n, $\Phi_u^{(b)}(n, m_u(\eta)) \le \Phi_j^{(b)}(n, M(\eta))$. Then for every $a < Cd(\mu_u, \lambda)$, $M(\eta) \ge \eta^{-a}$ for all sufficiently small η.*

With $m_u(\eta) = \lceil C\log(e/\eta)\rceil$ pulls of u, the comparison required for all sufficiently large n requires polynomially many pulls of the equal-mean suboptimal arm j (Appendix G.6). A global, path-dependent analysis may still establish finite-mean sufficient exploration with deterministic tie-breaking.

7 Experiments

We compare empirical mean stopping times with the $T^\star_\beta(\mu)$ predictor and examine sensitivity to the leader-sampling probability β, the guarded-pull fraction, and removal of the logarithmic count penalty. All experiments use Bernoulli rewards, the stopping threshold in (3.1), and smallest-index tie-breaking.

Protocol. We use three representative mean vectors: SEP5 $= [0.70, 0.50, 0.40, 0.30, 0.20]$ has well-separated arms, CLUSTER-CLOSE5 $= [0.70, 0.64, 0.63, 0.62, 0.61]$ has distinct but closely clustered suboptimal arms, and TIE-CLOSE4 $= [0.60, 0.56, 0.56, 0.56]$ has three equal-mean suboptimal arms. The evaluation grid combines these instances with $\beta \in \{0.3, 0.5, 0.7\}$ and $\delta \in \{10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}\}$. For each setting we run both the unguarded and guarded algorithms with 200 independent replications. We compare the empirical mean stopping time with the threshold-crossing predictor

$$t_{\mathrm{pred}}(\mu, \beta, \delta) := \min\{t \ge K : t/T^\star_\beta(\mu) > c(t, \delta)\}.$$

The minimum is over integer times: the predictor is the first round when the leading information term $t/T^\star_\beta(\mu)$ exceeds the stopping threshold. Appendix H gives the full grids, random-seed policy, numerical solver, uncertainty summaries, stopping caps, and additional diagnostics, together with the code and reproduction commands.

Predictor agreement and sensitivity to β. Figure 1 compares empirical mean stopping times with t_{pred} at $\beta = 0.5$. The empirical means track the predictor on both SEP5 and TIE-CLOSE4. Across the 45 guarded settings, the mean and maximum absolute relative errors of the empirical means with respect to t_{pred} are 1.53% and 3.86%; for the unguarded settings, they are 1.66% and 4.20%. None of the 18,000 penalized runs reaches the 700,000-round cap or recommends a wrong arm. Figure 2 also shows that the dependence on β is nonmonotone and instance-specific, as predicted by $T^\star_\beta(\mu)$; the mean gaps alone do not determine the best choice of β.

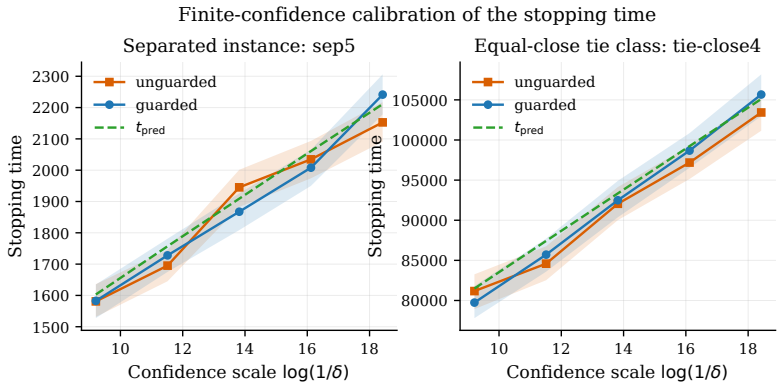
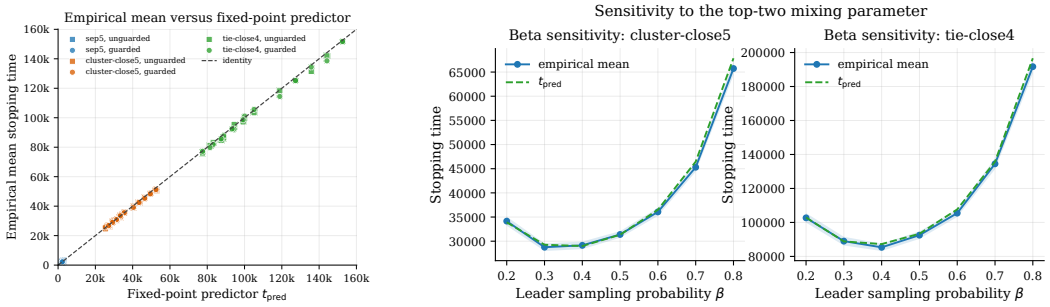


Figure 1: Agreement between observed stopping times and the predictor at $\beta = 0.5$. Shaded bands are bootstrap percentile 95% confidence intervals for the empirical mean over 200 replications per point. The dashed curve is the threshold-crossing predictor t_{pred}.



(a) Empirical mean versus t_{pred}.

(b) Sensitivity to the leader-sampling probability β.

Figure 2: Main experimental diagnostics. Left: empirical mean stopping time versus the threshold-crossing predictor over the complete evaluation grid for the guarded and unguarded algorithms. Right: β-sensitivity on CLUSTER-CLOSE5 and TIE-CLOSE4 with $\delta = 10^{-6}$.

Guard overhead and penalty ablation. On TIE-CLOSE4 with $\beta = 0.5$ and $\delta = 10^{-8}$, using 600 replications per penalized variant, the guard changes the median stopping time from 104,918 to 103,743 and the empirical 95th percentile from 134,140 to 130,792. The largest mean guarded-pull fraction $G^{\text{tot}}_{\tau_\delta}/\tau_\delta$ in the tested equal-mean instances is 0.0069%. Removing the penalty $\log N_{t,j}$ raises the mean recorded time from 105,593 for the unguarded penalized algorithm to 188,036 for the unpenalized ablation (500 replications). In the ablation, 12.8% of runs reach the cap and contribute the cap value to this mean. Appendix H.2 gives an idealized population-allocation reference that tracks a globally optimal allocation computed from the true means. It compares stopping-time scales rather than ranking algorithms and is distinct from empirical Track-and-Stop.

Scope of the comparisons. All variants share the stopping threshold and tie-breaking, so differences reflect the sampling rule alone; the results concern the tested Bernoulli instances and grid and do not claim superiority over other best-arm identification algorithms.

8 Conclusion

For the exact penalized challenger of Bernoulli β-EB-TCI, we proved linear growth of every suboptimal-arm count in the stabilized regime and stopping bounds with exact leading term $T^\star_\beta(\mu)\log(1/\delta)$: with high probability for the unguarded rule, in expectation under finite-mean sufficient exploration, and in expectation for every unique-best instance with the guard. Extensions to other reward distributions, adaptive β, and finite-mean stabilization of the unguarded rule with equal-mean suboptimal arms remain open.

Acknowledgments and Disclosure of Funding

We thank the anonymous reviewers and the Area Chair for their valuable feedback, which helped improve this work. We are also grateful to Marc Jourdan for helpful discussions and comments on the initial direction of this research.

Nam Nguyen is supported by grant No. VUNI.2526.CAIR.04.

The authors declare no competing interests.

References

[1] Jean-Yves Audibert, Sébastien Bubeck, and Rémi Munos. Best arm identification in multi-armed bandits. In Adam Tauman Kalai and Mehryar Mohri, editors, *Proceedings of the Twenty-Third Conference on Learning Theory*, pages 41–53. Omnipress, 2010. ISBN 978-0-9822529-2-5. URL https://www.learningtheory.org/colt2010/papers/59Audibert.pdf.

[2] Kapilan Balagopalan, Ngo Tuan Nguyen, Yao Zhao, and Kwang-Sung Jun. Fixing the loose brake: Exponential-tailed stopping time in best arm identification. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu, editors, *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 2603–2645. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/balagopalan25a.html.

[3] Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7(39):1079–1105, 2006. URL https://jmlr.org/papers/v7/evendar06a.html.

[4] Aurélien Garivier and Emilie Kaufmann. Optimal best arm identification with fixed confidence. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 998–1027, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL https://proceedings.mlr.press/v49/garivier16a.html.

[5] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020. doi: 10.1214/18-PS321. URL https://doi.org/10.1214/18-PS321.

[6] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021. doi: 10.1214/20-AOS1991. URL https://doi.org/10.1214/20-AOS1991.

[7] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil' UCB: An optimal exploration algorithm for multi-armed bandits. In Maria Florina Balcan, Vitaly Feldman, and Csaba Szepesvári, editors, *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pages 423–439, Barcelona, Spain, 13–15 Jun 2014. PMLR. URL https://proceedings.mlr.press/v35/jamieson14.html.

[8] Marc Jourdan and Rémy Degenne. Non-asymptotic analysis of a UCB-based top two algorithm. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *Advances in Neural Information Processing Systems*, volume 36, pages 68980–69020. Curran Associates, Inc., 2023. doi: 10.52202/075280-3019. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/d9b564716709357b4bccec9fc9ad04d2-Abstract-Conference.html.

[9] Marc Jourdan, Rémy Degenne, Dorian Baudry, Rianne de Heide, and Emilie Kaufmann. Top two algorithms revisited. In Sanmi Koyejo, S. Mohamed, Alekh Agarwal, Danielle Belgrave, Kyunghyun Cho, and Alice Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 26791–26803. Curran Associates, Inc., 2022. doi: 10.52202/068431-1943. URL https://proceedings.neurips.cc/paper_files/paper/2022/hash/ab5f5f22e3e09f4424592ffb06840ab0-Abstract-Conference.html.

[10] Shivaram Kalyanakrishnan, Ambuj Tewari, Peter Auer, and Peter Stone. PAC subset selection in stochastic multi-armed bandits. In John Langford and Joelle Pineau, editors, *Proceedings of the Twenty-Ninth International Conference on Machine Learning*, pages 655–662, New York, NY, USA, June–July 2012. Omnipress. ISBN 978-1-4503-1285-1. URL https://www.cs.utexas.edu/~pstone/Papers/bib2html/b2hd-ICML12-shivaram.html.

[11] Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1238–1246, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL https://proceedings.mlr.press/v28/karnin13.html.

[12] Emilie Kaufmann and Shivaram Kalyanakrishnan. Information complexity in bandit subset selection. In Shai Shalev-Shwartz and Ingo Steinwart, editors, *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pages 228–251, Princeton, NJ, USA, 12–14 Jun 2013. PMLR. URL https://proceedings.mlr.press/v30/Kaufmann13.html.

[13] Emilie Kaufmann and Wouter M. Koolen. Mixture martingales revisited with applications to sequential tests and confidence intervals. *Journal of Machine Learning Research*, 22(246):1–44, 2021. URL https://jmlr.org/papers/v22/18-798.html.

[14] 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): 1–42, 2016. URL https://jmlr.org/papers/v17/kaufman16a.html.

[15] Tor Lattimore and Csaba Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.

[16] Shie Mannor and John N. Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research*, 5:623–648, 2004. URL https://www.jmlr.org/papers/v5/mannor04b.html.

[17] Daniel Russo. Simple bayesian algorithms for best arm identification. In Vitaly Feldman, Alexander Rakhlin, and Ohad Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 1417–1418, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR. URL https://proceedings.mlr.press/v49/russo16.html.

[18] Xuedong Shang, Rianne de Heide, Pierre Ménard, Emilie Kaufmann, and Michal Valko. Fixed-confidence guarantees for bayesian best-arm identification. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 1823–1832. PMLR, 26–28 Aug 2020. URL https://proceedings.mlr.press/v108/shang20a.html.

Appendix to "Sharp Non-Asymptotic Analysis of the Penalized Challenger in β-EB-TCI for Bernoulli Bandits"

Table of Contents

A Technical Ingredients from Prior Work **14**

B Error Control and δ-Correctness **15**

B.1 Bernoulli Specialization of the Transportation Cost 15

B.2 Uniform Error Control 16

C Population Allocation Calculus and the Discrete Penalized Benchmark **17**

C.1 Bernoulli Two-Arm Calculus 17

C.2 Boundary Behavior and Inverse Maps 18

C.3 Discrete Rounding and Proof of Proposition 4.1 19

D Time-Uniform Pairwise GLR Control **21**

D.1 Armwise KL Confidence Sequences 21

D.2 Local Lipschitz Control of the Bernoulli GLR 22

D.3 Proof of Theorem 4.2 23

E Evidence Bounds in the Stabilized Regime **24**

E.1 Continuous Equalizers and Slope Bounds 24

E.2 Guard-Count Control 26

E.3 Uniform Challenger-Index Lower Bound 26

E.4 Proof of Proposition 5.1 28

E.5 Challenger Growth and Separation 29

F Leader-Fraction Control and Threshold Inversion **30**

F.1 Uniform Leader-Fraction Concentration 30

F.2 Entrance into a Compact Neighborhood of β 32

F.3 Centered Quantile Integration 32

F.4 Threshold Inversion 34

F.5 Proof of Theorem 5.2 35

G Guarantees Under a Unique Best Arm **37**

G.1 Separation Implies Permanent Leader Correctness 37

G.2 High-Probability Stabilization Bound for the Unguarded Algorithm 38

G.3 Finite-Mean Sufficient Exploration 40

G.4 Pairwise-Distinct Means 41

G.5 Guarded Expected Stopping-Time Bound 42

G.6 Limitation of One-Step Index Comparisons 42

H Experiment Details **43**

H.1 Algorithms and Implementation 43
H.2 Idealized Population-Allocation Reference 46
H.3 Supplementary Figures 46

Main notation.

Symbol	Meaning
K	number of arms, with $K \geq 2$
$\mu = (\mu_1, \ldots, \mu_K)$	vector of Bernoulli means
$i^\star$	unique best arm
δ	target recommendation-error probability
β	probability of sampling the empirical leader on unguarded rounds
$N_{t,i}$	number of pulls of arm i after t total pulls
$\widehat{\mu}_{t,i}$	empirical mean of arm i after t total pulls
B_{t+1}	empirical leader at round $t+1$
C_{t+1}	challenger on unguarded round $t+1$
$V_t(i,j)$	symmetric empirical pairwise Bernoulli GLR
$V_t^\mu(i,j)$	population counterpart of $V_t(i,j)$
$W_t(i,j)$	one-sided transportation cost used by the algorithm
$N_{\mathrm{eff},t}(i,j)$	effective sample size $N_{t,i}N_{t,j}/(N_{t,i}+N_{t,j})$
τ_δ	first threshold-crossing time
$c(t,\delta)$	stopping threshold
α_t	best-arm sampling fraction $N_{t,i^\star}/t$
t_U	time from which $B_{t+1} = i^\star$ and $\alpha_t \in U$ persist
$\Gamma^{\mathrm{disc}}_{n,s}(\mu)$	discrete penalized allocation benchmark
$T^\star_\alpha(\mu)$	characteristic time with best-arm fraction fixed to α
$T^\star(\mu)$	global characteristic time $\min_{\alpha\in(0,1)} T^\star_\alpha(\mu)$

A Technical Ingredients from Prior Work

We collect the results from prior work used in the analysis and explain their roles in our guarantees. Error control, concentration under adaptive sampling, and sufficient exploration serve distinct roles.

Stopping-rule error guarantee. For every $\delta \in (0,1)$, the bounded-reward result of Jourdan et al. [9, Lemma 2], specialized to the Bernoulli transportation cost and threshold (3.1), gives

$$\mathbb{P}_\mu\Big(\tau_\delta < \infty, \widehat{i}_{\tau_\delta} \neq i^\star\Big) \leq \delta,$$

for any adaptive sampling rule. This bounds the probability of stopping with an incorrect recommendation. It does not establish almost-sure termination, a stopping-time upper bound, or a conditional error bound given stopping. Appendix B verifies that our Bernoulli stopping statistic is covered by this result.

Time-uniform Bernoulli concentration. We use the one-dimensional KL confidence-sequence bound of Kaufmann and Koolen [13, Section 3.3.2, Equation (14)]. For every $\eta \in (0,1)$, Lemma D.1 applies this bound to each arm and evaluates it at the adaptive pull counts, yielding an event $\mathcal{K}_\eta$ with $\mathbb{P}_\mu(\mathcal{K}_\eta) \geq 1-\eta$ on which, for a positive finite universal constant c_0,

$$N_{t,i}\, d(\widehat{\mu}_{t,i}, \mu_i) \leq c_0\left(\log\frac{K}{\eta} + \log\log(et)\right) \quad \text{for all } t \geq K \text{ and } i \in [K].$$

This event underlies the empirical-to-population pairwise-GLR comparison. Entry into the stabilized regime, challenger growth, and stopping require further analysis of the algorithm.

Finite-mean sufficient exploration for pairwise-distinct means. Corollary 6.6 uses the finite-mean sufficient-exploration theorem of Jourdan et al. [9] for pairwise-distinct means and the unguarded variant specified in that corollary. The cited theorem is not used in the unconditional high-probability analysis or the guarded expected stopping-time bound.

Results established in this paper. We prove a deterministic approximation of the discrete penalized allocation benchmark. Using the concentration event above, we establish the simultaneous empirical-to-population pairwise-GLR comparison, linear challenger growth in the stabilized regime, and the uniform evidence lower bound. We also prove a high-probability bound on the time to reach that regime and invert the stopping-threshold inequality. The exploration guard gives deterministic sufficient exploration with sublinear pull overhead. The one-step index-comparison limitation concerns only the stated index-domination proof route; it is not an impossibility theorem for the algorithm.

Connection to the main theorems. The deterministic allocation-benchmark analysis and time-uniform concentration yield Theorem 5.2 and the challenger-growth results. The high-probability stabilization bound gives Corollary 6.4 and almost-sure termination of the unguarded algorithm under a unique best arm.

If the stabilization time also admits a summable upper-quantile bound, integration gives the expected stopping-time bound in Theorem 6.5. For pairwise-distinct means, the finite-mean sufficient-exploration theorem of Jourdan et al. [9] verifies this premise for the tie-breaking variant in Corollary 6.6. A high-probability bound on a random time alone need not give a finite expectation.

For the guarded algorithm, Proposition 6.7 verifies finite-mean sufficient exploration directly and leads to Theorem 6.8. Whenever our analysis establishes almost-sure termination, the stopping-rule error guarantee above gives δ-correctness.

B Error Control and δ-Correctness

We derive the Bernoulli transportation cost and apply the bounded-reward stopping-rule guarantee to control recommendation errors. Combined with almost-sure termination, this guarantee yields δ-correctness.

B.1 Bernoulli Specialization of the Transportation Cost

We identify the one-sided transportation cost with the symmetric pairwise GLR when the empirical means are ordered. In the bounded-reward framework of Jourdan et al. [9], let $\mathcal{M}$ be the class of all probability distributions supported on $[0,1]$, $F \in \mathcal{M}$ a reward distribution, and $u \in (0,1)$ a test mean. For a reward X with law G, define

$$K_{\inf}^{+}(F,u) := \inf\left\{\mathrm{KL}(F,G) : G \in \mathcal{M},\ \mathbb{E}_G[X] > u\right\},$$

$$K_{\inf}^{-}(F,u) := \inf\left\{\mathrm{KL}(F,G) : G \in \mathcal{M},\ \mathbb{E}_G[X] < u\right\}.$$

For distinct arms $i, j \in [K]$ and $t \geq K$, the one-sided transportation cost is

$$W_t^{\mathrm{gen}}(i,j) := \inf_{x \in (0,1)} \left\{N_{t,i} K_{\inf}^{-}(F_{t,i}, x) + N_{t,j} K_{\inf}^{+}(F_{t,j}, x)\right\},$$

where $F_{t,i}$ is arm i's empirical reward distribution after t pulls.

In our Bernoulli setting, $F_{t,i} = \mathrm{Ber}(\widehat{\mu}_{t,i})$. The next proposition shows that the information infima over $\mathcal{M}$ reduce to infima over Bernoulli mean parameters. Remark B.2 treats empirical means at 0 or 1 using the boundary convention in Section 3.

