Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

EXACT OPTIMAL FIXED WIDTH CONFIDENCE INTERVAL ESTIMATION FOR THE MEAN

Vikas Deep

Kellogg School of Management
Northwestern University
2211 Campus Dr
Evanston, IL 60208, USA

Achal Bassamboo

Kellogg School of Management
Northwestern University
2211 Campus Dr
Evanston, IL 60208, USA

Sandeep Juneja

School of Technology and Computer Science
Tata Institute of Fundamental Research
Dr Homi Bhabha Road, Navy Nagar
Colaba, Mumbai, Maharashtra 400005, INDIA

Assaf Zeevi

Graduate School of Buisness
Columbia University
665 W 130th St
New York, NY 10027, USA

ABSTRACT

We consider a classical problem in simulation/statistics - given i.i.d. samples of a rv, the goal is to arrive at a confidence interval (CI) of a pre-specified width ε, and with a coverage guarantee that the mean lies in the CI with probability at least $1-\delta$ for pre-specified $\delta \in (0,1)$. This problem has been well studied in an asymptotic regime as ε shrinks to zero. The novelty of our analysis is the derivation of the lower bound on the number of samples required by any algorithm to construct a CI of $\varepsilon-$width with *the coverage guarantee for fixed* $\varepsilon > 0$ and δ, and construction of an algorithm that, under mild assumptions, matches the lower bound. For simplicity, we present our results for rv belonging to a single parameter exponential family, and illustrate its efficacy through a numerical study.

1 INTRODUCTION AND LITERATURE REVIEW

The problem of constructing fixed-width CI of the mean of a distribution is well studied in the statistics and simulation literature. Most of the extant literature studies this problem with the fixed width ε and pre-specified confidence level $1-\delta$. In this paper, we aim to address three important elements absent from said literature.

a) Most of the existing algorithms have coverage guarantees only in an asymptotic regime, where the probability of the mean not lying in the CI tends to pre-specified δ when $\varepsilon \to 0$. However, this guarantee may not hold for fixed $\varepsilon > 0$.
b) Lower bound on the number of i.i.d. samples required for any algorithm to construct ε width CI, to achieve coverage guarantee for fixed $\varepsilon > 0$ and pre-specified δ is not developed in the literature. The asymptotic optimality of the algorithms presented in the literature in the regime where $\varepsilon \to 0$ is implicitly based on the fact that the estimator is asymptotically normally distributed. However, for fixed $\varepsilon > 0$, this clearly may not be true.
c) Typically the constructed CI's are symmetric around the sample mean. Such symmetric construction of CI stems from asymptotic normality. However, asymmetric CI of fixed width may be achieved it with fewer samples, leaving open the question of a more precise structure of the CI.

For $\delta \in (0,1)$, we address (a) above by presenting an algorithm that has guaranteed coverage for finite $\varepsilon > 0$, i.e., the probability of mean not lying in the CI is less than δ. Regarding (b), we derive a lower bound on the sample size of i.i.d. copies that are required by any algorithm that constructs CI of the desired width and has the above coverage guarantee. Regarding (c), we propose an algorithm that does not construct symmetric CI around the sample mean and has the coverage guarantee for fixed $\varepsilon > 0$ and pre-specified δ. Further, it matches the lower bound in an asymptotic regime where $\delta \to 0$. In fact, we prove that for most distributions any algorithm that constructs symmetric CI around the sample mean or in general any unbiased estimator would require a larger sample size for small δ. Our proposed algorithm, motivated by the lower bound, is based on the "plug-in" method used in the lower bound characterization.

For simplicity of analysis, we assume that the unknown distribution belongs to the single parameter exponential family. In our framework, we assume that the family of the distributions is known but the parameter is unknown and the mean of the distribution depends on the parameter. The class of single parameter exponential family distributions include the practically relevant Bernoulli distribution family as well as Gaussian distributions with known variance, Gamma distributions with known shape parameter and Poisson distributions. Our analysis can be extended to bounded random variables (see Section 6).

We also present the numerical study comparing our proposed algorithm with two algorithms in the existing literature, one proposed in (Chow and Robbins 1965) and the other proposed in (Hickernell, Jiang, Liu, and Owen 2013). The (Chow and Robbins 1965) algorithm relies on Normal approximation, and the coverage guarantee provided is asymptotic (as $\varepsilon \to 0$). We observe that in a reasonable setting with a fixed ε and δ, the probability that their CI does not contain the mean is much larger than δ. The (Hickernell, Jiang, Liu, and Owen 2013) algorithm based on the Berry-Esseen inequality does provide $(\varepsilon,\ \delta)-$coverage guarantee but as δ gets smaller, it takes more samples to construct CI as compared to our algorithm. This is in line with our theoretical result that algorithms that construct symmetric CI around the sample mean will fail to match the lower bound.

Brief relevant literature survey. In the statistics literature, fixed width interval analysis dates back to (Chow and Robbins 1965) that provides asymptotic $(\varepsilon,\ \delta)-$coverage guarantees (see also (Siegmund 1985) and (Yu 1989)). A great deal of research is also done in multi stage estimation algorithm for CI (Mukhopadhyay and Datta 1996) which has similar drawbacks.

Our work is closely related to (Hickernell, Jiang, Liu, and Owen 2013) which studies the construction of fixed-width CI coverage guarantee in non-asymptotic regime. Using the Berry-Esseen inequality, they propose an algorithm that utilizes the upper bound on the modified kurtosis. While we do not make such an assumption, we assume that the distributional form is known. Their algorithm constructs symmetric CI around the sample mean, hence requiring a larger sample size for constructing CI as compared to our algorithm for small δ (see our numerical results).

The rest of this paper is organized as follows: Section 2 gives the formal problem description, Section 3 considers the lower bound, and Section 4 provides asymptotically optimal estimation algorithm. Section 5 provides a numerical study to illustrate our findings and compares the proposed algorithm with algorithms presented in the literature. In Section 6, we discuss some potential extensions of our work, and Section 7 provides the proof of all the results. Section 8 provides the definition of a function used in the analysis.

2 PROBLEM DESCRIPTION

Suppose $X_1, X_2, \ldots$ are i.i.d. copies of a random variable X with distribution function F and mean of X is μ, i.e., $\mathbb{E}[X] = \mu$. Let $\mathbb{P}$ denote the probability measure induced by F. We further assume that the distribution of X, i.e., F belongs to canonical single parameter exponential family $\mathscr{P}$ (see (Garivier and Cappé 2011)). It is defined as follows:

$$\mathscr{P} = \left\{ p_\theta : \theta \in \Theta, \frac{dp_\theta}{d\xi} = \exp(\theta x - b(\theta) + h(x)) \right\},$$

where $\Theta \subset \mathbf{R}$, ξ is some reference measure on $\mathbf{R}$, $h(\cdot)$ is a real-valued function and $b(\cdot)$ is twice differentiable strictly convex function. As can be easily seen, mean of the distribution p_θ is $b'(\theta)$. Further, as is well known, each distribution $p_\theta \in \mathscr{P}$ can be parameterized either by θ or by its mean. We do not know the parameter associated with the distribution of X. Hence the mean of the distribution, which we denote by μ, i.e., $\mathbb{E}[X] = \mu$, is unknown. Our goal is to estimate the confidence interval of width $\varepsilon > 0$ which contains the mean μ, with high probability using the least number of i.i.d. copies of X.

For $n \geq 1$, let $\mathscr{F}_n$ denote the information contained in the σ-algebra generated by $\{X_k, k \leq n\}$. We aim to find a stopping time τ with respect to $\{\mathscr{F}_n : n = 1, 2, 3...\}$ such that $\mu_\tau^R - \mu_\tau^L \leq \varepsilon$, where $[\mu_s^L, \mu_s^R]$ be the estimated confidence interval after observing $X_1, X_2, \ldots, X_s$. Let $\mathscr{F}_\tau$ denote the σ algebra associated with the stopping time τ.

Our interest is in developing a fixed width CI that contains the mean with high probability. The definition below formalizes this notion.

Definition 1 Consider any random variable X whose distribution belongs to a single parameter exponential family $\mathscr{P}$. An algorithm consisting of a stopping time τ with respect to $\{\mathscr{F}_n : n = 1, 2, 3...\}$ and an estimated interval $[\mu_\tau^L, \mu_\tau^R]$ whose width is less than ε is said have **(ε, δ)$-$coverage guarantee**, if:

1. $\mathbb{P}\{\tau < \infty\} = 1$, and
2. estimated interval $[\mu_\tau^L, \mu_\tau^R]$ contains the mean $\mu = \mathbb{E}[X]$ with probability at least $1 - \delta$, i.e.,

$$\mathbb{P}\{\mu \notin [\mu_\tau^L, \mu_\tau^R]\} \leq \delta. \tag{1}$$

In the set of all (ε, δ)$-$coverage guarantee algorithms, we would like to identify an algorithm which uses the least number of samples of X, i.e, the one that minimizes $\mathbb{E}[\tau]$. To this end, for a given δ and ε, we first develop a lower bound on the expected number of samples, i.e., $\mathbb{E}[\tau]$, required for any (ε, δ)$-$coverage algorithm. We then construct a (ε, δ)$-$coverage guarantee algorithm inspired by the lower bound using the "plug-in" approach. Thereafter we prove that our algorithm is asymptotically optimal as $\delta \to 0$, i.e., the expected number of samples used by our algorithm matches the lower bound.