Proposition B.1 (Bernoulli specialization of the transportation cost)**.** *For any $p, u \in (0,1)$,*

$$K_{\inf}^{+}(\mathrm{Ber}(p), u) = \begin{cases} 0, & u < p, \\ d(p,u), & u \geq p, \end{cases} \qquad K_{\inf}^{-}(\mathrm{Ber}(p), u) = \begin{cases} d(p,u), & u \leq p, \\ 0, & u > p, \end{cases}.$$

Consequently, for any $p_i, p_j \in (0,1)$ with $p_i \geq p_j$ and any $N_i, N_j \geq 0$,

$$\inf_{x \in (0,1)} \left\{N_i K_{\inf}^{-}(\mathrm{Ber}(p_i), x) + N_j K_{\inf}^{+}(\mathrm{Ber}(p_j), x)\right\} = \inf_{x \in (0,1)} \left\{N_i d(p_i, x) + N_j d(p_j, x)\right\}.$$

Proof. We reduce the information infima to Bernoulli candidates and use KL monotonicity to restrict the ordered-pair minimization to the interval between the two means. The probability kernel $x \mapsto \mathrm{Ber}(x)$ maps any $G \in \mathcal{M}$ to $\mathrm{Ber}(\mathbb{E}_G[X])$ and leaves $\mathrm{Ber}(p)$ unchanged. The data-processing

inequality therefore gives $\mathrm{KL}(\mathrm{Ber}(p), G) \geq d(p, \mathbb{E}_G[X])$. Bernoulli candidates with means in $(0,1)$ attain the resulting infima or approach their boundary values, so

$$K^+_{\inf}(\mathrm{Ber}(p), u) = \inf_{q>u} d(p,q)\,, \qquad K^-_{\inf}(\mathrm{Ber}(p), u) = \inf_{q<u} d(p,q)\,,$$

where both infima are over $q \in (0,1)$. For $q \in (0,1)$,

$$\frac{\partial}{\partial q} d(p,q) = -\frac{p}{q} + \frac{1-p}{1-q} = \frac{q-p}{q(1-q)}.$$

The map $q \mapsto d(p,q)$ is decreasing on $(0,p]$, increasing on $[p,1)$, and attains its unique minimum 0 at $q = p$. The formulas for $K^+_{\inf}$ and $K^-_{\inf}$ follow, with limiting values at the strict constraint boundaries.

For the second claim, let $x \in (0,1)$ and define

$$\Psi(x) := N_i K^-_{\inf}(\mathrm{Ber}(p_i), x) + N_j K^+_{\inf}(\mathrm{Ber}(p_j), x), \qquad p_i \geq p_j.$$

If $x < p_j$, then

$$\Psi(x) = N_i d(p_i, x) \geq N_i d(p_i, p_j) = \Psi(p_j),$$

because $d(p_i, \cdot)$ is decreasing on $(0, p_i]$. If $x > p_i$, then

$$\Psi(x) = N_j d(p_j, x) \geq N_j d(p_j, p_i) = \Psi(p_i),$$

because $d(p_j, \cdot)$ is increasing on $[p_j, 1)$. The infimum is therefore attained on $[p_j, p_i]$, where the one-sided terms give

$$\Psi(x) = N_i d(p_i, x) + N_j d(p_j, x)\,.$$

The two infima agree. □

Remark B.2. For $x \in (0,1)$, $d(0,x) = -\log(1-x)$ and $d(1,x) = -\log x$. The same kernel argument and KL monotonicity give the one-sided formulas at $p = 0$ and $p = 1$. The ordered-pair argument then applies on $[p_j, p_i] \cap (0,1)$; if both means equal the same endpoint, both infima tend to zero as x approaches that endpoint. The cost identity therefore also covers endpoint empirical means. For distinct arms $i, j \in [K]$ and $t \geq K$, whenever $\widehat{\mu}_{t,i} \geq \widehat{\mu}_{t,j}$,

$$W^{\mathrm{gen}}_t(i,j) = \inf_{x \in (0,1)} \{N_{t,i} d(\widehat{\mu}_{t,i}, x) + N_{t,j} d(\widehat{\mu}_{t,j}, x)\} = V_t(i,j).$$

For $p_i, p_j \in [0,1]$ and $N_i, N_j \geq 0$ with $N_i + N_j > 0$, the pooled mean is

$$x^\star = \frac{N_i p_i + N_j p_j}{N_i + N_j}.$$

The derivative of $x \mapsto N_i d(p_i, x) + N_j d(p_j, x)$ on $(0,1)$ has the sign of $(N_i + N_j)x - N_i p_i - N_j p_j$. If $x^\star \in (0,1)$, it is the minimizer; if $x^\star = 0$ or 1, the infimum is approached as x tends to that endpoint and is attained on $[0,1]$ with the boundary convention in Section 3.

B.2 Uniform Error Control

The bounded-reward stopping-rule guarantee controls recommendation errors under either sampling rule; almost-sure termination is established separately.

Theorem B.3 (Uniform error control for the Bernoulli β-EB-TCI stopping rule)**.** *For $\delta \in (0,1)$, with the threshold*

$$c(t,\delta) = \log(1/\delta) + 2\log(1 + t/2) + 2 + \log(K-1),$$

Algorithm 1, with or without the optional exploration guard, satisfies

$$\mathbb{P}_\mu\left(\tau_\delta < \infty, \widehat{i}_{\tau_\delta} \neq i^\star\right) \leq \delta.$$

Proof. We apply Jourdan et al. [9, Lemma 2] with the bounded-reward class $\mathcal{M}$ defined above. By Remark B.2, the cost used by Algorithm 1 agrees with that lemma's transportation cost, including at endpoint empirical means. The stopping rule therefore matches the rule in the lemma, with recommendation

$$\widehat{i}_t \in \arg\max_{i \in [K]} \widehat{\mu}_{t,i}.$$

For this threshold and any adaptive sampling rule, the lemma bounds the probability of stopping with a wrong recommendation by δ.

The optional guard changes how arms are sampled; the statistic $W_t(\widehat{i}_t, j)$, threshold, and recommendation rule remain the same. The time-uniform error guarantee therefore applies to both the guarded and unguarded algorithms. □

Theorem B.3 controls errors on runs that stop; it does not establish almost-sure termination for an arbitrary sampling rule. Termination follows in the regimes of Corollary 6.4, Theorem 6.5, Corollary 6.6, and Theorem 6.8; combining termination with the displayed error bound gives δ-correctness in those regimes.

C Population Allocation Calculus and the Discrete Penalized Benchmark

We characterize the optimal continuous population allocation at a fixed best-arm fraction and establish its regularity as that fraction varies. Rounding this allocation gives the discrete penalized allocation benchmark approximation in Proposition 4.1.

C.1 Bernoulli Two-Arm Calculus

We first compute the pairwise information rate and its derivatives. Fix $\alpha \in (0,1)$, $j \in [K] \setminus \{i^\star\}$, and a scalar weight $\omega > 0$. Define the population pooled mean

$$m_j^\alpha(\omega) := \frac{\alpha\mu_{i^\star} + \omega\mu_j}{\alpha + \omega}.$$

By Remark B.2, the infimum in the definition of $I_j^\alpha(\omega)$ is uniquely attained at $x = m_j^\alpha(\omega)$, and

$$I_j^\alpha(\omega) = \alpha d\big(\mu_{i^\star}, m_j^\alpha(\omega)\big) + \omega d\big(\mu_j, m_j^\alpha(\omega)\big)\,.$$

Proposition C.1 (Bernoulli population-allocation calculus)**.** *The map I_j^α is C^∞ on $(0,\infty)$, strictly increasing, and strictly concave. More precisely,*

$$\frac{d}{d\omega} I_j^\alpha(\omega) = d\big(\mu_j, m_j^\alpha(\omega)\big) > 0,$$

and

$$\frac{d^2}{d\omega^2} I_j^\alpha(\omega) = -\frac{\alpha^2(\mu_{i^\star} - \mu_j)^2}{(\alpha+\omega)^3 m_j^\alpha(\omega)(1 - m_j^\alpha(\omega))} < 0.$$

Proof. We differentiate the two-arm objective at its pooled minimizer. For $x \in (0,1)$, let

$$f_\omega(x) := \alpha d(\mu_{i^\star}, x) + \omega d(\mu_j, x)\,.$$

For $p, x \in (0,1)$,

$$\frac{\partial}{\partial x} d(p,x) = -\frac{p}{x} + \frac{1-p}{1-x} = \frac{x-p}{x(1-x)}.$$

Applying this identity gives

$$f'_\omega(x) = \frac{(\alpha+\omega)x - (\alpha\mu_{i^\star} + \omega\mu_j)}{x(1-x)}.$$

Hence $f'_\omega(x) = 0$ if and only if $x = m_j^\alpha(\omega)$. Also,

$$f''_\omega(x) = \frac{\alpha\mu_{i^\star} + \omega\mu_j}{x^2} + \frac{\alpha(1-\mu_{i^\star}) + \omega(1-\mu_j)}{(1-x)^2} > 0,$$

so f_ω is strictly convex and the minimizer is unique.

Because $\omega \mapsto m_j^\alpha(\omega)$ is smooth, so is I_j^α. The envelope theorem gives

$$\frac{d}{d\omega} I_j^\alpha(\omega) = \frac{\partial}{\partial\omega} f_\omega(m_j^\alpha(\omega)) = d\big(\mu_j, m_j^\alpha(\omega)\big)\,.$$

Since $\mu_j < \mu_{i^\star}$, $m_j^\alpha(\omega) \in (\mu_j, \mu_{i^\star})$, and therefore $d\big(\mu_j, m_j^\alpha(\omega)\big) > 0$.

For $p, q \in (0,1)$,

$$\frac{d}{dq} d(p,q) = \frac{q-p}{q(1-q)}, \qquad \frac{d}{d\omega} m_j^\alpha(\omega) = \frac{\alpha(\mu_j - \mu_{i^\star})}{(\alpha+\omega)^2}.$$

Differentiating once more yields

$$\frac{d^2}{d\omega^2} I_j^\alpha(\omega) = \frac{m_j^\alpha(\omega) - \mu_j}{m_j^\alpha(\omega)(1 - m_j^\alpha(\omega))} \cdot \frac{\alpha(\mu_j - \mu_{i^\star})}{(\alpha+\omega)^2}.$$

Using

$$m_j^\alpha(\omega) - \mu_j = \frac{\alpha(\mu_{i^\star} - \mu_j)}{\alpha + \omega},$$

we obtain the stated second derivative, which is strictly negative. □

C.2 Boundary Behavior and Inverse Maps

The endpoint limits determine the ranges of the pairwise information rates and hence the domains of their inverse maps. These inverses characterize the unique fixed-α optimizer and its regularity.

Lemma C.2 (Boundary behavior)**.** *Fix $\alpha \in (0,1)$ and $j \in [K] \setminus \{i^\star\}$. Then*

$$\lim_{\omega \downarrow 0} I_j^\alpha(\omega) = 0, \qquad \lim_{\omega \uparrow \infty} I_j^\alpha(\omega) = \alpha d(\mu_{i^\star}, \mu_j)\,.$$

Consequently, I_j^α is a C^∞ strictly increasing bijection from $(0,\infty)$ onto $(0, \alpha d(\mu_{i^\star}, \mu_j))$.

Proof. We evaluate the two terms in the information rate as the competitor weight tends to zero or infinity. By Proposition C.1,

$$I_j^\alpha(\omega) = \alpha d\big(\mu_{i^\star}, m_j^\alpha(\omega)\big) + \omega d\big(\mu_j, m_j^\alpha(\omega)\big)\,.$$

As $\omega \downarrow 0$, $m_j^\alpha(\omega) \to \mu_{i^\star}$, so both terms vanish. As $\omega \uparrow \infty$, $m_j^\alpha(\omega) \to \mu_j$, and hence

$$\alpha d\big(\mu_{i^\star}, m_j^\alpha(\omega)\big) \to \alpha d(\mu_{i^\star}, \mu_j)\,.$$

Moreover,

$$m_j^\alpha(\omega) - \mu_j = \frac{\alpha(\mu_{i^\star} - \mu_j)}{\alpha + \omega} = O(\omega^{-1}),$$

and, for fixed $p \in (0,1)$, $d(p,q) = O((q-p)^2)$ as $q \to p$. Therefore $\omega d\big(\mu_j, m_j^\alpha(\omega)\big) \to 0$. The bijection claim follows from Proposition C.1 and the two endpoint limits. □

Proposition C.3 (Unique fixed-α optimizer)**.** *For every $\alpha \in (0,1)$ there exist a unique vector $w^\alpha(\mu) \in \Delta_{1-\alpha}$ and a unique scalar $\lambda_\alpha > 0$ such that*

$$I_j^\alpha(w_j^\alpha(\mu)) = \lambda_\alpha \qquad \forall j \neq i^\star.$$

Moreover, $\lambda_\alpha = T_\alpha^\star(\mu)^{-1}$. On compact subsets of $(0,1)$, the following map is C^1:

$$\alpha \mapsto \big(w^\alpha(\mu), T_\alpha^\star(\mu)^{-1}\big)\,.$$

Proof. We solve for the common information rate using the inverse maps, then apply the implicit function theorem to obtain regularity. Fix $\alpha \in (0,1)$ and define

$$D_j := d(\mu_{i^\star}, \mu_j) > 0, \qquad j \neq i^\star.$$

By Lemma C.2, for each fixed α the map $I_j^\alpha : (0,\infty) \to (0, \alpha D_j)$ is a strictly increasing C^∞ bijection. Define

$$G_j(\alpha, \lambda) := (I_j^\alpha)^{-1}(\lambda), \qquad 0 < \lambda < \alpha D_j.$$

To impose the total allocation constraint, define

$$D_{\min} := \min_{j \neq i^\star} D_j, \qquad \varphi(\alpha, \lambda) := \sum_{j \neq i^\star} G_j(\alpha, \lambda) - (1-\alpha),$$

for (α, λ) in the domain

$$\{(\alpha, \lambda) : \alpha \in (0,1),\ 0 < \lambda < \alpha D_{\min}\}.$$

For fixed α, the map $\lambda \mapsto \varphi(\alpha, \lambda)$ is continuous and strictly increasing. In addition,

$$\lim_{\lambda \downarrow 0} \varphi(\alpha, \lambda) = -(1-\alpha) < 0,$$

whereas

$$\lim_{\lambda \uparrow \alpha D_{\min}} \varphi(\alpha, \lambda) = +\infty,$$

because the inverse for an arm attaining the minimum directed KL diverges as $\lambda \uparrow \alpha D_{\min}$. Therefore there exists a unique $\lambda_\alpha \in (0, \alpha D_{\min})$ satisfying $\varphi(\alpha, \lambda_\alpha) = 0$. Set $w_j^\alpha(\mu) := G_j(\alpha, \lambda_\alpha)$. Then $w^\alpha(\mu) \in \Delta_{1-\alpha}$ and $I_j^\alpha(w_j^\alpha(\mu)) = \lambda_\alpha$ for every $j \neq i^\star$.

To prove optimality, take a vector $w \in \Delta_{1-\alpha}$ and write

$$m(w) := \min_{j \neq i^\star} I_j^\alpha(w_j).$$

Since each I_j^α is increasing,

$$w_j \geq G_j(\alpha, m(w)) \qquad \forall j \neq i^\star.$$

Summing over $j \neq i^\star$ gives

$$1 - \alpha = \sum_{j \neq i^\star} w_j \geq \sum_{j \neq i^\star} G_j(\alpha, m(w)).$$

By strict monotonicity of the sum of inverse maps and the definition of λ_α, this implies $m(w) \leq \lambda_\alpha$. Hence

$$\max_{w \in \Delta_{1-\alpha}} \min_{j \neq i^\star} I_j^\alpha(w_j) = \lambda_\alpha,$$

and therefore $\lambda_\alpha = T_\alpha^\star(\mu)^{-1}$. If w is any maximizer, then $m(w) = \lambda_\alpha$, so $w_j \geq G_j(\alpha, \lambda_\alpha) = w_j^\alpha(\mu)$ for every $j \neq i^\star$. Both vectors sum to $1-\alpha$, so equality holds in every coordinate, proving uniqueness.

For regularity, define

$$\Psi_j(\alpha, \omega, \lambda) := I_j^\alpha(\omega) - \lambda.$$

Since Ψ_j is C^1 and $\partial_\omega \Psi_j(\alpha, \omega, \lambda) = \partial_\omega I_j^\alpha(\omega) > 0$, the inverse function theorem implies that G_j is C^1 on its domain. The residual φ is therefore C^1, with

$$\partial_\lambda \varphi(\alpha, \lambda) = \sum_{j \neq i^\star} \partial_\lambda G_j(\alpha, \lambda) > 0.$$

Applying the implicit function theorem to $\varphi(\alpha, \lambda_\alpha) = 0$ shows that $\alpha \mapsto \lambda_\alpha$ is C^1 on compact subsets of $(0,1)$. Each coordinate $\alpha \mapsto w_j^\alpha(\mu) = G_j(\alpha, \lambda_\alpha)$ is therefore C^1 as well. □

Lemma C.4 (Lipschitz control of $T_\alpha^\star(\mu)^{-1}$ on compacts)**.** *Let $U \Subset (0,1)$ be compact. Then there exists a finite constant $L_\mu(U)$ such that for all $\alpha, \alpha' \in U$,*

$$\left| \frac{1}{T_\alpha^\star(\mu)} - \frac{1}{T_{\alpha'}^\star(\mu)} \right| \leq L_\mu(U) |\alpha - \alpha'|.$$

Proof. By Proposition C.3, the map $\alpha \mapsto T_\alpha^\star(\mu)^{-1}$ is C^1 on $(0,1)$. Its derivative is bounded on the compact convex hull of U, which lies in $(0,1)$. The interval between any two points of U lies in this hull, so the mean value theorem yields the displayed Lipschitz bound. □

C.3 Discrete Rounding and Proof of Proposition 4.1

The rounding lemma preserves the total count while changing each coordinate by at most one. Applied to the continuous optimizer, it bounds the cost of integer counts in the discrete penalized allocation benchmark.

Lemma C.5 (Integer rounding with a fixed sum)**.** *Let $a_1, \dots, a_{K-1} \geq 1$ be real numbers satisfying $\sum_{j=1}^{K-1} a_j = s \in \mathbb{N}$. Then there exist integers $m_1, \dots, m_{K-1} \geq 1$ such that*

$$\sum_{j=1}^{K-1} m_j = s, \qquad |m_j - a_j| \leq 1 \quad \forall j.$$

Proof. Let $\widetilde{m}_j := \lfloor a_j \rfloor$, so $\widetilde{m}_j \geq 1$ for all j, and define $r := s - \sum_{j=1}^{K-1} \widetilde{m}_j$. Because $\sum_j a_j = s$ and each fractional part satisfies $a_j - \lfloor a_j \rfloor < 1$, the remainder satisfies $r \in \{0, 1, \dots, K-2\}$. Add 1 to the r coordinates with largest fractional parts. The resulting integers satisfy the required properties. □

Proof of Proposition 4.1. We obtain an upper bound from the continuous optimum and a lower bound by rounding its optimizer and bounding the logarithmic count penalty. Fix a compact interval $U \Subset (0,1)$ and integers $n \geq 1$ and $s \geq K-1$ with $n/(n+s) \in U$. Set

$$\alpha := \frac{n}{n+s} \in U, \qquad t := n + s.$$

We show that there exists $C_\mu(U) < \infty$, independent of n and s, such that

$$\left| \Gamma^{\mathrm{disc}}_{n,s}(\mu) - \left(\frac{t}{T^\star_\alpha(\mu)} + \log t \right) \right| \leq C_\mu(U).$$

Let $m = (m_j)_{j \neq i^\star}$ be any feasible integer allocation and define $w_j := m_j/t$. Then $w \in \Delta_{1-\alpha}$ and

$$\Phi_j(n, m_j) = t I^\alpha_j(w_j) + \log t + \log w_j.$$

Since $0 < w_j < 1$, the residual penalty satisfies $\log w_j \leq 0$, so

$$\min_{j \neq i^\star} \Phi_j(n, m_j) \leq \log t + t \min_{j \neq i^\star} I^\alpha_j(w_j) \leq \log t + \frac{t}{T^\star_\alpha(\mu)}.$$

Maximizing over feasible m yields

$$\Gamma^{\mathrm{disc}}_{n,s}(\mu) \leq \frac{t}{T^\star_\alpha(\mu)} + \log t.$$

For the lower bound, let $w^\alpha(\mu)$ be the optimizer from Proposition C.3. Since $\alpha \mapsto w^\alpha(\mu)$ is continuous on U and all coordinates are positive,

$$c_U := \inf_{\alpha \in U} \min_{j \neq i^\star} w^\alpha_j(\mu) > 0.$$

Proposition C.1 also gives

$$L_U := \sup_{\alpha \in U} \sup_{j \neq i^\star} \sup_{\omega \in [c_U/2, 1]} \left| \frac{d}{d\omega} I^\alpha_j(\omega) \right| < \infty.$$

Assume first that $t \geq 2/c_U$. Then $a_j := t w^\alpha_j(\mu) \geq 2$, so Lemma C.5 gives integers $m_j \geq 1$ such that

$$\sum_{j \neq i^\star} m_j = s, \qquad |m_j - t w^\alpha_j(\mu)| \leq 1.$$

Hence

$$\left| \frac{m_j}{t} - w^\alpha_j(\mu) \right| \leq \frac{1}{t}, \qquad \frac{m_j}{t} \geq \frac{c_U}{2}.$$

By the mean value theorem,

$$\left| I^\alpha_j\left(\frac{m_j}{t}\right) - I^\alpha_j(w^\alpha_j(\mu)) \right| \leq \frac{L_U}{t}.$$

Since $I^\alpha_j(w^\alpha_j(\mu)) = T^\star_\alpha(\mu)^{-1}$,

$$t I^\alpha_j\left(\frac{m_j}{t}\right) \geq \frac{t}{T^\star_\alpha(\mu)} - L_U.$$

Moreover, $\log(m_j/t) \geq \log(c_U/2)$. Therefore

$$\Phi_j(n, m_j) \geq \frac{t}{T^\star_\alpha(\mu)} + \log t - L_U + \log(c_U/2) \qquad \forall j \neq i^\star.$$

Taking the minimum over j gives

$$\Gamma^{\mathrm{disc}}_{n,s}(\mu) \geq \frac{t}{T^\star_\alpha(\mu)} + \log t - C_{\mu,1}(U),$$

where $C_{\mu,1}(U) := L_U - \log(c_U/2)$.

If $t < 2/c_U$, only finitely many admissible triples (n, s, t) satisfy $n + s = t$ and $n/(n + s) \in U$, so their approximation errors can be absorbed into a finite constant $C_{\mu,0}(U)$. Taking $C_\mu(U) := \max\{C_{\mu,0}(U), C_{\mu,1}(U)\}$ completes the proof. □

D Time-Uniform Pairwise GLR Control

We prove Theorem 4.2 by combining armwise KL confidence sequences with a local Lipschitz bound for the Bernoulli GLR.

D.1 Armwise KL Confidence Sequences

We first control each arm's empirical mean uniformly over its pull count. These bounds then hold at the adaptive counts used by the algorithm.

Lemma D.1 (Armwise KL confidence sequences)**.** *There exists a positive finite universal constant c_0 such that for every $\eta \in (0, 1)$ there is an event $\mathcal{K}_\eta$ with $\mathbb{P}_\mu(\mathcal{K}_\eta) \geq 1 - \eta$ and, on $\mathcal{K}_\eta$, simultaneously for all $t \geq K$ and all arms $i \in [K]$,*

$$N_{t,i} d(\widehat{\mu}_{t,i}, \mu_i) \leq c_0 \left(\log \frac{K}{\eta} + \log\log(et) \right).$$

Consequently,

$$|\widehat{\mu}_{t,i} - \mu_i| \leq \sqrt{\frac{c_0 H_t(\eta)}{2N_{t,i}}}.$$

Proof. We combine the pull-indexed confidence sequences by a union bound, then evaluate them at the adaptive counts.

For each arm $i \in [K]$, let $(Y_{i,s})_{s \geq 1}$ be the reward sequence of arm i, indexed by pull number. These rewards are i.i.d. with law $\mathrm{Ber}(\mu_i)$. For each positive integer n, define

$$\overline{Y}_{i,n} := \frac{1}{n} \sum_{s=1}^{n} Y_{i,s}.$$

By construction,

$$\widehat{\mu}_{t,i} = \overline{Y}_{i,N_{t,i}} \qquad \text{for all } t \geq K.$$

Since the algorithm pulls each arm once initially, $N_{t,i} \geq 1$.

Apply the one-dimensional time-uniform KL confidence-sequence bound for one-parameter exponential families from Kaufmann and Koolen [13, Section 3.3.2, Equation (14)], specialized to Bernoulli distributions. For $\gamma \in (0, 1/2]$, the confidence-dependent term in that bound is at most a universal multiple of $\log(1/\gamma)$. There is therefore a positive finite universal constant c_0 such that, for every $\gamma \in (0, 1/2]$, there exists an event $\mathcal{E}_i(\gamma)$ with $\mathbb{P}_\mu(\mathcal{E}_i(\gamma)) \geq 1 - \gamma$. On $\mathcal{E}_i(\gamma)$, simultaneously for all $n \geq 1$,

$$n d\big(\overline{Y}_{i,n}, \mu_i\big) \leq c_0 \left(\log \frac{1}{\gamma} + \log\log(en) \right).$$

Choose $\gamma = \eta/K < 1/2$, using $K \geq 2$, and define $\mathcal{K}_\eta := \bigcap_{i=1}^{K} \mathcal{E}_i(\eta/K)$. A union bound gives $\mathbb{P}_\mu(\mathcal{K}_\eta) \geq 1 - \eta$.

Because each arm's bound holds for every sample size $n \geq 1$, it also holds at the adaptive count $n = N_{t,i}$. On $\mathcal{K}_\eta$, simultaneously for all $t \geq K$ and all $i \in [K]$,

$$N_{t,i} d(\widehat{\mu}_{t,i}, \mu_i) \leq c_0 \left(\log \frac{K}{\eta} + \log\log(eN_{t,i}) \right) \leq c_0 \left(\log \frac{K}{\eta} + \log\log(et) \right).$$

Pinsker's inequality, $d(p,q) \geq 2(p-q)^2$, yields the displayed radius bound. □

The pull-indexed construction of $\mathcal{K}_\eta$ above gives $\mathcal{K}_\eta \subseteq \mathcal{K}_\eta^{\mathrm{sep}}$, the event defined in Section 6. Indeed, Pinsker's inequality at $n = N_{t,i}$ bounds the estimation error by $\sqrt{c_0(\log(K/\eta) + \log\log(eN_{t,i}))/(2N_{t,i})} \leq r(N_{t,i}, \eta)$. Hence $\mathbb{P}_\mu(\mathcal{K}_\eta^{\mathrm{sep}}) \geq 1 - \eta$.

D.2 Local Lipschitz Control of the Bernoulli GLR

We next use an entropy representation of the two-arm GLR to bound its derivatives on interior mean intervals. The resulting Lipschitz bound scales with the effective sample size.

Lemma D.2 (Bregman representation and local Lipschitz control)**.** *Fix $m, \ell \geq 1$ and $p, q \in (0,1)$, and define*

$$V_{m,\ell}(p,q) := \inf_{x \in (0,1)} \{ m d(p,x) + \ell d(q,x) \}.$$

For the local weight $\alpha := m/(m+\ell)$, let

$$h(x) := x \log x + (1-x)\log(1-x) \qquad \text{for } x \in (0,1),$$

and extend h continuously to $[0,1]$. Then

$$V_{m,\ell}(p,q) = (m+\ell)\left[\alpha h(p) + (1-\alpha)h(q) - h(\alpha p + (1-\alpha)q)\right].$$

Moreover, for every $\varepsilon \in (0, 1/2)$ and all $p, q, u, v \in [\varepsilon, 1-\varepsilon]$,

$$|V_{m,\ell}(p,q) - V_{m,\ell}(u,v)| \leq \frac{2}{\varepsilon(1-\varepsilon)} \frac{m\ell}{m+\ell} \left(|p-u| + |q-v| \right).$$

Proof. We substitute the pooled minimizer into the GLR and bound its derivatives on the interior interval.

The minimizer in the definition of $V_{m,\ell}(p,q)$ is $r := \alpha p + (1-\alpha)q$. Substitution gives the Bregman representation.

For the Lipschitz bound, write

$$V_{m,\ell}(p,q) = (m+\ell)\left[\alpha h(p) + (1-\alpha)h(q) - h(r)\right], \qquad r = \alpha p + (1-\alpha)q.$$

Then

$$\partial_p V_{m,\ell}(p,q) = m(h'(p) - h'(r)), \qquad \partial_q V_{m,\ell}(p,q) = \ell(h'(q) - h'(r)).$$

Since $h''(x) = 1/[x(1-x)]$, for $x, y \in [\varepsilon, 1-\varepsilon]$,

$$|h'(x) - h'(y)| \leq \frac{|x-y|}{\varepsilon(1-\varepsilon)}.$$

Moreover,

$$|p - r| = (1-\alpha)|p-q| \leq 1-\alpha, \qquad |q-r| = \alpha|p-q| \leq \alpha.$$

It follows that

$$|\partial_p V_{m,\ell}(p,q)| \leq \frac{m(1-\alpha)}{\varepsilon(1-\varepsilon)} = \frac{m\ell}{m+\ell} \frac{1}{\varepsilon(1-\varepsilon)},$$

and the same bound holds for $|\partial_q V_{m,\ell}(p,q)|$. Applying the mean value theorem first in the p coordinate and then in the q coordinate keeps both segments inside the interior square and proves the stated inequality, with the factor 2 as slack. □

D.3 Proof of Theorem 4.2

We combine the armwise confidence radii with local Lipschitz control to prove the uniform comparison between the empirical GLR and its population counterpart.

Proof of Theorem 4.2. We bound both GLR statistics directly when one count is small and apply local Lipschitz control when both counts are large.

Fix $\eta \in (0,1)$. Work on the armwise confidence event $\mathcal{K}_\eta$ from Lemma D.1. This event has probability at least $1-\eta$. Fix $t \geq K$ and distinct arms $i, j \in [K]$, and set

$$m := N_{t,i}, \qquad \ell := N_{t,j}, \qquad H := H_t(\eta), \qquad N_{\text{eff}} := \frac{m\ell}{m+\ell}.$$

Throughout this proof, C_μ denotes a positive finite constant depending only on μ; its value may increase between inequalities. On $\mathcal{K}_\eta$,

$$|\widehat{\mu}_{t,i} - \mu_i| \leq \sqrt{\frac{c_0 H}{2m}}, \qquad |\widehat{\mu}_{t,j} - \mu_j| \leq \sqrt{\frac{c_0 H}{2\ell}}.$$

Let

$$\varepsilon_\mu := \frac{1}{2} \min_{a \in [K]} \{\mu_a, 1-\mu_a\}.$$

First suppose that

$$\min\{m, \ell\} \leq \frac{2c_0}{\varepsilon_\mu^2} H.$$

Assume without loss of generality that $m \leq \ell$. Evaluating the empirical infimum defining $V_t(i,j)$ at $x = \mu_j$ gives

$$V_t(i,j) \leq m d(\widehat{\mu}_{t,i}, \mu_j) + \ell d(\widehat{\mu}_{t,j}, \mu_j)\,.$$

The first term is at most $C_\mu m$ because $p \mapsto d(p, \mu_j)$ is bounded on $[0,1]$, and the second term is at most $c_0 H$ on $\mathcal{K}_\eta$. The small-count assumption therefore gives $V_t(i,j) \leq C_\mu H$. Evaluating the population infimum at $x = \mu_j$ gives

$$V_t^\mu(i,j) \leq m d(\mu_i, \mu_j) \leq C_\mu m \leq C_\mu H.$$

The two bounds imply

$$|V_t(i,j) - V_t^\mu(i,j)| \leq C_\mu H,$$

in the small-count regime.

Now suppose that

$$\min\{m, \ell\} > \frac{2c_0}{\varepsilon_\mu^2} H.$$

Then both empirical means belong to $[\varepsilon_\mu/2, 1-\varepsilon_\mu/2]$, and all true means belong to the same interval. Lemma D.2, applied with $\varepsilon = \varepsilon_\mu/2$, gives

$$|V_t(i,j) - V_t^\mu(i,j)| \leq C_\mu N_{\text{eff}} \left(|\widehat{\mu}_{t,i} - \mu_i| + |\widehat{\mu}_{t,j} - \mu_j|\right).$$

Using the armwise confidence radii,

$$|V_t(i,j) - V_t^\mu(i,j)| \leq C_\mu N_{\text{eff}} \left(\sqrt{\frac{H}{m}} + \sqrt{\frac{H}{\ell}}\right).$$

Since $N_{\text{eff}} \leq m$ and $N_{\text{eff}} \leq \ell$,

$$\frac{N_{\text{eff}}}{\sqrt{m}} \leq \sqrt{N_{\text{eff}}}, \qquad \frac{N_{\text{eff}}}{\sqrt{\ell}} \leq \sqrt{N_{\text{eff}}}.$$

We obtain

$$|V_t(i,j) - V_t^\mu(i,j)| \leq C_\mu \sqrt{N_{\text{eff}} H},$$

in the large-count regime. The two regimes together give the bound in Theorem 4.2. Since the armwise bounds hold on the same event $\mathcal{K}_\eta$ for all arms and all pull counts, the conclusion holds simultaneously over all $t \geq K$ and all pairs $i \neq j$. Taking $\mathcal{G}_\eta := \mathcal{K}_\eta$ completes the proof. □

E Evidence Bounds in the Stabilized Regime

We compare the population challenger indices with the discrete penalized allocation benchmark to derive the evidence and count-growth bounds, then bound the separation time. Algorithmic time indices are integers on the continued sampling process of Section 5. Generic constants may change between inequalities, with their dependencies indicated by subscripts and arguments.

For $t \geq K$, let $G_{t+1} \in \{0, 1\}$ indicate whether the exploration guard is used at round $t + 1$; for the unguarded algorithm, $G_{t+1} \equiv 0$. We write

$$\mathcal{J} := [K] \setminus \{i^\star\}, \qquad d_K := |\mathcal{J}| = K - 1.$$

For each positive integer n and $j \in \mathcal{J}$, extend the population penalized index to real $m > 0$ by

$$\Phi_j(n, m) := \inf_{x \in (0,1)} \{nd(\mu_{i^\star}, x) + md(\mu_j, x)\} + \log m.$$

For integer budgets $S \geq d_K$, the discrete benchmark is

$$\Gamma^{\mathrm{disc}}_{n,S}(\mu) = \max_{m_j \in \{1,2,\dots\} : \sum_{j \in \mathcal{J}} m_j = S} \min_{j \in \mathcal{J}} \Phi_j(n, m_j).$$

E.1 Continuous Equalizers and Slope Bounds

We first bound the slopes of the population indices and construct a real allocation at which all indices agree. Rounding this continuous equalizer relates its value to the discrete benchmark.

Lemma E.1 (Uniform slope bounds for Φ_j)**.** *Fix a compact interval $U \Subset (0, 1)$. There exist constants $0 < c_\mu(U) \leq C_\mu(U) < \infty$ such that, whenever $\alpha = n/(n + S) \in U$, $1 \leq m \leq S + 1$, and $j \in \mathcal{J}$,*

$$c_\mu(U) \leq \partial_m \Phi_j(n, m) \leq C_\mu(U).$$

Proof. We differentiate the index and use the compact interval of allocation fractions to keep its pooled mean separated from the suboptimal-arm mean. By the envelope theorem,

$$\partial_m \Phi_j(n, m) = d(\mu_j, x_j(n, m)) + \frac{1}{m}, \qquad x_j(n, m) := \frac{n\mu_{i^\star} + m\mu_j}{n + m}.$$

If $a_U := \inf U$, then for $1 \leq m \leq S + 1$,

$$\frac{n}{n + m} \geq \frac{n}{n + S + 1} = \frac{\alpha(n + S)}{n + S + 1} \geq \frac{a_U}{2},$$

since $n + S \geq 1$. Hence $x_j(n, m)$ stays uniformly separated from μ_j. Since $\mu_{i^\star} > \mu_j$, the term $d(\mu_j, x_j(n, m))$ has a positive lower bound depending only on (μ, U). The upper bound follows from $x_j(n, m) \in [\mu_j, \mu_{i^\star}]$ and $1/m \leq 1$. □

Lemma E.2 (Continuous equalizer and small-budget cutoff)**.** *Fix a compact interval $U \Subset (0, 1)$. For each integer $n \geq 1$, define*

$$\gamma_{\min}(n) := \max_{k \in \mathcal{J}} \Phi_k(n, 1), \qquad S_0(n, \mu) := \sum_{j \in \mathcal{J}} \Phi_j^{-1}(n, \gamma_{\min}(n)),$$

where $\Phi_j^{-1}(n, \gamma)$ is the inverse of $m \mapsto \Phi_j(n, m)$ on $[1, \infty)$. There exists $\overline{S}_\mu < \infty$ such that $S_0(n, \mu) \leq \overline{S}_\mu$ for all $n \geq 1$. For every real budget $S \geq S_0(n, \mu)$, there is a unique pair $(\gamma_{n,S}, r_{n,S})$, $r_{n,S} = (r_{j,n,S})_{j \in \mathcal{J}}$, satisfying

$$\sum_{j \in \mathcal{J}} r_{j,n,S} = S, \qquad r_{j,n,S} \geq 1, \qquad \Phi_j(n, r_{j,n,S}) = \gamma_{n,S} \quad \forall j \in \mathcal{J}.$$

Moreover, if $t = n + S$, $\alpha = n/t \in U$, and $S \geq S_0(n, \mu)$, then

$$\gamma_{n,S} = \frac{t}{T^\star_\alpha(\mu)} + \log t + O_{\mu,U}(1), \qquad r_{j,n,S} = tw^\alpha_j(\mu) + O_{\mu,U}(1),$$

uniformly over $\alpha \in U$, S, and $j \in \mathcal{J}$. If S is an integer, then

$$|\gamma_{n,S} - \Gamma^{\mathrm{disc}}_{n,S}(\mu)| \leq C_\mu(U).$$

Proof. We invert the increasing indices to construct the equalizer, compare it with the fixed-fraction optimizer, and round its coordinates for integer budgets. For fixed n, the map $m \mapsto \Phi_j(n,m)$ is continuous on $(0,\infty)$ and strictly increasing. Indeed,

$$\partial_m \Phi_j(n,m) = d(\mu_j, x_j(n,m)) + \frac{1}{m} > 0, \qquad x_j(n,m) = \frac{n\mu_{i^\star} + m\mu_j}{n+m}.$$

Since $\log m \to +\infty$, we have $\Phi_j(n,m) \to +\infty$ as $m \to \infty$. The inverse on $[1,\infty)$ is therefore well defined for every $\gamma \geq \Phi_j(n,1)$.

Let

$$D^{\max}_{\mu,\star} := \max_{k \in \mathcal{J}} d(\mu_k, \mu_{i^\star}).$$

Evaluating $\Phi_k(n,1)$ at $x = \mu_{i^\star}$ gives $\gamma_{\min}(n) \leq D^{\max}_{\mu,\star}$. Since $\Phi_j(n,m) \geq \log m$, any $m \geq \exp(D^{\max}_{\mu,\star})$ satisfies $\Phi_j(n,m) \geq \gamma_{\min}(n)$. Hence

$$\Phi_j^{-1}(n, \gamma_{\min}(n)) \leq \exp(D^{\max}_{\mu,\star}), \qquad S_0(n,\mu) \leq \overline{S}_\mu := d_K \exp(D^{\max}_{\mu,\star}).$$

For $S \geq S_0(n,\mu)$, define

$$F_n(\gamma) := \sum_{j \in \mathcal{J}} \Phi_j^{-1}(n,\gamma), \qquad \gamma \geq \gamma_{\min}(n).$$

Then F_n is continuous and strictly increasing, $F_n(\gamma_{\min}(n)) = S_0(n,\mu)$, and $F_n(\gamma) \to \infty$ as $\gamma \to \infty$. Therefore $F_n(\gamma) = S$ has a unique solution $\gamma_{n,S}$, and the corresponding coordinates $r_{j,n,S} := \Phi_j^{-1}(n, \gamma_{n,S})$ give the equalizer. This equalizer maximizes the real-valued max-min problem: if every index is at least γ, then each coordinate must satisfy $m_j \geq \Phi_j^{-1}(n,\gamma)$. Feasibility therefore requires $\sum_{j \in \mathcal{J}} \Phi_j^{-1}(n,\gamma) \leq S$; the largest such γ is exactly $\gamma_{n,S}$.

Let $\lambda_\alpha = T^\star_\alpha(\mu)^{-1}$ and let $w^\alpha(\mu)$ be the unique continuous optimizer from Proposition C.3. Writing $q_j = r_{j,n,S}/t$, the equalizer identity gives

$$\gamma_{n,S} = t I_j^\alpha(q_j) + \log t + \log q_j \qquad \forall j \in \mathcal{J}.$$

For the upper bound, choose j minimizing $I_j^\alpha(q_j)$. Since $\sum_j q_j = 1 - \alpha$, Proposition C.3 gives $\min_j I_j^\alpha(q_j) \leq \lambda_\alpha$, and $\log q_j \leq 0$. Hence $\gamma_{n,S} \leq t\lambda_\alpha + \log t$. For the lower bound, use $m_j = t w_j^\alpha(\mu)$. Since $\inf_{\alpha \in U, j} w_j^\alpha(\mu) > 0$, these coordinates are at least one for all sufficiently large t, so the allocation is feasible and

$$\gamma_{n,S} \geq \min_j \left\{ t I_j^\alpha(w_j^\alpha(\mu)) + \log t + \log w_j^\alpha(\mu) \right\} \geq t\lambda_\alpha + \log t - C_\mu(U).$$

This proves the estimate for $\gamma_{n,S}$ for all sufficiently large t.

We next control $r_{j,n,S}$. Let $c_U := \inf_{\alpha \in U, j} w_j^\alpha(\mu) > 0$. If $q_j < c_U/2$, then uniformly in $\alpha \in U$ the strict monotonicity of I_j^α gives $I_j^\alpha(q_j) \leq \lambda_\alpha - \rho_U$ for some $\rho_U > 0$. For large t, this contradicts the lower bound on $\gamma_{n,S}$. Thus, for all sufficiently large t, $q_j \geq c_U/2$. Since $\log q_j = O_{\mu,U}(1)$, the equalizer identity and the estimate for $\gamma_{n,S}$ imply

$$I_j^\alpha(q_j) = \lambda_\alpha + O_{\mu,U}(1/t).$$

On $[c_U/2, 1]$, the derivative $\partial_\omega I_j^\alpha(\omega)$ is uniformly bounded away from zero. The mean value theorem yields $|q_j - w_j^\alpha(\mu)| \leq C_\mu(U)/t$, the asserted estimate for $r_{j,n,S}$ for all sufficiently large t. For any fixed upper bound on t, only finitely many integer n occur, and their admissible real budgets S form compact sets. Since $1 \leq r_{j,n,S} \leq S$, continuity of $\Phi_j(n,\cdot)$ bounds $\gamma_{n,S} = \Phi_j(n, r_{j,n,S})$ on these sets. The quantities $t\lambda_\alpha$, $\log t$, and $t w_j^\alpha(\mu)$ are bounded there as well. Enlarging the remainder constants therefore gives both estimates uniformly over the small-budget range.