3 LOWER BOUND

Recall $\mathscr{P} = \{p_\theta, \theta \in \Theta\}$, where $\Theta \subset \mathbf{R}$. Let $KL(p_\theta, p_{\tilde{\theta}})$ represent the KL divergence of p_θ with respect to $p_{\tilde{\theta}}$. Since there is a one-to-one mapping between the mean of the distribution and the parameter θ, we can define a divergence function which takes two means μ, $\tilde{\mu}$ as inputs and maps this pair to the KL divergence between the two distributions p_{θ_μ} to $p_{\theta_{\tilde{\mu}}}$ in $\mathscr{P}$, that is,

$$d(\mu, \tilde{\mu}) \triangleq KL(p_{\theta(\mu)}, p_{\theta(\tilde{\mu})}) = b(\tilde{\theta}) - b(\theta) - b'(\theta)(\tilde{\theta} - \theta),$$

such that $b'(\theta_\mu) = \mu$ and $b'(\theta_{\tilde{\mu}}) = \tilde{\mu}$.

Let S denote the support of $d(\mu, \cdot)$ and denote $\sup S = \overline{S}$ and $\inf S = \underline{S}$. It is well known that $d(\mu, \cdot)$ is a strict quasi-convex function, $d(\mu, \tilde{\mu}) > 0 \forall \tilde{\mu} \neq \mu$ and $d(\mu, \mu) = 0$.

To start the analysis of the lower bound we need to define an alternate probability measure $\tilde{\mathbb{P}}$ which is induced by distribution of X, $\tilde{F} \in \mathscr{P}$ in which $\mathbb{E}[X] = \tilde{\mu} \neq \mu$. The literature presents a non-asymptotic inequality for $\delta-$correct (equivalent of (ε, δ)$-$coverage guarantee) in the multi-arm bandit literature context that gives lower bounds on the expected number of samples generated by each arm when there are finitely many of such arms (see, e.g., (Lattimore and Szepesvári 2020), (Kaufmann, Cappé, and Garivier 2016)). If we compare our problem with bandit literature, we have a single arm bandit problem. Using Lemma 1 from (Kaufmann, Cappé, and Garivier 2016), it follows that for a stopping time τ that is almost surely finite, with probability measures $\mathbb{P}$ and $\tilde{\mathbb{P}}$ as above,

$$\mathbb{E}_{\mathbb{P}}[\tau] d(\mu, \tilde{\mu}) \geq \sup_{\mathscr{E} \in \mathscr{F}_\tau} \phi(\mathbb{P}(\mathscr{E}), \tilde{\mathbb{P}}(\mathscr{E})), \tag{2}$$

where $\phi(p_1,p_2) \triangleq p_1 \log \frac{p_1}{p_2} + (1-p_1)\log\left(\frac{1-p_1}{1-p_2}\right)$ for $p_1,p_2 \in (0,1)$. $\mathbb{E}_{\mathbb{P}}[\cdot]$ denotes the expectation operator under probability measure $\mathbb{P}$. For notational ease, we omit $\mathbb{P}$ from $\mathbb{E}_{\mathbb{P}}[\cdot]$ and denote it by $\mathbb{E}[\cdot]$.

Using (2), we get the following proposition.

Proposition 1 Let $\delta \in (0,1)$ and stopping time τ_δ denote the total number of samples used in estimating confidence interval of width ε. For any $(\varepsilon,\ \delta)-$coverage guarantee algorithm with an almost surely finite stopping time τ_δ, we have:

$$\frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{\min_{\tilde{\mu}\in\{\mu-\varepsilon,\mu+\varepsilon\}} d(\mu,\tilde{\mu})}. \tag{3}$$

Notice that it may be possible that $\mu+\varepsilon > \overline{S}$, in that case, we define $d(\mu,\mu+\varepsilon) = \infty$. Similarly, if $\mu-\varepsilon < \underline{S}$, we define $d(\mu,\mu-\varepsilon) = \infty$.

Proof of the above lower bound relies on the fact that when the algorithm observes i.i.d. samples from a distribution with mean μ, it needs to see enough samples to be convinced with probability at least $1-\delta$ that the samples are not coming from a distribution with mean $\tilde{\mu}$ for $\{\tilde{\mu} \geq \mu+\varepsilon\} \cup \{\tilde{\mu} \leq \mu-\varepsilon\}$.