Finally, assume S is an integer. The continuous equalizer maximizes the real-valued max-min problem, so $\Gamma^{\mathrm{disc}}_{n,S}(\mu) \leq \gamma_{n,S}$. Conversely, Lemma C.5 rounds $r_{n,S}$ to integers $\overline{r}_j \geq 1$ with $\sum_j \overline{r}_j = S$ and $|\overline{r}_j - r_{j,n,S}| \leq 1$. Lemma E.1 gives $\Phi_j(n, \overline{r}_j) \geq \gamma_{n,S} - C_\mu(U)$, and hence $\Gamma^{\mathrm{disc}}_{n,S}(\mu) \geq \gamma_{n,S} - C_\mu(U)$. □

E.2 Guard-Count Control

The number of guard activations grows at most as the square root of time. We prove this by counting activations within epochs of the guard threshold.

Lemma E.3 (Deterministic bound on the number of guard activations)**.** *For the enabled guard in Algorithm 1, $G_{t+1} = 1$ if and only if $\min_{i\in[K]} N_{t,i} < \sqrt{t/K}$, and each guarded pull selects a least-sampled arm. Then for all $t \geq K$,*

$$G_t^{\mathrm{tot}} := \sum_{u=K}^{t-1} G_{u+1} \leq K\left(1 + \left\lfloor \sqrt{t/K} \right\rfloor\right) \leq \sqrt{Kt} + K.$$

Proof. There are at most K guard activations before the threshold leaves each epoch. For each integer $r \geq 0$, define

$$\mathcal{I}_r := \left\{u \geq K : r \leq \sqrt{u/K} < r+1\right\}.$$

We first show by induction that every arm has count at least r at time Kr^2, for each integer $r \geq 1$. Initialization gives the case $r = 1$. Assume the claim at Kr^2. Within $\mathcal{I}_r$, every guarded pull raises an arm's count from r to $r+1$, so at most K such pulls are possible. After the first round of this epoch, the threshold strictly exceeds r; while any count remains r, every subsequent round is guarded. The epoch has $K(2r+1)$ rounds, so all counts reach $r+1$ before the next epoch begins. This proves the induction and the bound of K guard pulls per epoch. Once all counts reach $r+1$, the guard cannot activate again in $\mathcal{I}_r$. There are at most $1 + \left\lfloor \sqrt{t/K} \right\rfloor$ epochs up to time t, which proves the claim. □

E.3 Uniform Challenger-Index Lower Bound

We track deficits in the challenger counts to bound the smallest population index relative to the discrete benchmark. Let $U \Subset (0,1)$ be a compact interval containing β, and let t_U be a U-stabilization time:

$$B_{t+1} = i^\star, \qquad \alpha_t := \frac{N_{t,i^\star}}{t} \in U, \qquad \forall t \geq t_U.$$

Here the local term U-stabilization time means a random time $t_U \in \{K, K+1, \dots\} \cup \{\infty\}$ satisfying these conditions, even if β is an endpoint of U. The main-text convention additionally requires β to lie in the interior of U; all applications to Proposition 5.1 use that convention. For $s \geq t_U$, define

$$n_s := N_{s,i^\star}, \qquad S_s := s - n_s, \qquad m_{s,j} := N_{s,j}, \quad j \in \mathcal{J}.$$

Then $\sum_{j\in\mathcal{J}} m_{s,j} = S_s$ and $Q_s(j) = \Phi_j(n_s, m_{s,j})$. Define the empirical penalized index and recall its population counterpart:

$$P_t(j) := V_t(i^\star, j) + \log N_{t,j}, \qquad Q_t(j) = V_t^\mu(i^\star, j) + \log N_{t,j}.$$

For $\eta \in (0,1)$, define the uniform error bound

$$\varepsilon_t := C_\mu\left(\sqrt{tH_t(\eta)} + H_t(\eta)\right).$$

On the event $\mathcal{G}_\eta$ from Theorem 4.2, simultaneously for all $t \geq K$ and all $j \in \mathcal{J}$,

$$|P_t(j) - Q_t(j)| = |V_t(i^\star, j) - V_t^\mu(i^\star, j)| \leq \varepsilon_t.$$

Lemma E.4 (Challenger-budget uniform lower bound)**.** *Fix a compact interval $U \Subset (0,1)$ containing β. There exists a constant $C_\mu(U) < \infty$ such that, on $\mathcal{G}_\eta$, for every U-stabilization time t_U and every $s \geq t_U$,*

$$\min_{j\in\mathcal{J}} Q_s(j) \geq \Gamma^{\mathrm{disc}}_{n_s,S_s}(\mu) - C_\mu(U)\left(t_U + G_s^{\mathrm{tot}} + \varepsilon_s + \sup_{t_U \leq u \leq s} u|\alpha_u - \beta| + 1\right). \tag{E.1}$$

For the unguarded rule, $G_s^{\mathrm{tot}} \equiv 0$.

Proof. We show that unguarded challenger pulls serve a sufficiently large deficit, then bound the total positive deficit to control every population index. Let $\overline{S}_\mu$ be the uniform cutoff from Lemma E.2 and set

$$\overline{\alpha}_U := \sup U < 1, \qquad s_{\mathrm{eq}} := \left\lceil \frac{\overline{S}_\mu}{1-\overline{\alpha}_U} \right\rceil.$$

If $u \geq s_{\mathrm{eq}}$ and $\alpha_u \in U$, then $S_u = (1-\alpha_u)u \geq \overline{S}_\mu$, so the continuous equalizer $(\gamma_{n_u,S_u}, r_{n_u,S_u})$ exists. If $s < s_{\mathrm{eq}}$, only finitely many count configurations are possible; after enlarging $C_\mu(U)$, (E.1) holds because the factor in parentheses is at least one. We therefore assume $s \geq s_\bullet := \max\{t_U, s_{\mathrm{eq}}\}$. Then the equalizer exists for every $u \in [s_\bullet, s]$.

Define the fixed-β challenger target

$$p_j^\beta := \frac{w_j^\beta(\mu)}{1-\beta}, \qquad j \in \mathcal{J},$$

so that $p_j^\beta > 0$ and $\sum_j p_j^\beta = 1$. For $u \in [s_\bullet, s]$, set

$$D_{u,j} := p_j^\beta S_u - N_{u,j}, \qquad j \in \mathcal{J}.$$

Write $D_u := (D_{u,j})_{j\in\mathcal{J}}$; then $\sum_j D_{u,j} = 0$. Lemma E.2 and the identity $S_u = (1-\alpha_u)u$ give

$$r_{j,n_u,S_u} - p_j^\beta S_u = u\left(w_j^{\alpha_u}(\mu) - \frac{1-\alpha_u}{1-\beta} w_j^\beta(\mu)\right) + O_{\mu,U}(1).$$

Consider the map

$$\alpha \mapsto w_j^\alpha(\mu) - \frac{1-\alpha}{1-\beta} w_j^\beta(\mu).$$

By Proposition C.3, its derivative $\partial_\alpha w_j^\alpha(\mu) + p_j^\beta$ is uniformly bounded for $\alpha, \beta \in U$, since $p_j^\beta \leq 1$. The map vanishes at $\alpha = \beta$, so, uniformly over $u \in [s_\bullet, s]$,

$$|r_{j,n_u,S_u} - p_j^\beta S_u| \leq C_\mu(U)(u|\alpha_u - \beta| + 1). \tag{E.2}$$

We first relate the selected arm's deficit to the largest deficit. Let $u+1 \leq s$ be an unguarded challenger-pull round with $u \geq s_\bullet$, and let $a = C_{u+1}$. Put $M_u := \max_{j\in\mathcal{J}} D_{u,j}$. We claim that, for constants $\kappa_\mu(U) \in (0,1]$ and $A_\mu(U) < \infty$,

$$D_{u,a} \geq \kappa_\mu(U) M_u - A_\mu(U)(\varepsilon_u + u|\alpha_u - \beta| + 1). \tag{E.3}$$

Let $j^+ \in \arg\max_{j\in\mathcal{J}} D_{u,j}$ and define $e_{u,j} := r_{j,n_u,S_u} - N_{u,j}$. By (E.2),

$$e_{u,j^+} \geq M_u - C_\mu(U)(u|\alpha_u - \beta| + 1). \tag{E.4}$$

Choose the threshold constant in the following case split larger than both the error coefficient in (E.4) and the error coefficient in the lower bound for $\gamma_{n_u,S_u} - Q_u(a)$ below divided by $c_\mu(U)$. These coefficients are fixed by the slope bounds and (E.2). If $M_u \leq C_\mu(U)(\varepsilon_u + u|\alpha_u - \beta| + 1)$, then (E.3) follows after increasing $A_\mu(U)$, because the zero-sum property implies $D_{u,a} \geq -d_K M_u$. Otherwise $e_{u,j^+} > 0$, and the slope lower bound in Lemma E.1 gives

$$Q_u(j^+) \leq \gamma_{n_u,S_u} - c_\mu(U) M_u + C_\mu(U)(u|\alpha_u - \beta| + 1).$$

Since a minimizes P_u and $|P_u(j) - Q_u(j)| \leq \varepsilon_u$ on $\mathcal{G}_\eta$,

$$Q_u(a) \leq Q_u(j^+) + 2\varepsilon_u.$$

The difference $\gamma_{n_u,S_u} - Q_u(a)$ is bounded below by $c_\mu(U)M_u - C_\mu(U)(\varepsilon_u + u|\alpha_u - \beta| + 1)$. In the present case the choice of threshold makes this lower bound positive, so $e_{u,a} > 0$, and the slope upper bound yields

$$e_{u,a} \geq c'_\mu(U) M_u - C'_\mu(U)(\varepsilon_u + u|\alpha_u - \beta| + 1).$$

Combining this with (E.2) proves (E.3).

For a real x, write $(x)_+ := \max\{x, 0\}$. Define $D_u^+ := \sum_{j\in\mathcal{J}} (D_{u,j})_+$. Since the deficits sum to zero, $M_u \leq D_u^+$. Fix the final time s and set

$$\Xi_s := \varepsilon_s + \sup_{t_U \leq v \leq s} v|\alpha_v - \beta| + 1, \qquad \Lambda_s := \frac{A_\mu(U)\Xi_s + 1}{\kappa_\mu(U)}.$$

Since ε_u is nondecreasing, (E.3) gives the following bound for every unguarded challenger pull $u+1 \le s$ with $u \ge s_\bullet$:

$$D_{u,C_{u+1}} \ge \kappa_\mu(U) M_u - A_\mu(U)\Xi_s.$$

If $u+1$ is a leader pull, then $D_{u+1} = D_u$. If it is a challenger pull of arm a, then

$$D_{u+1,j} = D_{u,j} + p_j^\beta - \mathbf{1}\{j = a\}.$$

Consider an unguarded challenger pull. If $M_u > \Lambda_s$, then $D_{u,a} > 1$. Adding $p^\beta := (p_j^\beta)_{j\in\mathcal{J}}$ can increase D^+ by at most $\sum_j p_j^\beta = 1$, while subtracting one from coordinate a decreases the positive mass by exactly one, since that coordinate remains positive. Hence $D_{u+1}^+ \le D_u^+$. If $M_u \le \Lambda_s$, then $D_u^+ \le d_K \Lambda_s$, and a single challenger update can increase D^+ by at most one. Therefore every unguarded challenger pull satisfies

$$D_{u+1}^+ \le \max\{D_u^+, d_K\Lambda_s + 1\}.$$

A guard pull need not serve a large deficit: if it pulls $i^\star$, the deficit vector D is unchanged; if it pulls a suboptimal arm, the same update holds, increasing D^+ by at most one.

Induction from $s_\bullet$ to s gives

$$D_s^+ \le \max\{D_{s_\bullet}^+, d_K\Lambda_s + 1\} + G_s^{\mathrm{tot}}.$$

Moreover,

$$D_{s_\bullet}^+ \le \sum_{j\in\mathcal{J}} p_j^\beta S_{s_\bullet} = S_{s_\bullet} \le s_\bullet \le t_U + s_{\mathrm{eq}},$$

and s_{eq} is absorbed into the constant. Hence

$$D_s^+ \le C_\mu(U)\left(t_U + G_s^{\mathrm{tot}} + \varepsilon_s + \sup_{t_U \le u \le s} u|\alpha_u - \beta| + 1\right).$$

By (E.2), this implies

$$\max_{j\in\mathcal{J}}(r_{j,n_s,S_s} - N_{s,j})_+ \le C_\mu(U)\left(t_U + G_s^{\mathrm{tot}} + \varepsilon_s + \sup_{t_U \le u \le s} u|\alpha_u - \beta| + 1\right).$$

Let the right-hand side be E_s. Then $N_{s,j} \ge r_{j,n_s,S_s} - E_s$ for all j. If $N_{s,j} \ge r_{j,n_s,S_s}$, monotonicity gives $Q_s(j) \ge \gamma_{n_s,S_s}$. Otherwise Lemma E.1 gives $Q_s(j) \ge \gamma_{n_s,S_s} - C_\mu(U)E_s$. Thus

$$\min_{j\in\mathcal{J}} Q_s(j) \ge \gamma_{n_s,S_s} - C_\mu(U)E_s.$$

Using $\gamma_{n_s,S_s} \ge \Gamma^{\mathrm{disc}}_{n_s,S_s}(\mu) - C_\mu(U)$ from Lemma E.2 proves the lemma. □

E.4 Proof of Proposition 5.1

We combine Lemma E.4 with leader-fraction concentration to prove the uniform evidence bound.

Proof of Proposition 5.1. We absorb guard and allocation fluctuations into the confidence remainder, then cancel the logarithmic count penalty using the discrete benchmark. Fix the compact interval U from Proposition 5.1, with β in its interior, and $\eta \in (0,1)$. Recall

$$\widetilde{H}_s(\eta) := \log\frac{K}{\eta} + \log(es).$$

Let $\mathcal{A}_\eta$ be the event from Proposition F.1. Define

$$\mathcal{E}_{\mathrm{env}}(\eta) := \mathcal{G}_\eta \cap \mathcal{A}_\eta.$$

By Theorem 4.2 and Proposition F.1, $\mathbb{P}_\mu(\mathcal{E}_{\mathrm{env}}(\eta)) \ge 1 - 2\eta$. On $\mathcal{E}_{\mathrm{env}}(\eta)$, Lemma E.4 gives, for every $s \ge t_U$,

$$\min_{j\in\mathcal{J}} Q_s(j) \ge \Gamma^{\mathrm{disc}}_{n_s,S_s}(\mu) - C_\mu(U)\left(t_U + G_s^{\mathrm{tot}} + \varepsilon_s + \sup_{t_U \le u \le s} u|\alpha_u - \beta| + 1\right).$$

On $\mathcal{A}_\eta$, Proposition F.1 applied with $m = t_U$ gives, uniformly over $u \in [t_U, s]$,

$$u|\alpha_u - \beta| \le C_\beta t_U + C_\beta \sqrt{s\widetilde{H}_s(\eta)} + G_s^{\mathrm{tot}}.$$

If the guard is disabled, $G_s^{\mathrm{tot}} = 0$. If it is enabled, Lemma E.3 gives $G_s^{\mathrm{tot}} \le \sqrt{Ks} + K \le C_K \sqrt{s\widetilde{H}_s(\eta)}$. Since

$$\varepsilon_s \le C_\mu \left(\sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta) \right),$$

collecting terms gives the population-index bound with constants $C_\mu(U)$ and $\overline{C}_\mu(U)$ as in Proposition 5.1.

To pass from population to empirical indices, on $\mathcal{G}_\eta$ we have

$$|P_s(j) - Q_s(j)| \le C_\mu \left(\sqrt{sH_s(\eta)} + H_s(\eta) \right) \le C_\mu \left(\sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta) \right).$$

Hence

$$\min_j P_s(j) \ge \min_j Q_s(j) - C_\mu \left(\sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta) \right).$$

After entry into the stabilized regime, $B_{s+1} = i^\star$, so $\widehat{\mu}_{s,i^\star} \ge \widehat{\mu}_{s,j}$ for every $j \ne i^\star$. Therefore $W_s(i^\star, j) = V_s(i^\star, j)$ for every $j \ne i^\star$, and

$$W_s(i^\star, j) = P_s(j) - \log N_{s,j} \ge P_s(j) - \log s.$$

Combining this with Proposition 4.1,

$$\Gamma^{\mathrm{disc}}_{n_s, S_s}(\mu) \ge \frac{s}{T^\star_{\alpha_s}(\mu)} + \log s - C_\mu(U),$$

cancels the logarithmic count penalty $\log s$ and proves the transportation-cost bound after enlarging the same two constants if needed. □

E.5 Challenger Growth and Separation

The evidence bound forces every suboptimal-arm count to grow linearly. We use this growth to bound the time at which all separation thresholds are met.

Proof of Proposition 5.3. We compare the population-index lower bound with an upper bound in terms of the suboptimal-arm count. Fix $\eta \in (0,1)$ and work on $\mathcal{E}_{\mathrm{env}}(\eta)$ from Proposition 5.1. For any U-stabilization time t_U, $s \ge t_U$, and $j \ne i^\star$, that proposition gives

$$Q_s(j) \ge \frac{s}{T^\star_{\alpha_s}(\mu)} - C_\mu(U) \left(t_U + \sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta) \right) - \overline{C}_\mu(U).$$

Since $\alpha_s \in U$, the leading term is at least $\lambda_{\mu,U} s$. On the other hand, evaluating the population transportation cost at $x = \mu_{i^\star}$ yields

$$Q_s(j) = V^\mu_s(i^\star, j) + \log N_{s,j} \le d(\mu_j, \mu_{i^\star})\, N_{s,j} + \log s.$$

Combining the two inequalities gives the claim. □

Lemma E.5 (Logarithmic separation thresholds)**.** *There exists $A_\mu < \infty$ such that for all $\eta \in (0, 1/2]$,*

$$n_\star(\eta) \le A_\mu \left(\log \frac{K}{\eta} + \log\log \frac{e}{\eta} + 1 \right).$$

Proof. We compare the logarithmic confidence terms with a sample count proportional to the confidence scale. For each $j \ne i^\star$, let $\Delta_j := \mu_{i^\star} - \mu_j > 0$. The condition defining $n_j(\eta)$ is

$$\frac{c_0}{2n} \left(\log \frac{8K^2}{\eta} + \log\log(en) \right) \le \frac{\Delta_j^2}{16}.$$

For a fixed positive constant A to be chosen, set

$$n = \left\lceil A \left(\log \frac{K}{\eta} + \log\log \frac{e}{\eta} + 1 \right) \right\rceil .$$

At this chosen count, $n \le C_K(A+1)\log(e/\eta)$, and hence $\log\log(en) \le C_K + \log(A+1) + \frac{1}{2}\log(e/\eta)$ uniformly over $\eta \in (0, 1/2]$. The numerator in the defining condition is therefore bounded by a constant multiple of the sum of the confidence scale above and $\log(A+1)$. Since A dominates $\log(A+1)$ as A increases, a sufficiently large A makes the condition hold for every $j \neq i^\star$. By the minimum defining $n_j(\eta)$, $n_j(\eta) \le n$; no estimate at larger counts is needed. Taking the maximum over j and absorbing the ceiling into A_μ proves the claim. □

Proof of Corollary 6.1. We choose a time proportional to the stabilization time and the confidence scale, then apply the count-growth bound to meet every separation threshold. Fix $\eta \in (0, 1/4]$ and work on $\mathcal{E}_{\mathrm{env}}(\eta)$ for any U-stabilization time t_U. Let

$$R_\eta := \log \frac{K}{\eta} + \log\log \frac{e}{\eta} + 1.$$

By Lemma E.5, $n_\star(\eta) \le A_\mu R_\eta$. Choose

$$s = \lceil C_{\mu,\beta}(U)(t_U + R_\eta) \rceil ,$$

with the constant to be fixed. For any $\rho > 0$, the elementary inequalities $\sqrt{s\widetilde{H}_s(\eta)} \le \rho s + C_\rho \widetilde{H}_s(\eta)$ and $\log(es) \le \rho s + C_\rho$ imply

$$\sqrt{s\widetilde{H}_s(\eta)} + \widetilde{H}_s(\eta) + \log s \le \rho s + C_{\mu,\rho}(U) R_\eta.$$

Taking ρ small in Proposition 5.3 gives, for each $j \neq i^\star$,

$$N_{s,j} \ge c_{\mu,j}(U)s - C'_{\mu,j}(U)t_U - C''_{\mu,j}(U)R_\eta.$$

Choosing $C_{\mu,\beta}(U)$ sufficiently large ensures $N_{s,j} \ge n_j(\eta)$ for all $j \neq i^\star$. Since $s \ge t_U$, entry into the stabilized regime gives $N_{s,i^\star} = \alpha_s s \ge (\inf U)s$, and the same choice of $C_{\mu,\beta}(U)$ ensures $N_{s,i^\star} \ge n_\star(\eta)$. Therefore $s \ge t_{\mathrm{sep}}(\eta)$, which proves the claim. □

F Leader-Fraction Control and Threshold Inversion

We bound deviations of the best-arm fraction after permanent leader correctness and invert the stopping threshold. Together with the evidence envelope, these estimates prove the stopping bounds in Theorem 5.2.

F.1 Uniform Leader-Fraction Concentration

We bound fluctuations of the sampling coins simultaneously over all start times, so the bound also applies at a random stabilization time. All algorithmic time indices refer to the continued sampling process defined in Section 5.

For the start-time range in Proposition F.1, extend the indicators to initialization rounds by setting $G_{u+1} := 1$ and $B_{u+1} := I_{u+1}$ for $1 \le u < K$. Here I_{u+1} is the prescribed initialization pull, so B_{u+1} is only an auxiliary arm label in this range. This convention counts initialization pulls in $G_t^{\mathrm{tot}}(m)$ when $m < K$; the algorithm and the total guard count G_t^{tot}, which starts after initialization, are unchanged. In this proposition, also set $\alpha_t := N_{t,i^\star}/t$ for every integer $t \ge 1$.

Proposition F.1 (Leader-fraction control, uniform in the start time)**.** *There exists a constant $C_\beta < \infty$ such that for every $\eta \in (0,1)$ there exists an event $\mathcal{A}_\eta$ with $\mathbb{P}_\mu(\mathcal{A}_\eta) \ge 1 - \eta$ on which the following holds. For all integers $m \ge 1$ and $t \ge m$, if $B_{u+1} = i^\star$ for every integer $u \ge m$, then*

$$|(N_{t,i^\star} - N_{m,i^\star}) - \beta(t-m)| \le C_\beta \sqrt{(t-m+1)\log\left(\frac{etm^2}{\eta}\right)} + G_t^{\mathrm{tot}}(m),$$

where $G_t^{\mathrm{tot}}(m) := \sum_{u=m}^{t-1} G_{u+1}$ is the number of guard activations and any initialization pulls in $\{m+1, \dots, t\}$. Consequently,

$$|\alpha_t - \beta| \le \frac{(1+\beta)m}{t} + C_\beta \sqrt{\frac{1}{t} \log\left(\frac{etm^2}{\eta}\right)} + \frac{G_t^{\mathrm{tot}}(m)}{t}.$$

Proof. We apply Azuma–Hoeffding to the centered sampling coins and take a union bound over horizons and start times. Guarded and initialization pulls contribute at most their number to the error. Fix $\eta \in (0,1)$. For each $m \ge 1$, set $\eta_m := 6\eta/(\pi^2 m^2)$ so that $\sum_{m \ge 1} \eta_m \le \eta$.

For the analysis, draw an independent $U_{u+1} \sim \mathrm{Ber}(\beta)$ at every round, including guarded and initialization rounds where the coin is unused. Unused coins do not affect sampling decisions, so this construction leaves the algorithm's law unchanged. Let $\mathcal{F}_u$ be the sigma-field generated by the sampled arms, observed rewards, and strategy randomization through round u, including unused coins through that round.

Fix an integer $m \ge 1$. For integers $u \ge m$, define

$$Z_{u+1}^{(m)} := (U_{u+1} - \beta)\mathbf{1}\{G_{u+1} = 0\}.$$

Since G_{u+1} is $\mathcal{F}_u$-measurable and U_{u+1} is independent of $\mathcal{F}_u$ with mean β, one has

$$\mathbb{E}_\mu\left[Z_{u+1}^{(m)} \mid \mathcal{F}_u\right] = 0, \qquad |Z_{u+1}^{(m)}| \le 1.$$

Thus

$$S_t^{(m)} := \sum_{u=m}^{t-1} Z_{u+1}^{(m)}, \qquad t \ge m,$$

is a martingale with bounded increments under $\mathbb{P}_\mu$.

For $n \ge 1$, put $t = m + n$. Azuma–Hoeffding gives

$$\mathbb{P}_\mu\left(|S_{m+n}^{(m)}| \ge \sqrt{2n \log\left(\frac{12n^2}{\eta_m}\right)}\right) \le \frac{\eta_m}{6n^2}.$$

A union bound over $n \ge 1$ gives, with probability at least $1 - \eta_m$,

$$|S_t^{(m)}| \le C_\beta \sqrt{(t-m+1)\log\left(\frac{etm^2}{\eta}\right)} \qquad \forall t \ge m,$$

after increasing the numerical constant C_β.

Suppose that $B_{u+1} = i^\star$ for every integer $u \ge m$. On every unguarded round $u+1$ in this range, the algorithm pulls $i^\star$ if and only if $U_{u+1} = 1$. Initialization rounds have $G_{u+1} = 1$ and therefore contribute only to the second sum below. For every $t \ge m$,

$$\begin{aligned}
&(N_{t,i^\star} - N_{m,i^\star}) - \beta(t-m) \\
&= \sum_{u=m}^{t-1} (U_{u+1} - \beta)\mathbf{1}\{G_{u+1} = 0\} + \sum_{u=m}^{t-1} \left(\mathbf{1}\{I_{u+1} = i^\star\} - \beta\right)\mathbf{1}\{G_{u+1} = 1\}.
\end{aligned}$$

The first sum is $S_t^{(m)}$. The second sum has absolute value at most $G_t^{\mathrm{tot}}(m)$, since each guarded or initialization pull contributes a term of absolute value at most one. Therefore,

$$|(N_{t,i^\star} - N_{m,i^\star}) - \beta(t-m)| \le C_\beta \sqrt{(t-m+1)\log\left(\frac{etm^2}{\eta}\right)} + G_t^{\mathrm{tot}}(m).$$

Intersecting the events over all $m \ge 1$ and using $\sum_{m \ge 1} \eta_m \le \eta$ gives $\mathcal{A}_\eta$. The fraction bound follows by adding the initial discrepancy $N_{m,i^\star} - \beta m$ and dividing by t. □

F.2 Entrance into a Compact Neighborhood of β

We use the fraction bound to control when α_t remains near β. Let $t_{\mathrm{lead}} \geq K$ denote an integer time from which $B_{u+1} = i^\star$ for every integer $u \geq t_{\mathrm{lead}}$ on the continued sampling process. This time requires only permanent leader correctness; t_U also requires $\alpha_t \in U$ thereafter.

Lemma F.2 (Entrance into a fixed compact neighborhood of β)**.** *Fix $\eta \in (0,1)$ and a compact interval $U = [\beta - \rho, \beta + \rho] \Subset (0,1)$ with $\rho > 0$. There exists $C_\beta(U) < \infty$ such that the following holds. On the event of Proposition F.1, if $t_{\mathrm{lead}} \leq L$ with $L \geq 1$, then*

$$\alpha_t \in U \qquad \forall t \geq C_\beta(U) \left(L + K + \log \frac{e(L+1)^2}{\eta} \right).$$

When the guard is disabled, the additive K term can be dropped.

Proof. We bound the initial discrepancy, coin fluctuations, and guarded-pull fraction uniformly beyond a time proportional to L and the confidence scale. If the guard is enabled, Proposition F.1 yields for all $t \geq t_{\mathrm{lead}}$,

$$|\alpha_t - \beta| \leq \frac{(1+\beta)L}{t} + C_\beta \sqrt{\frac{1}{t} \log \left(\frac{et(L+1)^2}{\eta} \right)} + \frac{G_t^{\mathrm{tot}}(t_{\mathrm{lead}})}{t}.$$

By Lemma E.3, the last term is at most $\sqrt{K/t} + K/t$, and it is absent in the unguarded case. Let

$$A := L + K + \log \frac{e(L+1)^2}{\eta}.$$

For $t \geq CA$, the terms for the initial discrepancy and guarded pulls are each at most $\rho/3$ once C is large enough. The coin-fluctuation term is also at most $\rho/3$: the map $t \mapsto \log(et(L+1)^2/\eta)/t$ is decreasing above an absolute constant, and at $t = CA$ its numerator is bounded by a constant multiple of A. Increasing C if necessary proves the claim. □

F.3 Centered Quantile Integration

The next results convert upper-quantile bounds into expectation bounds by integrating the remainder above a deterministic leading term. In this subsection, $\mathbb{P}$ and $\mathbb{E}$ denote a generic probability law and its expectation.

Theorem F.3 (Centered quantile integration)**.** *Let $X \geq 0$ be a random variable. Let $A \geq 0$ and let $\rho_X : (0,1] \to [0,\infty)$ satisfy*

$$\mathbb{P}(X \leq A + \rho_X(\eta)) \geq 1 - \eta \qquad \forall \eta \in (0,1].$$

Define the monotone envelope

$$\rho_X^{\downarrow}(\eta) := \sup_{\xi \in [\eta, 1]} \rho_X(\xi).$$

Then

$$\mathbb{E}[X] \leq A + \int_0^1 \rho_X^{\downarrow}(u)\, du.$$

Moreover, if

$$\sum_{k \geq 0} 2^{-k} \rho_X^{\downarrow}(2^{-k}) < \infty,$$

then

$$\mathbb{E}[X] \leq A + \sum_{k \geq 0} 2^{-k} \rho_X^{\downarrow}(2^{-k}).$$

Proof. We integrate the upper-tail quantile, then bound the remainder integral on dyadic intervals. Let

$$q_X(\eta) := \inf\{x \geq 0 : \mathbb{P}(X \leq x) \geq 1 - \eta\}, \qquad \eta \in (0,1],$$

be the upper-tail quantile of X. The assumed upper-quantile inequality gives

$$q_X(\eta) \leq A + \rho_X(\eta) \leq A + \rho_X^{\downarrow}(\eta) \qquad \forall \eta \in (0,1].$$

The quantile representation of expectation for nonnegative random variables therefore yields

$$\mathbb{E}[X] = \int_0^1 q_X(u)\,du \le A + \int_0^1 \rho_X^{\downarrow}(u)\,du.$$

For the dyadic bound, $\rho_X^{\downarrow}$ is nonincreasing in η, so for $u \in [2^{-k-1}, 2^{-k}]$,

$$\rho_X^{\downarrow}(u) \le \rho_X^{\downarrow}(2^{-k-1}).$$

Consequently,

$$\int_0^1 \rho_X^{\downarrow}(u)\,du \le \sum_{k\ge 0} 2^{-k-1}\rho_X^{\downarrow}(2^{-k-1}) \le \sum_{k\ge 0} 2^{-k}\rho_X^{\downarrow}(2^{-k}).$$

Combining this estimate with the first part proves the theorem. □

Lemma F.4 (Monotone summable upper-quantile bound)**.** *Let $T \ge 0$ be a random time admitting a summable upper-quantile bound u_T, namely $\mathbb{P}(T \le u_T(\eta)) \ge 1-\eta, \forall \eta \in (0,1]$, and $\sum_{k\ge 0} 2^{-k} u_T(2^{-k}) < \infty$. Define the upper-tail quantile*

$$q_T(\eta) := \inf\{x \ge 0 : \mathbb{P}(T \le x) \ge 1-\eta\}, \qquad \eta \in (0,1].$$

Then q_T is a nonincreasing upper-quantile bound for T, and $\sum_{k\ge 0} 2^{-k} q_T(2^{-k}) < \infty$.

Proof. We verify the quantile's confidence property and monotonicity, then compare its dyadic series with that of u_T. By definition of q_T, one has

$$\mathbb{P}(T \le q_T(\eta)) \ge 1-\eta \qquad \forall \eta \in (0,1].$$

Thus q_T is an upper-quantile bound. If $0 < \eta_1 \le \eta_2 \le 1$, then $1-\eta_1 \ge 1-\eta_2$, so the confidence requirement at level η_1 is at least as stringent as the one at level η_2. Therefore

$$q_T(\eta_1) \ge q_T(\eta_2),$$

so q_T is nonincreasing in η.

Since u_T is also an upper-quantile bound, the definition of q_T gives

$$q_T(\eta) \le u_T(\eta) \qquad \forall \eta \in (0,1].$$

Hence

$$\sum_{k\ge 0} 2^{-k} q_T(2^{-k}) \le \sum_{k\ge 0} 2^{-k} u_T(2^{-k}) < \infty.$$

□

Lemma F.5 (Finite mean gives summable dyadic quantiles)**.** *Let $X \ge 0$ satisfy $\mathbb{E}[X] < \infty$, and let*

$$q_X(\eta) := \inf\{x \ge 0 : \mathbb{P}(X \le x) \ge 1-\eta\}, \qquad \eta \in (0,1].$$

Then

$$\sum_{k\ge 0} 2^{-k} q_X(2^{-k}) < \infty.$$

Proof. We use quantile monotonicity to bound each dyadic term by a quantile integral. The map $\eta \mapsto q_X(\eta)$ is nonincreasing. For each $k \ge 0$ and $u \in [2^{-k-1}, 2^{-k}]$, one has $q_X(u) \ge q_X(2^{-k})$. Hence

$$2^{-k} q_X(2^{-k}) \le 2 \int_{2^{-k-1}}^{2^{-k}} q_X(u)\,du.$$

Summing over k and using the quantile representation of expectation gives

$$\sum_{k\ge 0} 2^{-k} q_X(2^{-k}) \le 2 \int_0^1 q_X(u)\,du = 2\mathbb{E}[X] < \infty.$$

□

F.4 Threshold Inversion

We give a deterministic threshold-crossing bound that is uniform in the entrance delay L. Here T is a positive scalar parameter, and time arguments in c and $\widetilde{H}$ are real.

Theorem F.6 (Threshold inversion with a linear entrance delay)**.** *Fix constants $T > 0$, $C_0 \geq 0$, and $a, b, d \geq 0$. For $\eta \in (0, 1]$ and $t \geq 1$, define*

$$\widetilde{H}_t(\eta) := \log \frac{K}{\eta} + \log(et).$$

There exists a finite constant A_{inv}, depending only on T, C_0, a, b, d and K, such that the following holds. For every $\delta \in (0, 1/2]$ and $\eta \in (0, 1/4]$, set

$$L_\delta := \log \frac{1}{\delta}, \qquad M_{\delta,\eta} := \log \frac{K}{\eta} + \log\log \frac{e}{\delta},$$

and

$$s_{\delta,\eta} := \left\lceil TL_\delta + A_{\mathrm{inv}} \sqrt{L_\delta M_{\delta,\eta}} + A_{\mathrm{inv}} M_{\delta,\eta} + A_{\mathrm{inv}} \right\rceil.$$

Then, for every $L \geq 0$ and

$$t = (1 + C_0)L + s_{\delta,\eta},$$

one has

$$\frac{L + s_{\delta,\eta}}{T} - a\sqrt{t\widetilde{H}_t(\eta)} - b\widetilde{H}_t(\eta) - d > c(t, \delta).$$

Proof. We choose A_{inv} so that the threshold residual is positive at $L = 0$ and increasing for $L \geq 0$. Write

$$\overline{c} := 1 + C_0, \qquad s := s_{\delta,\eta}, \qquad M := M_{\delta,\eta}, \qquad L_\delta := \log \frac{1}{\delta}.$$

Because $\delta \in (0, 1/2]$ and $\eta \in (0, 1/4]$, there is a positive constant $m_0 = m_0(K)$ such that $M \geq m_0$.

We first record two elementary bounds. Put

$$R := \sqrt{L_\delta M} + M + 1.$$

For every fixed $A \geq 1$ and every $1 \leq s \leq TL_\delta + AR + 2$, there exists a constant $C_1 = C_1(T, K)$ such that

$$\widetilde{H}_s(\eta) \leq C_1(1 + \log A)M.$$

By AM–GM, $R \leq (L_\delta + 3M)/2 + 1 \leq 2(L_\delta + M + 1)$. Absorbing the factor two into C_T gives $\log(es) \leq C_T + \log A + \log(L_\delta + M + 2)$. Since $\log(L_\delta + M + 2) \leq C_K + \log(1 + L_\delta) + \log(1 + M)$ and $M = \log(K/\eta) + \log(1 + L_\delta)$, the displayed bound follows after using $M \geq m_0$. Consequently, for the value $A = A_{\mathrm{inv}}$ and the associated s,

$$\widetilde{H}_s(\eta) \leq C_1(1 + \log A_{\mathrm{inv}})M. \tag{F.1}$$

Moreover, because $s \geq A_{\mathrm{inv}} M$,

$$\frac{\widetilde{H}_s(\eta)}{s} \leq \frac{C_1(1 + \log A_{\mathrm{inv}})}{A_{\mathrm{inv}}}. \tag{F.2}$$

Define

$$\psi(L) := \frac{L + s}{T} - a\sqrt{(\overline{c}L + s)\widetilde{H}_{\overline{c}L+s}(\eta)} - b\widetilde{H}_{\overline{c}L+s}(\eta) - d - c(\overline{c}L + s, \delta).$$

For $t \geq 1$,

$$\frac{d}{dt}\widetilde{H}_t(\eta) = \frac{1}{t},$$

and

$$\frac{d}{dt}\sqrt{t\widetilde{H}_t(\eta)} = \frac{\widetilde{H}_t(\eta) + 1}{2\sqrt{t\widetilde{H}_t(\eta)}} \leq \sqrt{\frac{\widetilde{H}_t(\eta)}{t}},$$

because $\widetilde{H}_t(\eta) \geq 1$. Also

$$\frac{d}{dt}c(t,\delta) = \frac{1}{1+t/2} \leq \frac{2}{t}.$$

Hence

$$\psi'(L) \geq \frac{1}{T} - a\overline{c}\sqrt{\frac{\widetilde{H}_{\overline{c}L+s}(\eta)}{\overline{c}L+s}} - (b+2)\frac{\overline{c}}{\overline{c}L+s}.$$

The map $t \mapsto \widetilde{H}_t(\eta)/t$ is decreasing for $t \geq 1$ because $\widetilde{H}_t(\eta) \geq 1$. Using (F.2), we can choose A_{inv} large enough that

$$a\overline{c}\sqrt{\frac{\widetilde{H}_s(\eta)}{s}} + (b+2)\frac{\overline{c}}{s} \leq \frac{1}{2T},$$

uniformly in δ and η. For such a choice, $\psi'(L) \geq 1/(2T) > 0$ for all $L \geq 0$.

It remains to make $\psi(0) > 0$. Since

$$c(s,\delta) \leq L_\delta + 2\widetilde{H}_s(\eta) + C_2(K),$$

one has

$$\psi(0) \geq \frac{s}{T} - L_\delta - a\sqrt{s\widetilde{H}_s(\eta)} - (b+2)\widetilde{H}_s(\eta) - d - C_2(K).$$

The definition of s gives

$$\frac{s}{T} - L_\delta \geq \frac{A_{\mathrm{inv}}}{T}R.$$

Using (F.1) and $s \leq TL_\delta + A_{\mathrm{inv}}R + 2$, we get

$$s\widetilde{H}_s(\eta) \leq C_3(1+\log A_{\mathrm{inv}})M\left(TL_\delta + A_{\mathrm{inv}}R + 2\right).$$

Since $L_\delta M \leq R^2$, $RM \leq R^2$, and $M \leq R^2$, this implies

$$\sqrt{s\widetilde{H}_s(\eta)} \leq C_4\sqrt{A_{\mathrm{inv}}(1+\log A_{\mathrm{inv}})}\,R.$$

Also, (F.1) gives $\widetilde{H}_s(\eta) \leq C_5(1+\log A_{\mathrm{inv}})R$. Therefore

$$\psi(0) \geq \left[\frac{A_{\mathrm{inv}}}{T} - aC_4\sqrt{A_{\mathrm{inv}}(1+\log A_{\mathrm{inv}})} - (b+2)C_5(1+\log A_{\mathrm{inv}}) - d - C_2(K)\right]R.$$

Because

$$\frac{A}{\sqrt{A(1+\log A)}} \to \infty \quad \text{and} \quad \frac{A}{1+\log A} \to \infty \qquad \text{as } A \to \infty,$$

we can increase A_{inv} so that the bracket is strictly positive. The residual therefore satisfies $\psi(L) \geq \psi(0) > 0$ for every $L \geq 0$, which proves the desired inequality. □

F.5 Proof of Theorem 5.2

The preceding estimates yield the high-probability and expected stopping-time bounds in the stabilized regime.

Proof of Theorem 5.2. We combine the evidence envelope and fraction bound to obtain a threshold crossing, then integrate the centered stopping-time quantiles. Fix $\delta \in (0, 1/2]$ and $\eta \in (0, 1/4]$, and use $\widetilde{H}_t(\eta) = \log(K/\eta) + \log(et)$ as above.