However, the above lower bound does not provide insights into the construction of CI of the mean. To gain further insight that may help develop an algorithm, we consider the asymptotic regime where $\delta \to 0$. Before stating the asymptotic lower bound in the $\delta \to 0$ regime, we further restrict our attention to *stable* estimation algorithms defined below.

Assumption 1 Let τ_δ be the stopping time of a $(\varepsilon,\delta)-$coverage algorithm with CI $[\mu^L_{\tau_\delta},\mu^R_{\tau_\delta}]$, where $\mu^R_{\tau_\delta} - \mu^L_{\tau_\delta} \leq \varepsilon$. The algorithm is called **stable** if $\mu^L_{\tau_\delta} \xrightarrow{p} a$ and $\mu^R_{\tau_\delta} \xrightarrow{p} b$ as $\delta \to 0$, where a and b are constants.

Assumption 1 is satisfied by most of the existing algorithms. Now we state one of the key results of the paper.

Theorem 2 For any $(\varepsilon,\delta)-$coverage and stable algorithm with an almost surely finite stopping time τ_δ, we have

$$\liminf_{\delta\to 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{d(\mu,\mu^L)}, \tag{4}$$

where, for a given μ, μ^L uniquely solves,

$$d(\mu,\mu^L) = d(\mu,\mu^R), \text{ and } \mu^L+\varepsilon = \mu^R. \tag{5}$$

Remark. To see the tightness of the lower bound provided in Theorem 2, consider the Gaussian distribution (with known variance). Here, μ^L and μ^R correspond to $\mu-\frac{\varepsilon}{2}$ and $\mu+\frac{\varepsilon}{2}$, respectively. It follows that $d(\mu,\mu^L) = \frac{4}{\varepsilon^2}$ and $d(\mu,\mu-\varepsilon) = d(\mu,\mu+\varepsilon) = \frac{2}{\varepsilon^2}$. Hence the right hand side of (4) is $\frac{4}{\varepsilon^2}$, while the right hand side of (3) is $\frac{2}{\varepsilon^2}$.

Corollary 3 For any $(\varepsilon,\delta)-$coverage and stable algorithm with a finite stopping time τ_δ which constructs symmetric CI around an unbiased estimator of μ, we have,

$$\liminf_{\delta\to 0} \frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{\min_{\tilde{\mu}\in\{\mu-\varepsilon/2,\mu+\varepsilon/2\}} d(\mu,\tilde{\mu})}. \tag{6}$$

Remark. It is worth noticing that lower bound given in (6) is larger than the one given in (4) as $\min_{\tilde{\mu}\in\{\mu-\varepsilon/2,\mu+\varepsilon/2\}} d(\mu,\tilde{\mu}) \leq d(\mu,\mu^L)$. Hence to construct optimal, i.e., using least i.i.d. samples, $(\varepsilon,\delta)-$coverage and stable algorithm, one should not use symmetric CI around any unbiased estimator (unless their is symmetry which yields $d(\mu,\mu-\varepsilon/2) = d(\mu,\mu^L)$). The proposition below formalizes it.

Proposition 4 Any $(\varepsilon,\delta)-$coverage and stable algorithm, with a finite stopping time τ_δ, can match the lower bound, i.e., $\lim_{\delta\to 0}\left(\frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)}\right) = \frac{1}{d(\mu,\mu^L)}$, only when $\mu^L_{\tau_\delta} \xrightarrow{p} \mu^L$ and $\mu^R_{\tau_\delta} \xrightarrow{p} \mu^R$, where μ^L and μ^R satisfy (5).

The above proposition suggests asymptotically CI should be symmetric around μ w.r.t. KL divergence. In the next section, using this insight we construct our proposed algorithm. Since the algorithm proposed in (Hickernell, Jiang, Liu, and Owen 2013) is symmetric around the sample mean, Corollary 3 states that that algorithm will at best match the bound provided in (6) which in general fall short of being asymptotically optimal.

4 PROPOSED ASYMPTOTICALLY OPTIMAL STABLE ALGORITHM P1

Our algorithm proceeds sequentially, at stage n the algorithm observes n samples and constructs a confidence interval as follows. Inspired by Lemma 4, we construct the confidence interval around the sample average $\hat{\mu}_n$ in such a way so that it is symmetric around $\hat{\mu}_n$ w.r.t. KL divergence, but the width of the confidence interval monotonically decreases with n at some appropriate rate to ensure $(\varepsilon,\delta)-$coverage guarantee.