Let $\mathcal{E}_{\mathrm{env}}(\eta)$ be the event from Proposition 5.1, with probability at least $1-2\eta$, and let $\mathcal{A}_\eta$ be the event from Proposition F.1, with probability at least $1-\eta$. We work on

$$\mathcal{E} := \mathcal{E}_{\mathrm{env}}(\eta) \cap \mathcal{A}_\eta, \qquad \mathbb{P}_\mu(\mathcal{E}) \geq 1 - 3\eta.$$

For all $t \geq t_U$, entry into the stabilized regime gives $B_{t+1} = i^\star$ and $\alpha_t \in U$. Hence the stopping condition is met at time t as soon as

$$\min_{j \neq i^\star} W_t(i^\star, j) > c(t,\delta).$$

On $\mathcal{E}_{\mathrm{env}}(\eta)$, Proposition 5.1 gives, for all $t \geq t_U$,

$$\min_{j \neq i^\star} W_t(i^\star, j) \geq \frac{t}{T^\star_{\alpha_t}(\mu)} - B_{\mu,0}(U) t_U - B_{\mu,1}(U)\sqrt{t\widetilde{H}_t(\eta)} - B_{\mu,2}(U)\widetilde{H}_t(\eta) - B_{\mu,3}(U),$$

for positive finite constants $B_{\mu,\ell}(U)$, $\ell \in \{0,1,2,3\}$.

By Lemma C.4, for all $\alpha \in U$,

$$\frac{1}{T^\star_\alpha(\mu)} \geq \frac{1}{T^\star_\beta(\mu)} - L_\mu(U)|\alpha - \beta|.$$

On $\mathcal{A}_\eta$, Proposition F.1 with start time $m = t_U$ implies, uniformly for $t \geq t_U$,

$$t|\alpha_t - \beta| \leq C_\beta t_U + C_\beta\sqrt{t\widetilde{H}_t(\eta)} + G^{\mathrm{tot}}_t(t_U).$$

The proof of Proposition F.1 permits a numerical C_β uniform in β; the initial discrepancy also uses only $1 + \beta \leq 2$. For the unguarded rule the last term is zero. With the guard enabled, Lemma E.3 gives $G^{\mathrm{tot}}_t(t_U) \leq \sqrt{Kt} + K \leq C_K\sqrt{t\widetilde{H}_t(\eta)}$. Thus

$$\frac{t}{T^\star_{\alpha_t}(\mu)} \geq \frac{t}{T^\star_\beta(\mu)} - A_{\mu,0}(U) t_U - A_{\mu,1}(U)\sqrt{t\widetilde{H}_t(\eta)}.$$

Combining the envelope and fraction bounds gives constants $a_\mu(U), b_\mu(U), d_\mu(U)$ and $C_{\mu,0}(U) \geq 0$ such that, for all $t \geq t_U$,

$$\min_{j \neq i^\star} W_t(i^\star, j) \geq \frac{t}{T^\star_\beta(\mu)} - \frac{C_{\mu,0}(U)}{T^\star_\beta(\mu)} t_U - a_\mu(U)\sqrt{t\widetilde{H}_t(\eta)} - b_\mu(U)\widetilde{H}_t(\eta) - d_\mu(U). \tag{F.3}$$

To choose these coefficients uniformly over $\beta \in U$, set $T_{\max} := \max_{\alpha \in U} T^\star_\alpha(\mu) < \infty$ and take $C_{\mu,0}(U) = T_{\max}(B_{\mu,0}(U) + A_{\mu,0}(U))$. The remaining coefficients already depend only on (μ, U). In the inversion proof, use $T_{\max}$ in the upper bound on s and $1/T \geq 1/T_{\max}$ in the derivative and base-value bounds. A single A_{inv} then works for all $T = T^\star_\beta(\mu)$ with $\beta \in U$.

Apply Theorem F.6 with

$$T = T^\star_\beta(\mu), \qquad C_0 = C_{\mu,0}(U), \qquad a = a_\mu(U), \qquad b = b_\mu(U), \qquad d = d_\mu(U).$$

Let $s_{\delta,\eta}$ be the time given by this theorem and set

$$t^\star := \lceil (1 + C_{\mu,0}(U)) t_U + s_{\delta,\eta} \rceil .$$

Evaluating (F.3) at $t = t^\star$ gives

$$\min_{j \neq i^\star} W_{t^\star}(i^\star, j) \geq \frac{t^\star - C_{\mu,0}(U) t_U}{T^\star_\beta(\mu)} - a_\mu(U)\sqrt{t^\star \widetilde{H}_{t^\star}(\eta)} - b_\mu(U)\widetilde{H}_{t^\star}(\eta) - d_\mu(U).$$

To apply Theorem F.6 at this integer time, take $L = (t^\star - s_{\delta,\eta})/(1 + C_{\mu,0}(U)) \geq t_U$. Then $t^\star = (1 + C_{\mu,0}(U))L + s_{\delta,\eta}$ and $t^\star - C_{\mu,0}(U) t_U \geq L + s_{\delta,\eta}$. The theorem therefore makes the right-hand side strictly larger than $c(t^\star, \delta)$. Hence $\tau_\delta \leq t^\star$ on $\mathcal{E}$, and the stated high-probability bound follows by expanding $s_{\delta,\eta}$ and absorbing both ceiling errors, each less than one, into the bounded remainder.

Now assume that t_U admits a summable upper-quantile bound u_U. The expected-bound coefficients may also depend on u_U. Let

$$L_\delta := \log \frac{1}{\delta}, \qquad A_\delta := T^\star_\beta(\mu) L_\delta.$$

The high-probability bound gives constants $D_{\mu,k}(U) < \infty$, $k \in \{0,1,2,3\}$, such that, for every $\zeta \in (0,1]$,

$$\mathbb{P}_\mu\left(\tau_\delta \leq A_\delta + R_\delta(\zeta)\right) \geq 1 - \zeta,$$

where

$$R_\delta(\zeta) = D_{\mu,0}(U)u_U(\zeta/4) + D_{\mu,1}(U)\sqrt{L_\delta\left(\log\frac{4K}{\zeta} + \log\log\frac{e}{\delta}\right)} + D_{\mu,2}(U)\left(\log\frac{4K}{\zeta} + \log\log\frac{e}{\delta}\right) + D_{\mu,3}(U).$$

This follows by applying the high-probability bound with $\eta = \zeta/4$ and intersecting it with the event $\{t_U \le u_U(\zeta/4)\}$.

Let $R_\delta^\downarrow$ be the monotone envelope of R_δ. Theorem F.3 gives

$$\mathbb{E}_\mu[\tau_\delta] \le A_\delta + \sum_{k\ge 0} 2^{-k} R_\delta^\downarrow(2^{-k}).$$

By Lemma F.4, we may take u_U to be the exact upper-tail quantile of t_U. Then R_δ is nonincreasing, so $R_\delta^\downarrow = R_\delta$. The shifted dyadic series also satisfies

$$\sum_{k\ge 0} 2^{-k} u_U(2^{-k-2}) < \infty.$$

In the remaining estimates, C denotes a finite constant independent of δ and may change between inequalities. The term $C\sqrt{L_\delta}$ is dominated by $C\sqrt{L_\delta \log\log(e/\delta)}$ for $\delta \in (0, 1/2]$. Similarly,

$$\sum_{k\ge 0} 2^{-k}\left(\log(4K) + k\log 2 + \log\log\frac{e}{\delta}\right) \le C\log\log\frac{e}{\delta} + C.$$

Combining these bounds yields

$$\mathbb{E}_\mu[\tau_\delta] \le T_\beta^\star(\mu)\log(1/\delta) + C'_{\mu,1}(U)\sqrt{\log(1/\delta)\log\log(e/\delta)} + C'_{\mu,2}(U)\log\log(e/\delta) + C'_{\mu,3}(U).$$

□

G Guarantees Under a Unique Best Arm

We bound the time to enter the stabilized regime under a unique best arm, then derive the high-probability and expected stopping-time guarantees. All algorithmic time indices refer to the continued sampling process in Section 5. Throughout, t_U is the earliest U-stabilization time defined in Section 6, with value ∞ if none exists. The countable persistence conditions make this time measurable and independent of the confidence level. Each horizon below bounds this same t_U. For each compact interval $U \Subset (0,1)$ with β in its interior, we apply Lemma F.2 to a symmetric compact subinterval centered at β. Stabilization in this subinterval implies stabilization in U.

G.1 Separation Implies Permanent Leader Correctness

The separation counts ensure that the empirical best-arm mean remains strictly above every suboptimal-arm mean.

Proof of Proposition 6.2. We compare the armwise confidence radii with one quarter of each best-arm gap. Fix $\eta \in (0,1)$. The real extension of $r(n,\eta)^2$ is decreasing for $n \ge 1$: apart from the positive factor $c_0/(2n^2)$, its derivative is $1/\log(en) - \log(8K^2/\eta) - \log\log(en) < 0$, since $K \ge 2$. Thus the radius bound at $n_j(\eta)$ holds at every larger count. Fix an integer $t \ge t_{\mathrm{sep}}(\eta)$ and an arm $j \in [K] \setminus \{i^\star\}$. By definition, $N_{t,i^\star} \ge n_\star(\eta) \ge n_j(\eta)$ and $N_{t,j} \ge n_j(\eta)$. Hence

$$r(N_{t,i^\star},\eta) \le \frac{\mu_{i^\star} - \mu_j}{4}, \qquad r(N_{t,j},\eta) \le \frac{\mu_{i^\star} - \mu_j}{4}.$$

On $\mathcal{K}_\eta^{\mathrm{sep}}$, these radius bounds yield

$$\widehat{\mu}_{t,i^\star} \ge \mu_{i^\star} - \frac{\mu_{i^\star} - \mu_j}{4}, \qquad \widehat{\mu}_{t,j} \le \mu_j + \frac{\mu_{i^\star} - \mu_j}{4}.$$

The bounds imply $\widehat{\mu}_{t,i^\star} > \widehat{\mu}_{t,j}$. Since $t \ge t_{\mathrm{sep}}(\eta)$ and $j \ne i^\star$ were arbitrary, the empirical leader is correct at every time from $t_{\mathrm{sep}}(\eta)$ onward. □

G.2 High-Probability Stabilization Bound for the Unguarded Algorithm

We bound the penalized index of any nonleader whose count is below its separation threshold. This bound ensures progress toward separation; a martingale argument then bounds the time needed to reach the separation counts.

Lemma G.1 (Polynomial-horizon index bound for separation-deficient nonleaders)**.** *Let* $p := \min\{\beta, 1-\beta\}$. *There exist positive finite constants* $\widetilde{A}_\mu, \widetilde{B}_\mu$ *such that the following holds. Fix exponents* $a > \widetilde{A}_\mu$ *and* $b > \widetilde{B}_\mu + 1$. *Then there exists a constant* $C_{\mu,\beta,a,b} < \infty$ *with the property that, for every* $\eta \in (0, 1/4]$, *if*

$$M_\eta = \left\lceil C_{\mu,\beta,a,b}\eta^{-a}(\log(e/\eta))^b \right\rceil, \qquad T_\eta = K + \left\lceil \frac{2K}{p}(M_\eta - 1) + \frac{16}{p^2}\log\frac{6}{\eta} \right\rceil,$$

then on the event $\mathcal{G}_{\eta/6} \cap \mathcal{K}_{\eta/6}$ *the following implication holds for every integer* $K \le t \le T_\eta$. *If* $j \in [K] \setminus \{B_{t+1}\}$ *and*

$$N_{t,j} < \begin{cases} n_\star(\eta/6), & j = i^\star, \\ n_j(\eta/6), & j \neq i^\star, \end{cases},$$

then

$$W_t(B_{t+1}, j) + \log N_{t,j} \le \log(M_\eta - 1).$$

Consequently, whenever such a separation-deficient nonleader exists, the challenger selected by the unguarded algorithm has count strictly smaller than M_η.

Proof. We bound the population GLR by the deficient arm's count, then apply the uniform GLR comparison and logarithmic separation-count bound. Let $j \neq B_{t+1}$ satisfy the displayed separation-deficiency condition. Then $N_{t,j} \le n_\star(\eta/6)$. Since B_{t+1} is an empirical leader, $\widehat{\mu}_{t,B_{t+1}} \ge \widehat{\mu}_{t,j}$, and therefore $W_t(B_{t+1}, j) = V_t(B_{t+1}, j)$. Write

$$D^{\max}_{\mu,\mathrm{pair}} := \max_{a,b\in[K]:a\neq b} d(\mu_a, \mu_b) < \infty.$$

Evaluating the population symmetric GLR at $x = \mu_{B_{t+1}}$ gives

$$V_t^\mu(B_{t+1}, j) \le D^{\max}_{\mu,\mathrm{pair}} n_\star(\eta/6).$$

On $\mathcal{G}_{\eta/6}$, Theorem 4.2 and $N_{\mathrm{eff},t}(B_{t+1}, j) \le N_{t,j} \le n_\star(\eta/6)$ imply

$$W_t(B_{t+1}, j) \le D^{\max}_{\mu,\mathrm{pair}} n_\star(\eta/6) + C_\mu\left(\sqrt{n_\star(\eta/6)H_{T_\eta}(\eta/6)} + H_{T_\eta}(\eta/6)\right).$$

We track the dependence on the cap constant before choosing it. Put $x := \log(e/\eta)$. Lemma E.5 gives a constant $S_\mu \ge 1$, independent of a, b, p and the cap constant, such that $n_\star(\eta/6) \le S_\mu x$ for all $\eta \in (0, 1/4]$. Write $c := \max\{1, C_\mu\}$ and $d := D^{\max}_{\mu,\mathrm{pair}}$, and choose

$$\widetilde{A}_\mu := (d + c/2)S_\mu + 3c/2, \qquad \widetilde{B}_\mu := 3c/2 + 1.$$

Fix $a > \widetilde{A}_\mu$ and $b > \widetilde{B}_\mu + 1$. For a candidate $C := C_{\mu,\beta,a,b} \ge 2S_\mu$, let $F_\eta := \eta^{-a}x^b$. Since $a > 0$ and $b > 1$, $F_\eta \ge x \ge 1$. The ceiling bounds $M_\eta - 1 \le CF_\eta$ and $\lceil u \rceil \le u + 1$, together with $\log(6/\eta) \le (\log 6)x$, give

$$T_\eta \le D_{K,p}CF_\eta, \qquad D_{K,p} := K + 1 + \frac{2K}{p} + \frac{16\log 6}{p^2}.$$

Define $\Lambda_C := 1 + \log D_{K,p} + \log C + a + b$. Using $\log(1/\eta) \le x$ and $\log x \le x$, we obtain

$$\log(eT_\eta) \le \Lambda_C x, \qquad H_{T_\eta}(\eta/6) \le \log(1/\eta) + \log x + \log(6K) + \log\Lambda_C.$$

Applying $\sqrt{nH} \le (n + H)/2$ to the preceding GLR bound and using $\log N_{t,j} \le \log(S_\mu x)$ now yields

$$W_t(B_{t+1}, j) + \log N_{t,j} \le \widetilde{A}_\mu \log(1/\eta) + \widetilde{B}_\mu \log x + R_\mu + \frac{3c}{2}\log\Lambda_C,$$

where $R_\mu := (d + c/2)S_\mu + (3c/2)\log(6K) + \log S_\mu$. For the fixed a, b, p, $\log\Lambda_C = o(\log C)$ as $C \to \infty$. We can therefore choose $C \ge 2S_\mu$ so that

$$\log C \ge R_\mu + \log 2 + \frac{3c}{2}\log\Lambda_C.$$

This choice is independent of η. Since $CF_\eta \geq 2$, the lower ceiling bound gives $M_\eta - 1 \geq CF_\eta - 1 \geq CF_\eta/2$, and hence

$$\log(M_\eta - 1) \geq \log C - \log 2 + a \log(1/\eta) + b \log x.$$

The last two displays and the choices of a, b prove the required index bound. They also give $M_\eta - 1 \geq CF_\eta/2 \geq S_\mu x \geq n_\star(\eta/6)$ uniformly in η.

For the challenger-count conclusion, any arm k with $N_{t,k} \geq M_\eta$ has penalized index at least $\log M_\eta$, since transportation costs are nonnegative. A separation-deficient nonleader has index at most $\log(M_\eta - 1)$, so a minimizer of the challenger index cannot have count at least M_η. □

Proof of Proposition 6.3. We bound the separation time using a lower bound on the conditional probability of pulls below the count cap and Azuma–Hoeffding, then control the best-arm fraction. Fix $\eta \in (0, 1/4]$ and write $\overline{\eta} := \eta/6$. Let M_η, T_η be the quantities from Lemma G.1, and set $p := \min\{\beta, 1 - \beta\}$. For integers $t \geq K$, define the progress indicator and total count deficit by

$$X_{t+1} := \mathbf{1}\{N_{t,I_{t+1}} < M_\eta\}, \qquad \mathcal{D}_t := \sum_{i=1}^{K} (M_\eta - N_{t,i})_+, \qquad t \geq K.$$

Since each arm is pulled once initially, $\mathcal{D}_K = K(M_\eta - 1)$ and $\mathcal{D}_{t+1} = \mathcal{D}_t - X_{t+1}$. Let $\tau_M := \inf\{t \geq K : \mathcal{D}_t = 0\}$. At time τ_M, all arms have at least M_η pulls. By Lemma G.1, $M_\eta \geq n_\star(\overline{\eta})$, so $t_{\mathrm{sep}}(\overline{\eta}) \leq \tau_M$.

Define the good event

$$\mathcal{E}_{\overline{\eta}} := \mathcal{G}_{\overline{\eta}} \cap \mathcal{K}_{\overline{\eta}}.$$

We use the filtration $\mathcal{F}_s$ defined in Appendix F.1. All martingale statements below are under $\mathbb{P}_\mu$, without conditioning on the good event. For $s = K, \ldots, T_\eta - 1$, define the martingale differences

$$Y_{s+1} := X_{s+1} - \mathbb{E}_\mu[X_{s+1} \mid \mathcal{F}_s].$$

Then $\mathbb{E}_\mu[Y_{s+1} \mid \mathcal{F}_s] = 0$ and $|Y_{s+1}| \leq 1$.

We first show that the conditional probability of progress is bounded below on the good event. Fix an integer $K \leq s < T_\eta \wedge \tau_M \wedge t_{\mathrm{sep}}(\overline{\eta})$ and suppose that $\mathcal{E}_{\overline{\eta}}$ holds. Since $s < t_{\mathrm{sep}}(\overline{\eta})$, at least one arm is separation-deficient. If some separation-deficient arm is not the current leader B_{s+1}, then Lemma G.1 implies that the challenger selected by the unguarded algorithm has count $N_{s,C_{s+1}} < M_\eta$. Thus $X_{s+1} = 1$ whenever the challenger is pulled. If no separation-deficient nonleader exists, the current leader is the only separation-deficient arm. In that case $N_{s,B_{s+1}} < n_\star(\overline{\eta}) \leq M_\eta$, so $X_{s+1} = 1$ whenever the leader is pulled. In both cases,

$$\mathbb{E}_\mu[X_{s+1} \mid \mathcal{F}_s] \geq p \quad \text{on } \mathcal{E}_{\overline{\eta}} \cap \{s < T_\eta \wedge \tau_M \wedge t_{\mathrm{sep}}(\overline{\eta})\}. \tag{G.1}$$

On $\mathcal{E}_{\overline{\eta}} \cap \{t_{\mathrm{sep}}(\overline{\eta}) > T_\eta\}$, one also has $\tau_M > T_\eta$, because $\tau_M \leq T_\eta$ would imply $t_{\mathrm{sep}}(\overline{\eta}) \leq \tau_M \leq T_\eta$. Therefore (G.1) applies for every $s = K, \ldots, T_\eta - 1$. Moreover, $\mathcal{D}_{T_\eta} > 0$, so

$$\sum_{s=K}^{T_\eta - 1} X_{s+1} = \mathcal{D}_K - \mathcal{D}_{T_\eta} < \mathcal{D}_K = K(M_\eta - 1).$$

Consequently,

$$\mathcal{E}_{\overline{\eta}} \cap \{t_{\mathrm{sep}}(\overline{\eta}) > T_\eta\} \subseteq \left\{ \sum_{s=K}^{T_\eta - 1} Y_{s+1} \leq -a_\eta \right\}, \qquad a_\eta := p(T_\eta - K) - K(M_\eta - 1). \tag{G.2}$$

By the definition of T_η,

$$T_\eta - K \geq \frac{2K}{p}(M_\eta - 1) + \frac{16}{p^2} \log \frac{6}{\eta},$$

so $a_\eta \geq p(T_\eta - K)/2$. Azuma–Hoeffding gives

$$\mathbb{P}_\mu \left(\sum_{s=K}^{T_\eta - 1} Y_{s+1} \leq -a_\eta \right) \leq \exp\left(-\frac{a_\eta^2}{2(T_\eta - K)} \right) \leq \exp\left(-\frac{p^2(T_\eta - K)}{8} \right) \leq \eta/6.$$

Since $\mathbb{P}_\mu(\mathcal{E}_{\overline{\eta}}^c) \le 2\overline{\eta} = \eta/3$, (G.2) yields

$$\mathbb{P}_\mu\left(t_{\mathrm{sep}}(\eta/6) > T_\eta\right) \le \eta/2.$$

The polynomial form of T_η in Lemma G.1 gives the claimed separation-time bound.

To bound the stabilization time, set $\overline{\eta} := \eta/12$ and define the rescaled horizon $\widehat{T}_\eta := T_{6\overline{\eta}} = T_{\eta/2}$. Applying the preceding separation-time bound at confidence level $\eta/2$ gives

$$\mathbb{P}_\mu\left(t_{\mathrm{sep}}(\overline{\eta}) \le \widehat{T}_\eta\right) \ge 1 - \eta/4,$$

because $(\eta/2)/6 = \overline{\eta}$. Intersect this event with the KL confidence event $\mathcal{K}_{\overline{\eta}}$, which has probability at least $1 - \eta/12$ and is contained in $\mathcal{K}_{\overline{\eta}}^{\mathrm{sep}}$ by Appendix D.1, and with the leader-fraction event $\mathcal{A}_{\eta/4}$ of Proposition F.1, which has probability at least $1 - \eta/4$. On the intersection, Proposition 6.2 implies

$$B_{t+1} = i^\star \qquad \forall t \ge \widehat{T}_\eta.$$

Lemma F.2, applied at confidence level $\eta/4$, bounds the earliest stabilization time by

$$t_U \le C_\beta(U)\left(\widehat{T}_\eta + K + \log\frac{e(\widehat{T}_\eta + 1)^2}{\eta}\right).$$

Here $C_\beta(U)$ absorbs the factor 4 inside the logarithm and upward rounding of the resulting horizon to an integer. The failure probability of the intersection is at most

$$\eta/4 + \eta/12 + \eta/4 < \eta.$$

Since $\widehat{T}_\eta$ has the same polynomial form in $1/\eta$ after changing constants, the claimed upper-quantile bound follows. Letting $\eta \downarrow 0$ shows that this t_U is almost surely finite. □

Proof of Corollary 6.4. Combine Proposition 6.3 with Theorem 5.2. The upper-quantile bound for t_U is polynomial in $1/\eta$, which yields the claimed high-probability stopping bound. This upper-quantile bound need not be summable, so it does not by itself yield an expected stopping-time bound. □

G.3 Finite-Mean Sufficient Exploration

Finite-mean sufficient exploration yields a summable upper-quantile bound for the stabilization time, allowing us to apply the conditional expectation bound.

Proof of Theorem 6.5. We construct an upper-quantile bound for stabilization from quantiles of T_{se} and verify dyadic summability before applying Theorem 5.2. Fix the compact interval $U \Subset (0,1)$ from the theorem, with β in its interior.

Let

$$q_{\mathrm{se}}(\xi) := \inf\{x \ge 0 : \mathbb{P}_\mu(T_{\mathrm{se}} \le x) \ge 1 - \xi\}, \qquad \xi \in (0,1],$$

be the upper-tail quantile of T_{se}. By Lemma F.5,

$$\sum_{k\ge 0} 2^{-k} q_{\mathrm{se}}(2^{-k}) < \infty.$$

Recall the separation scale from Appendix E.5: for $\xi \in (0, 1/2]$,

$$R_\xi := \log\frac{K}{\xi} + \log\log\frac{e}{\xi} + 1.$$

Lemma E.5 gives $n_\star(\xi) \le A_\mu R_\xi$. For $\eta \in (0, 1/4]$, define

$$L_{\mathrm{se}}(\eta) := q_{\mathrm{se}}(\eta/3) + \left\lceil \frac{K}{c_{\mathrm{se}}^2} n_\star(\eta/3)^2 \right\rceil + 1.$$

The $+1$ slack in this horizon ensures

$$\lfloor L_{\mathrm{se}}(\eta) \rfloor \ge q_{\mathrm{se}}(\eta/3), \qquad \lfloor L_{\mathrm{se}}(\eta) \rfloor \ge \left\lceil \frac{K}{c_{\mathrm{se}}^2} n_\star(\eta/3)^2 \right\rceil.$$

On the event $\{T_{\mathrm{se}} \le q_{\mathrm{se}}(\eta/3)\}$, every integer $t \ge \lfloor L_{\mathrm{se}}(\eta) \rfloor$ satisfies $t \ge T_{\mathrm{se}}$ and

$$\min_i N_{t,i} \ge c_{\mathrm{se}} \sqrt{t/K} \ge n_\star(\eta/3).$$

Hence $t_{\mathrm{sep}}(\eta/3) \le \lfloor L_{\mathrm{se}}(\eta) \rfloor \le L_{\mathrm{se}}(\eta)$. Intersecting this event with the KL confidence event $\mathcal{K}_{\eta/3}$, which is contained in $\mathcal{K}^{\mathrm{sep}}_{\eta/3}$ by Appendix D.1, Proposition 6.2 implies

$$B_{t+1} = i^\star \qquad \forall t \ge \lfloor L_{\mathrm{se}}(\eta) \rfloor .$$

Now also intersect with the event $\mathcal{A}_{\eta/3}$ of Proposition F.1. Since the rule is unguarded, the guard term in Lemma F.2 is absent. Applying the lemma with the integer leader time $\lfloor L_{\mathrm{se}}(\eta) \rfloor$ and its upper bound $L_{\mathrm{se}}(\eta)$ gives

$$t_U \le u_U(\eta) := C_\beta(U) \left(L_{\mathrm{se}}(\eta) + \log \frac{e(L_{\mathrm{se}}(\eta)+1)^2}{\eta/3} \right).$$

The intersection of the three events has probability at least $1-\eta$. Thus u_U is an upper-quantile bound on $(0, 1/4]$; extend it to $(0,1]$ by setting $u_U(\eta) = u_U(1/4)$ for $\eta > 1/4$. As above, upward rounding is absorbed into $C_\beta(U)$, and letting $\eta \downarrow 0$ shows that t_U is almost surely finite.

It remains to check dyadic summability. Since $n_\star(\eta) \le A_\mu R_\eta$,

$$L_{\mathrm{se}}(\eta) \le q_{\mathrm{se}}(\eta/3) + C_{\mu, c_{\mathrm{se}}} R^2_{\eta/3}.$$

For dyadic levels $\eta_k = 2^{-k} \le 1/4$, monotonicity of q_{se} gives $q_{\mathrm{se}}(2^{-k}/3) \le q_{\mathrm{se}}(2^{-k-2})$, and therefore

$$\sum_{k \ge 0} 2^{-k} q_{\mathrm{se}}(2^{-k}/3) < \infty.$$

The dyadic sums of $R^2_{2^{-k}/3}$ and $\log(e/2^{-k})$ are finite. The finitely many larger error levels do not affect summability. Finally,

$$\log \frac{e(L_{\mathrm{se}}(\eta)+1)^2}{\eta/3} \le C \left(L_{\mathrm{se}}(\eta) + \log(e/\eta) + 1 \right),$$

for a finite constant C independent of η, using $\log(1+x) \le x$. Hence

$$\sum_{k \ge 0} 2^{-k} u_U(2^{-k}) < \infty.$$

For the stopping-bound coefficients, repeat this construction with the larger compact interval $U^+ := [(\inf U)/2, (1 + \sup U)/2]$. For every $\beta \in U$, the symmetric interval centered at β with radius $\frac{1}{2}\min\{\inf U, 1 - \sup U\}$ lies in U^+. Thus Lemma F.2 permits an entrance coefficient depending only on U, since the coin coefficient is numerical and $1 + \beta \le 2$. The expectation part of Theorem 5.2, applied on U^+, yields the claimed coefficients $C_{\mu,k}(U)$, with the permitted dependence on c_{se} and the quantiles of T_{se}. □

G.4 Pairwise-Distinct Means

For pairwise-distinct means, the cited finite-mean sufficient-exploration result supplies the premise of Theorem 6.5 for the rule in Corollary 6.6.

Proof of Corollary 6.6. We apply the cited sufficient-exploration result and then Theorem 6.5. In Jourdan et al. [9, Lemmas 7, 17, and 21, Appendices C.3 and D.1], Lemma 17 verifies the leader property for the empirical leader, and Lemma 21 verifies the challenger property for the TCI index. Bernoulli rewards belong to their bounded-reward model, and pairwise-distinct means satisfy their Assumption 2. With tie-breaking as in Corollary 6.6, Lemma 7 therefore gives an almost surely finite random time T_{se} with $\mathbb{E}_\mu[T_{\mathrm{se}}] < \infty$ such that

$$\min_{i \in [K]} N_{t,i} \ge \sqrt{t/K} \qquad \forall t \ge T_{\mathrm{se}}.$$

Thus Theorem 6.5 applies with $c_{\mathrm{se}} = 1$ to the rule in that corollary. □

G.5 Guarded Expected Stopping-Time Bound

We prove the exploration and stabilization bounds below, then derive the expected stopping-time bound.

Proposition 6.7 (restated). Assume the exploration guard in Algorithm 1 is enabled. Then for all $t \geq K$, $\min_{i\in[K]} N_{t,i} \geq \left\lfloor \sqrt{t/K} \right\rfloor - 1$. Consequently, for every $\eta \in (0,1)$, $t_{\mathrm{sep}}(\eta) \leq K\,(n_\star(\eta)+1)^2$. For every compact interval $U \Subset (0,1)$ containing β in its interior, there exists a constant $C_{\mu,\beta}(U) < \infty$ such that the earliest U-stabilization time t_U satisfies, for every $\eta \in (0, 1/4]$,

$$\mathbb{P}_\mu\big(t_U \leq C_{\mu,\beta}(U)(\log(e/\eta))^2\big) \geq 1-\eta.$$

Proof of Proposition 6.7. We prove the count bound by induction over square-root epochs, then combine separation and leader-fraction control to bound stabilization. Let $m \geq 1$ be an integer and define epoch boundaries $t_m := Km^2$. By initialization, $N_{t_1,i} = 1$ for all arms. Assume inductively that $N_{t_m,i} \geq m$ for all arms. For $t_m < t < t_{m+1}$, we have $m < \sqrt{t/K} < m+1$. Thus every round after the epoch's first round is guarded while any arm still has count m. Each guard pull raises a least-sampled arm's count from m to $m+1$, so at most K such pulls suffice. There are $K(2m+1)-1 \geq K$ rounds after the first, and hence all counts are at least $m+1$ at t_{m+1}. This proves the epoch invariant and the stated lower bound. Evaluating that bound at $K(n_\star(\eta)+1)^2$ gives the claimed bound on $t_{\mathrm{sep}}(\eta)$.

For stabilization, fix U as in the proposition and $\eta \in (0,1/4]$, and set $L := K(n_\star(\eta/2)+1)^2$. On $\mathcal{K}_{\eta/2} \cap \mathcal{A}_{\eta/2}$, an event of probability at least $1-\eta$, the inclusion $\mathcal{K}_{\eta/2} \subseteq \mathcal{K}^{\mathrm{sep}}_{\eta/2}$ and Proposition 6.2 give permanent leader correctness from L. Apply Lemma F.2 to a symmetric compact subinterval of U centered at β, with a constant $C_\beta(U) \geq 1$. Since the following integer horizon is at least L, it bounds the same earliest stabilization time:

$$t_U \leq \left\lceil C_\beta(U)\left(L + K + \log\frac{e(L+1)^2}{\eta/2}\right)\right\rceil.$$

Write $x := \log(e/\eta) \geq 1$. Lemma E.5 and $\log(x+\log 2) \leq x$ give $n_\star(\eta/2) \leq S_\mu x$, where $S_\mu := \max\{1, A_\mu(2+\log(2K))\}$. Thus $L \leq D_\mu x^2$ with $D_\mu := K(S_\mu+1)^2$, and

$$\log\frac{e(L+1)^2}{\eta/2} \leq x + \log 2 + 2\log(D_\mu+1) + 4\log x \leq (5+\log 2+2\log(D_\mu+1))\,x^2.$$

The ceiling adds at most $1 \leq x^2$, so the constant

$$C_{\mu,\beta}(U) := C_\beta(U)\,(D_\mu + K + 5 + \log 2 + 2\log(D_\mu+1)) + 1$$

depends only on μ, β, U (with K included in the instance μ) and satisfies

$$\mathbb{P}_\mu\big(t_U \leq C_{\mu,\beta}(U)(\log(e/\eta))^2\big) \geq 1-\eta \qquad \forall \eta \in (0,1/4].$$

Letting $\eta \downarrow 0$ shows that t_U is almost surely finite. An upper-quantile bound on $(0,1]$ is $u_U(\eta) := C_{\mu,\beta}(U)(\log(e/\min\{\eta,1/4\}))^2$. Its dyadic series is finite because $u_U(2^{-k}) = O_{\mu,\beta,U}((k+1)^2)$ and $\sum_{k\geq 0} 2^{-k}(k+1)^2 < \infty$. □

Proof of Theorem 6.8. By Proposition 6.7 and Lemma F.2, the guard yields a polylogarithmic upper-quantile bound for a U-stabilization time t_U. In particular, this bound satisfies both the probability requirement and the dyadic summability condition in Theorem 5.2. Applying the expectation part of Theorem 5.2 completes the proof. □

G.6 Limitation of One-Step Index Comparisons

We first determine the limiting population evidence as a provisional leader's count grows while a competitor's count stays fixed. We then use this limit to prove the one-step index-comparison limitation.

Lemma G.2 (Asymptotics with a fixed leader identity and growing count)**.** *Fix distinct arms $b, u \in [K]$ and a positive integer count m. The following map is nondecreasing for positive integer counts n:*

$$n \mapsto \inf_{x\in(0,1)} \{nd(\mu_b, x) + md(\mu_u, x)\}.$$

As $n \to \infty$, it converges to $md(\mu_u, \mu_b)$.

Proof. Monotonicity in n follows because the objective contains the nonnegative term $nd(\mu_b, x)$. Evaluating at $x = \mu_b$ gives the upper bound $md(\mu_u, \mu_b)$. Conversely, let $x_n := (n\mu_b + m\mu_u)/(n+m)$ be the exact pooled minimizer, as in Remark B.2. If $d(\mu_b, x_n)$ failed to converge to zero, then $nd(\mu_b, x_n)$ would be unbounded along a subsequence, contradicting the finite upper bound at $x = \mu_b$. Hence $x_n \to \mu_b$, and continuity gives the matching lower bound. □

Proof of Proposition 6.9. We compare the exact index of the equal-mean arm with the limiting index of the higher-mean arm as the provisional leader's count grows. Because $\mu_j = \mu_b = \lambda$, choosing $x = \lambda$ gives

$$\Phi_j^{(b)}(n, M(\eta)) = \log M(\eta),$$

for every n. Fix any $a < Cd(\mu_u, \lambda)$ and choose $\varepsilon > 0$ so that $a < Cd(\mu_u, \lambda) - 2\varepsilon$. By Lemma G.2, for every sufficiently small η there exists n_η such that for all $n \geq n_\eta$,

$$\inf_{x\in(0,1)} \{nd(\lambda, x) + m_u(\eta)d(\mu_u, x)\} \geq m_u(\eta)d(\mu_u, \lambda) - \varepsilon \log(e/\eta).$$

Since $m_u(\eta) \geq C \log(e/\eta)$ for small η,

$$\Phi_u^{(b)}(n, m_u(\eta)) \geq (Cd(\mu_u, \lambda) - \varepsilon) \log(e/\eta) + \log m_u(\eta).$$

By assumption, for all sufficiently small η and all $n \geq n_\eta$,

$$\log M(\eta) = \Phi_j^{(b)}(n, M(\eta)) \geq \Phi_u^{(b)}(n, m_u(\eta)) \geq (Cd(\mu_u, \lambda) - \varepsilon) \log(e/\eta).$$

Hence

$$M(\eta) \geq \eta^{-(Cd(\mu_u,\lambda)-\varepsilon)}.$$

Because $a < Cd(\mu_u, \lambda) - 2\varepsilon$, this implies $M(\eta) \geq \eta^{-a}$ for all sufficiently small η. □

H Experiment Details

This section describes the experiments in Section 7, including the grids, recorded statistics, and supplementary diagnostics. All experiments use synthetic Bernoulli rewards and no external datasets.

H.1 Algorithms and Implementation

All variants use the same GLR stopping rule, the threshold $c(t, \delta) = \log(1/\delta) + 2\log(1 + t/2) + 2 + \log(K - 1)$, and smallest-index tie-breaking. We compare three sampling variants, each initialized by one pull per arm:

- **Unguarded penalized**: the unguarded algorithm, with challenger $C_{t+1} \in \arg\min_{j\in[K]\setminus\{B_{t+1}\}} \{W_t(B_{t+1}, j) + \log N_{t,j}\}$.
- **Guarded penalized**: the same rule, except that if $\min_{i\in[K]} N_{t,i} < \sqrt{t/K}$, the algorithm pulls a least-sampled arm.
- **Unpenalized ablation**: the unguarded algorithm whose challenger minimizes $W_t(B_{t+1}, j)$ instead of $W_t(B_{t+1}, j) + \log N_{t,j}$; the stopping rule, threshold, initialization, and tie-breaking are unchanged.

Challenger comparisons treat index values within an absolute tolerance of 10^{-12} of the current minimum as ties, so the smaller index is retained. The implementation evaluates the Bernoulli two-sample GLR at the pooled mean, uses Numba-accelerated simulation kernels, and records every run in a raw CSV with the recorded time, final recommendation, cap indicator, guard count, final arm counts, and optimal population weights.

Experiment	Grid	Runs
Main evaluation	3 instances $\times$ 3 values of $\beta \times 5$ values of $\delta \times 2$ variants $\times$ 200 replications	18,000
β-sweep	2 instances $\times$ 7 values of $\beta \times 200$ replications	2,800
Stress survival	2 penalized variants $\times$ 600 replications	1,200
Penalty ablation	1 unpenalized variant $\times$ 500 replications	500
Guard overhead	$K \in \{4, 8, 12, 16\}$ with 200, 200, 100, 50 replications	550

Table 1: Experiment grids. The three main evaluation instances are SEP5, CLUSTER-CLOSE5, and TIE-CLOSE4. The evaluation confidence levels are $10^{-4}, 10^{-5}, 10^{-6}, 10^{-7}, 10^{-8}$.

Code and reproduction. The supplementary code archive `eb_tci_neurips_submission_code.zip` contains the directory `eb_tci_neurips_submission_code/`; all paths below are relative to this directory. The setup instructions are in `README.md`. The simulation kernels are in `src/eb_tci/sim_numba.py`, and the per-run data are in `results/raw_full.csv` and `results/raw_oracle_tracking_full.csv`. From the package root, the full experiments and plots are reproduced by:

```
python scripts/run_experiments.py --profile full --out results
python scripts/run_oracle_tracking_baseline.py --profile full --out results
python scripts/make_plots.py
```

Population benchmark and threshold-crossing predictor. For each instance and each β, we compute $T^\star_\beta(\mu)$ from the Bernoulli equalization system in Proposition C.3. For fixed $\alpha \in (0, 1)$ and suboptimal arm $j \in [K] \setminus \{i^\star\}$, recall

$$I^\alpha_j(\omega) := \inf_{x \in (0,1)} \{\alpha d(\mu_{i^\star}, x) + \omega d(\mu_j, x)\}, \qquad \omega > 0.$$

The optimizer is characterized by a scalar λ_α such that

$$I^\alpha_j(w^\alpha_j(\mu)) = \lambda_\alpha \quad \forall j \in [K] \setminus \{i^\star\}, \qquad \sum_{j \neq i^\star} w^\alpha_j(\mu) = 1 - \alpha,$$

and $T^\star_\alpha(\mu) = 1/\lambda_\alpha$. Numerically, we solve this by bisection on λ_α, evaluating each inverse $(I^\alpha_j)^{-1}$ by bisection of the strictly increasing rate I^α_j. The inverse-map domain is $0 < \lambda_\alpha < \alpha \min_{j \neq i^\star} d(\mu_{i^\star}, \mu_j)$, as established in Proposition C.3.

The implementation in `src/eb_tci/oracle.py`, function `solve_t_star`, uses the outer bracket $[0, (1 - 10^{-12})\alpha \min_{j \neq i^\star} d(\mu_{i^\star}, \mu_j)]$. Each inverse search starts with the weight bracket $[0, 1]$ and doubles its upper endpoint until the rate reaches the trial value of λ_α; a safeguard ends expansion if that endpoint exceeds 10^{12}. The outer and inner searches allow at most 140 and 120 bisection iterations, respectively. Both stop when the bracket width is at most 10^{-12} times the larger of one and the upper endpoint. The returned suboptimal weights are rescaled to sum to $1 - \alpha$, and the information rate is recomputed as the minimum of their pairwise rates. No separate residual test is performed. The threshold-crossing predictor retains the exact finite-confidence threshold and minimizes over integer times:

$$t_{\mathrm{pred}}(\mu, \beta, \delta) := \min\{t \geq K : t/T^\star_\beta(\mu) > c(t, \delta)\}.$$

Experiment grids and replication counts. The five grids in Table 1 comprise 23,050 independent simulation runs.

The main evaluation instances are

$$\text{SEP5} = [0.70, 0.50, 0.40, 0.30, 0.20], \qquad \text{CLUSTER-CLOSE5} = [0.70, 0.64, 0.63, 0.62, 0.61],$$

and

$$\text{TIE-CLOSE4} = [0.60, 0.56, 0.56, 0.56].$$

The β-sweep uses CLUSTER-CLOSE5 and TIE-CLOSE4, $\beta \in \{0.2, 0.3, \ldots, 0.8\}$, $\delta = 10^{-6}$, and the guarded penalized variant. The guard-overhead family uses $[0.60, 0.56, \ldots, 0.56]$ with $K \in \{4, 8, 12, 16\}$, $\beta = 0.5$, and $\delta = 10^{-6}$. The stress experiments use TIE-CLOSE4, $\beta = 0.5$, and $\delta = 10^{-8}$.

Variant	Settings	Median ARE	Mean ARE	Maximum ARE
Guarded penalized	45	1.39%	1.53%	3.86%
Unguarded penalized	45	1.60%	1.66%	4.20%

Table 2: Absolute relative error between empirical mean stopping time and t_{pred} over the complete evaluation grid. No penalized evaluation setting hits the simulation cap or yields a Monte Carlo misidentification.

Randomness, caps, and uncertainty summaries. The seed-to-setting map is defined by `build_specs` in `scripts/run_experiments.py`. The full profile starts at 20260424 and enumerates the main evaluation, β-sweep, guard overhead, stress survival, and penalty ablation in that order. Within the main evaluation, the nested loops run over instances (SEP5, CLUSTER-CLOSE5, TIE-CLOSE4), β $(0.3, 0.5, 0.7)$, δ $(10^{-4}, \ldots, 10^{-8})$, and variant (unguarded, then guarded), with replication as the innermost index. The β-sweep visits CLUSTER-CLOSE5, then TIE-CLOSE4, with β increasing from 0.2 to 0.8. Guard overhead visits $K = 4, 8, 12, 16$; stress survival visits the unguarded, then guarded variant. Each setting uses consecutive seeds and advances the seed counter by its replication count, except that every guard-overhead setting reserves 200 seeds, including those with 100 or 50 replications. The separately reported population-reference runs use 200 consecutive seeds starting at 20268000 for CLUSTER-CLOSE5 and 20269000 for TIE-CLOSE4, as specified in `src/eb_tci/oracle_tracking.py`. Both raw CSV files record the seed of each run.

All runs use a fail-safe cap of 700,000 rounds. Each run records the smaller of its stopping time τ_δ and the cap. A cap hit therefore contributes 700,000 to recorded-time summaries, while its stopping time remains censored. Only the unpenalized ablation hits the cap; its cap-hit rate is reported explicitly.

For a scalar statistic Z, let $Z^{(r)}$ be its value in replication $r \in \{1, \ldots, R\}$, where $R \geq 2$. The empirical mean is $\overline{Z} := R^{-1} \sum_{r=1}^{R} Z^{(r)}$, and its empirical standard error over independent runs is

$$\widehat{\mathrm{se}}(\overline{Z}) = \left(\frac{1}{R(R-1)} \sum_{r=1}^{R} (Z^{(r)} - \overline{Z})^2 \right)^{1/2}.$$

Bootstrap percentile 95% confidence intervals use 1000 resamples with an independent bootstrap seed 314159. The finite-confidence scaling figure uses bootstrap percentile bands for the empirical mean; the β-sensitivity figure uses ± 1.96 empirical standard errors.

Metrics. The main predictor-agreement metric is the absolute relative error

$$\mathrm{ARE}(\mu, \beta, \delta) := \left| \frac{\overline{\tau}_\delta}{t_{\mathrm{pred}}(\mu, \beta, \delta)} - 1 \right|,$$

where $\overline{\tau}_\delta$ is the Monte Carlo mean stopping time for the uncapped penalized evaluation runs. We also report the median, 10th, 90th, and 95th percentiles of the recorded times, the empirical error rate, the cap-hit rate, and final allocation diagnostics. For the guarded algorithm, recall the guard count at stopping:

$$G_{\tau_\delta}^{\mathrm{tot}} = \sum_{t=K}^{\tau_\delta - 1} \mathbf{1}\{ \min_{i \in [K]} N_{t,i} < \sqrt{t/K} \}.$$

We report the guarded-pull fraction $G_{\tau_\delta}^{\mathrm{tot}} / \tau_\delta$ and the normalization $G_{\tau_\delta}^{\mathrm{tot}} / \sqrt{K \tau_\delta}$ used in the theoretical guard-count bound. No guarded run hits the cap.

Aggregate results. Table 2 reports predictor agreement over the full $3 \times 3 \times 5$ grid. Table 3 summarizes the penalized stress tests and unpenalized ablation. Table 4 gives the equal-close guard-overhead family.

β-sensitivity numerical table. Table 5 gives the numerical values underlying the right-hand panel of Figure 2. The empirical mean tracks the nonmonotone dependence of $T_\beta^\star(\mu)$ on the leader sampling probability. The relative error is signed and equals $\overline{\tau}_\delta / t_{\mathrm{pred}} - 1$.

Variant	Replications	Mean	Median	95th percentile
Unguarded penalized	600	105,593	104,918	134,140
Guarded penalized	600	104,022	103,743	130,792
Unpenalized ablation	500	188,036	108,036	700,000

Table 3: Stress-test recorded-time summaries on TIE-CLOSE4 with $\beta = 0.5$ and $\delta = 10^{-8}$. The unpenalized ablation hits the 700,000-round cap in 12.8% of runs; those runs contribute the cap value to the mean and quantiles.

Instance	K	Replications	Mean τ_δ	Mean $G^{\mathrm{tot}}_{\tau_\delta}/\tau_\delta$	95th percentile $G^{\mathrm{tot}}_{\tau_\delta}/\sqrt{K\tau_\delta}$
EQUAL-CLOSE4	4	200	92,922	0.0058%	1.80%
EQUAL-CLOSE8	8	200	196,872	0.0065%	1.80%
EQUAL-CLOSE12	12	100	301,114	0.0069%	1.89%
EQUAL-CLOSE16	16	50	417,489	0.0062%	1.41%

Table 4: Guard overhead on the equal-close family with $\beta = 0.5$ and $\delta = 10^{-6}$. Both the mean guarded-pull fraction and the 95th percentile of the normalized guard count are shown as percentages.

H.2 Idealized Population-Allocation Reference

We use an idealized population-allocation reference to compare stopping-time scales. From the true means, it computes a full K-arm allocation

$$w^\star(\mu) \in \operatorname*{arg\,max}_{\substack{w\in(0,1)^K \\ \sum_{i=1}^K w_i = 1}} \min_{j\in[K]\setminus\{i^\star\}} I_j^{w_{i^\star}}(w_j).$$

The implementation in `src/eb_tci/oracle_tracking.py` initializes with one pull per arm and then deterministically C-tracks the fixed allocation: after checking the stopping rule at time t, it chooses

$$I_{t+1} \in \operatorname*{arg\,max}_{i\in[K]} \{(t+1)w_i^\star(\mu) - N_{t,i}\},$$

with smallest-index tie-breaking. It uses no additional projection or forced-exploration step. The GLR statistic and threshold $c(t,\delta)$ are the same as in the β-EB-TCI runs. This construction is motivated by the allocation principle underlying Track-and-Stop [4], but it is not an empirical Track-and-Stop implementation: the reference is given population information unavailable to a learning algorithm and never recomputes an empirical optimal allocation.

For exact reproduction of the reported reference runs, the default command uses the stored weights in `PRECOMPUTED_GLOBAL_ALLOCATIONS` in `src/eb_tci/oracle_tracking.py`. Recomputing these weights can change the deterministic pull order through differences in their last digits. The optional `--recompute-allocation` flag calls `solve_global_optimal_allocation`: it evaluates the fixed-fraction solver at 181 equally spaced values of α in $[0.05, 0.95]$, selects the first maximizing grid point, and refines between the grid points at most two positions on either side using 40 golden-section iterations. Among the two final trial points, it returns the one with the larger computed information rate, choosing the smaller α in a tie. Each fixed-fraction evaluation uses the bisection settings described above.

Table 6 compares this reference with the closest guarded β-EB-TCI point from the β-sweep. We use $\beta = 0.4$ because it is the closest grid value to the optimal best-arm fraction $\alpha^\star := w^\star_{i^\star}(\mu)$. The optimal fraction is $\alpha^\star = 0.357$ on CLUSTER-CLOSE5 and $\alpha^\star = 0.367$ on TIE-CLOSE4. Because the methods use different best-arm fractions and the reference knows μ, the table is a diagnostic comparison of finite-sample scales, not evidence for an algorithm ranking.

H.3 Supplementary Figures

Figure 3 reports the guard diagnostics that are summarized in the main text. The left panel compares stopping-time survival curves on TIE-CLOSE4. The right panel reports the normalized guard overhead across the equal-close family.

Instance	β	Mean τ_δ	t_{pred}	Signed relative error
CLUSTER-CLOSE5	0.2	34,151	33,901	0.74%
CLUSTER-CLOSE5	0.3	28,757	29,267	−1.74%
CLUSTER-CLOSE5	0.4	29,141	29,001	0.48%
CLUSTER-CLOSE5	0.5	31,390	31,334	0.18%
CLUSTER-CLOSE5	0.6	36,063	36,505	−1.21%
CLUSTER-CLOSE5	0.7	45,304	46,458	−2.48%
CLUSTER-CLOSE5	0.8	65,735	67,856	−3.13%
TIE-CLOSE4	0.2	102,704	102,907	−0.20%
TIE-CLOSE4	0.3	88,998	88,819	0.20%
TIE-CLOSE4	0.4	85,348	87,269	−2.20%
TIE-CLOSE4	0.5	92,510	93,342	−0.89%
TIE-CLOSE4	0.6	105,470	107,680	−2.05%
TIE-CLOSE4	0.7	134,493	135,786	−0.95%
TIE-CLOSE4	0.8	191,672	196,628	−2.52%

Table 5: β-sensitivity with $\delta = 10^{-6}$, the guarded penalized rule, and 200 replications per point.

Instance	Method	β or $\alpha^\star$	Replications	Mean τ_δ	SE	95th percentile
CLUSTER-CLOSE5	Guarded β-EB-TCI	0.400	200	29,141	401	38,530
CLUSTER-CLOSE5	Idealized allocation tracking	0.357	200	34,808	411	43,978
TIE-CLOSE4	Guarded β-EB-TCI	0.400	200	85,348	1,107	111,139
TIE-CLOSE4	Idealized allocation tracking	0.367	200	102,085	1,373	132,011

Table 6: Idealized population-allocation reference at $\delta = 10^{-6}$. The idealized reference rows use the population-optimal Bernoulli allocation and the same GLR threshold as the β-EB-TCI experiments. SE is the empirical standard error of the mean. No run in this table hits the cap or yields a Monte Carlo misidentification.

Figure 4 contains the remaining diagnostics. The penalty ablation shows the effect of removing $\log N_{t,j}$ from the challenger index. Its empirical survival curves summarize recorded times: cap hits contribute the cap value, so the curves do not estimate the uncensored stopping-time tail at or beyond the cap.

For the allocation panel, define the full count vector $N_t := (N_{t,i})_{i\in[K]}$ and the full fixed-β target $v^\beta(\mu)$ by

$$v^\beta_{i^\star}(\mu) := \beta, \qquad v^\beta_j(\mu) := w^\beta_j(\mu) \quad (j \in [K] \setminus \{i^\star\}).$$

The left allocation panel compares mean final sampling fractions with this target; the right panel reports the per-run distance $\|N_{\tau_\delta}/\tau_\delta - v^\beta(\mu)\|_1$. These are measurements at stopping; no per-run stabilization criterion is reported.

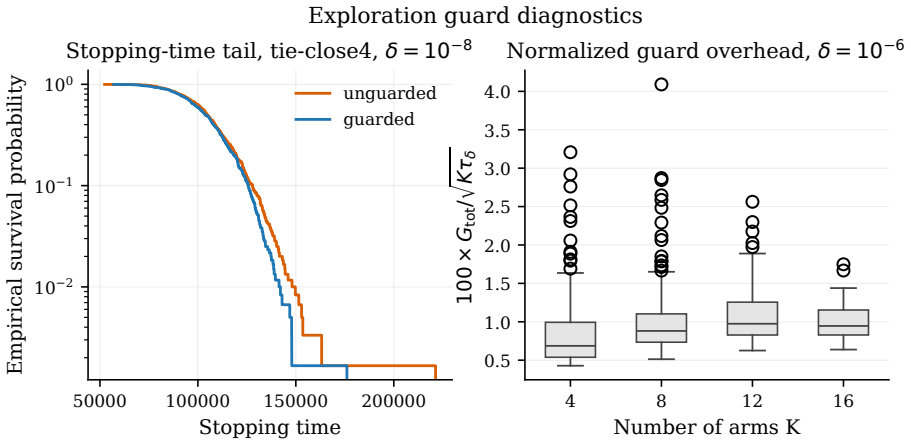
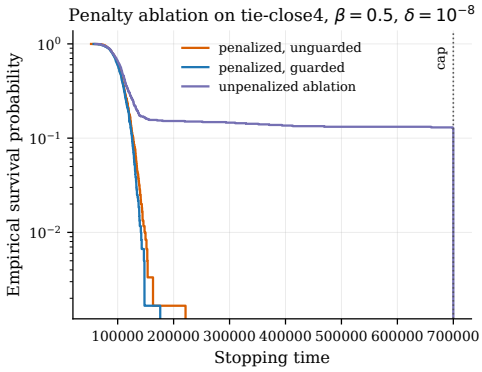
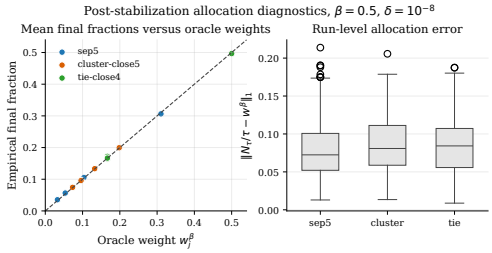


Figure 3: Exploration guard diagnostics. Left: empirical survival curves on TIE-CLOSE4 with $\beta = 0.5$ and $\delta = 10^{-8}$, using 600 replications per penalized variant. Right: normalized guard overhead on the equal-close family with $K \in \{4, 8, 12, 16\}$ and $\delta = 10^{-6}$.



(a) Penalty ablation on the equal-close stress instance.



(b) Final allocation diagnostics.

Figure 4: Additional diagnostics. Left: the unpenalized challenger removes $\log N_{t,j}$ but leaves the stopping rule and all other details unchanged. Right: final sampling fractions for the guarded rule with $\beta = 0.5$ and $\delta = 10^{-8}$, compared with the optimal population weights.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes].

   Justification: The abstract and introduction distinguish the stabilized-regime result (Theorem 5.2), the unguarded high-probability guarantee for every unique-best Bernoulli instance (Corollary 6.4), the unguarded expected-time results under additional exploration and tie-breaking assumptions (Theorem 6.5 and Corollary 6.6), and the guarded expected-time guarantee (Proposition 6.7 and Theorem 6.8). They state the fixed-β scope and unresolved unguarded case with tied suboptimal means; Section 7 provides synthetic finite-confidence diagnostics.

   Guidelines:

   - The answer [N/A] means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A [No] or [N/A] answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes].

   Justification: The analysis assumes interior Bernoulli means, a unique best arm, and fixed β. The sharp unguarded expected-time guarantee additionally requires finite-mean sufficient exploration; its pairwise-distinct corollary uses the tie-breaking convention of Jourdan et al. [9], not the paper's deterministic default. Section 6 and Section 8 identify the unresolved unguarded case with tied suboptimal means and extensions beyond Bernoulli rewards or fixed β. Section 7 confines empirical conclusions to the tested synthetic instances and does not claim superiority over other algorithms.

   Guidelines:

   - The answer [N/A] means that the paper has no limitation while the answer [No] means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.

- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes].

   Justification: The results specify the Bernoulli and unique-best-arm assumptions, the compact stabilization interval $U \Subset (0, 1)$, and the extra sufficient-exploration or tie-breaking assumptions where relevant. The optional guard is defined explicitly in Algorithm 1. Appendices B–G provide the derivations and proofs for the stopping rule, allocation benchmark, concentration, stabilized-regime guarantees, and whole-run bounds; Appendix A identifies the results imported from prior work.

   Guidelines:

   - The answer [N/A] means that the paper does not include theoretical results.
   - All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
   - All assumptions should be clearly stated or referenced in the statement of any theorems.
   - The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
   - Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
   - Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes].

   Justification: Section 3 specifies the sampling and stopping rules. Section 7 and Appendix H document the Bernoulli instances, parameter grids, seeds, replication counts, stopping cap, oracle calculation, numerical tolerances, uncertainty summaries, and metrics. Appendix H.2 clarifies that the population-allocation reference is not empirical Track-and-Stop; Appendix H.1 gives executable reproduction commands.

   Guidelines:

   - The answer [N/A] means that the paper does not include experiments.
   - If the paper includes experiments, a [No] answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.

- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

   Answer: [Yes].

   Justification: The anonymized supplementary archive `eb_tci_neurips_submission_code.zip` has been submitted with the paper. As documented in Appendix H.1, it contains the README and dependency specification, simulation and oracle implementations, raw result CSVs, and scripts reproducing the experiment grids, idealized allocation reference, figures, and tables. Appendix H.1 also provides reproduction commands. All rewards are generated synthetically, so no external dataset is required.

   Guidelines:

   - The answer [N/A] means that paper does not include experiments requiring code.
   - Please see the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - While we encourage the release of code and data, we understand that this might not be possible, so [No] is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
   - The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (`https://neurips.cc/public/guides/CodeSubmissionPolicy`) for more details.
   - The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
   - The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
   - At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
   - Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer) necessary to understand the results?

   Answer: [Yes].

   Justification: There are no training/test splits: the evaluations are Monte Carlo simulations of Bernoulli bandit instances. Section 7 and Appendix H.1 specify the mean vectors, β and

δ grids, sampling variants, deterministic tie-breaking and numerical tie tolerance, seeds, replication counts, GLR stopping threshold, 700,000-round cap, and the bisection procedure and tolerances used to compute $T^{\star}_{\beta}(\mu)$.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes].

Justification: Figure 1 shows bootstrap percentile 95% confidence bands for empirical mean stopping times from 200 independent replications per point. The β-sensitivity panel of Figure 2 uses ± 1.96 empirical standard errors. Appendix H.1 defines the standard-error calculation, 1,000 bootstrap resamples and their seed, and the replication counts; Appendix H (Table 6) also reports standard errors for the population-allocation reference.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The authors should answer [Yes] if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g., negative error rates).
- If error bars are reported in tables or plots, the authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [No].

Justification: Appendix H.1 documents the Numba-accelerated simulation implementation and replication counts, but the paper does not report the processor configuration, memory and storage requirements, per-experiment wall-clock runtimes, total compute consumption, or compute used for preliminary experiments. These resources cannot be quantified from the reported protocol alone.

Guidelines:

- The answer [N/A] means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.

- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics `https://neurips.cc/public/EthicsGuidelines`?

   Answer: [Yes].

   Justification: The study analyzes a theoretical fixed-confidence best-arm identification problem and evaluates algorithms using synthetic Bernoulli rewards (Sections 3 and 7). It involves no human participants, personal data, scraped datasets, or deployment of potentially high-risk models.

   Guidelines:

   - The answer [N/A] means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer [No], they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [N/A].

    Justification: The contribution is foundational theory for fixed-confidence best-arm identification, supported by synthetic simulations, rather than a deployed application. The paper does not identify a direct societal impact; possible benefits or harms would depend on downstream applications of bandit methods, which are not studied here.

    Guidelines:

    - The answer [N/A] means that there is no societal impact of the work performed.
    - If the authors answer [N/A] or [No], they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
    - The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate Deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
    - The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
    - If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pre-trained language models, image generators, or scraped datasets)?

Answer: [N/A].

Justification: The paper releases no high-misuse-risk dataset, pretrained or generative model, or scraped material. The only new asset discussed is code for reproducing synthetic Bernoulli simulations (Appendix H.1); special safeguards for high-risk model or dataset release are therefore not applicable.

Guidelines:

- The answer [N/A] means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [N/A].

Justification: The experiments use generated Bernoulli rewards and do not reuse external datasets, pretrained models, or third-party benchmark assets. Prior theoretical methods and algorithms are credited through the related-work discussion and references. General-purpose software dependencies serve as implementation tools, not as externally sourced research datasets or models.

Guidelines:

- The answer [N/A] means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [Yes].

Justification: The submitted supplementary reproduction-code package is documented in Appendix H.1, including its README, dependency specification, simulation kernels, oracle and predictor computation, experiment drivers, plotting scripts, and raw result files. It introduces no new dataset or pretrained model, and no human-derived data are involved.

Guidelines:

- The answer [N/A] means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [N/A].

    Justification: The research contains no crowdsourcing, human-subject experiment, user study, annotation task, or participant data; all observations used in the experiments are synthetic Bernoulli rewards.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [N/A].

    Justification: No human participants or crowdsourced workers are involved, so participant-risk disclosures and institutional review board (or equivalent) approval are not applicable.

    Guidelines:

    - The answer [N/A] means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

    Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does *not* impact the core methodology, scientific rigor, or originality of the research, declaration is not required.

    Answer: [N/A].

Justification: LLMs are not a component of the proposed algorithm, its proof methodology, or its experimental sampling rule. The described assistance was limited to writing and editing, and routine help with experiment scripts and reproduction materials; it is not an important, original, or non-standard methodological use requiring declaration under the NeurIPS 2026 policy.

Guidelines:

- The answer [N/A] means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy in the NeurIPS handbook for what should or should not be described.