Formally at stage n, first we estimate μ by $\hat{\mu}_n$, then we define μ_n^L and μ_n^R as follows:

$$\mu_n^R \triangleq \max\{q > \hat{\mu}_n : d(\hat{\mu}_n, q) \leq \beta(n,\delta)\} \text{ and } \mu_n^L \triangleq \min\{q < \hat{\mu}_n : d(\hat{\mu}_n, q) \leq \beta(n,\delta)\},$$

where, $\beta(n,\delta)$ is $O\left(\frac{\log(1/\delta)}{n}\right)$ and explicitly defined in §8 and is chosen to ensure the $(\varepsilon,\ \delta)-$coverage guarantee of the algorithm. Our confidence interval at the end of stage n, will be $[\mu_n^L, \mu_n^R]$.

After the construction of the confidence interval in stage n the algorithm decides to stop if $\mu_n^R - \mu_n^L \leq \varepsilon$. In that case stopping time is $\tau(\delta) = n$. Otherwise we sample X_{n+1} and move to stage $n+1$.

We make the following technical assumption which will be required in proving the asymptotic optimality of the proposed algorithm **P1**.

Assumption 2 The following holds:

$$\lim_{\mu\uparrow\overline{S}} d(\mu, \mu-\varepsilon) < \infty, \text{ and } \lim_{\mu\downarrow\underline{S}} d(\mu, \mu+\varepsilon) < \infty.$$

It can easily be verified that the above assumption is satisfied for all practically relevant distributions in $\mathscr{P}$ such as Gaussian (known variance), Bernoulli, Gamma (known shape parameter) and Poisson. Before stating our main result, we state a supporting lemma as a conclusion of assumption 2.

Lemma 5 Under Assumption 2, $sup_{\mu\in S} d(\mu, \mu^L) < \infty$, where μ^L satisfies (5).

Now we state our main result of this section, which characterizes the performance of our proposed algorithm **P1** and proves its asymptotic optimality.

Theorem 6 *(*$(\varepsilon,\ \delta)-$**coverage guarantee and stable nature of P1)** Under Assumption 2, for algorithm **P1**, the following holds:

a): For a given $\delta \in (0,1)$, τ_δ is finite almost surely.

b): P1 has $(\varepsilon,\ \delta)-$coverage guarantee, that is, $\mathbb{P}\left(\mu \notin [\mu_{\tau_\delta}^L, \mu_{\tau_\delta}^R]\right) \leq \delta$.

c): P1 is a stable algorithm. Formally, $\lim_{\delta\to 0}\mu_{\tau_\delta}^L = \mu^L$ and $\lim_{\delta\to 0}\mu_{\tau_\delta}^R = \mu^R$ almost surely.

Theorem 7 (Asymptotic optimality of P1) Under Assumption 2, for algorithm **P1**, the following holds:

$$\mathbb{P}\left(\lim_{\delta\to 0}\frac{\tau_\delta}{\log(1/\delta)} = \frac{1}{d(\mu,\mu^L)}\right) = 1 \text{ and } \lim_{\delta\to 0}\left(\frac{\mathbb{E}[\tau_\delta]}{\log(1/\delta)}\right) = \frac{1}{d(\mu,\mu^L)}. \tag{7}$$

where, μ^L uniquely solves, $d(\mu,\mu^L) = d(\mu, \mu^L + \varepsilon)$.

5 NUMERICAL EXPERIMENTS

For the first experiment, we simulate our proposed algorithm **P1**. We use exponential distribution for all our numerical experiments. We set the mean, i.e., μ to be 1. We seek the CI width, i.e., ε to be 0.1. We run

P1 for error probability $\delta = .05$, .03, .01 and .005. We run 2000 replications of **P1**. Results are presented in Table 1. We find that CI estimated by **P1** for all of the replications and for levels of error probability δ always contains the mean even though we allow for error in the probability that CI contains the mean to be $\delta > 0$, but it takes more i.i.d samples to construct CI for positive values of δ as suggested by lower bound but as δ diminishes to zero, as predicted by theory, the performance of the algorithm **P1** improves in terms of the ratio of the expected number of samples used by **P1** and lower bound provided in (4).

We run the algorithm given in (Chow and Robbins 1965) for the same set-up, referred to as **C1**, to compare the performance with our proposed algorithm **P1**. We provide the results in Table 1. We observe that the probability that CI obtained by **C1** misses out on the mean is higher than δ at the cost of stopping early (even faster than the lower bound on $(\varepsilon,\ \delta)-$coverage guarantee algorithms). This is supported by theory too as **C1** has the $(\varepsilon,\ \delta)-$coverage guarantee only when $\varepsilon \to 0$. To compare the performance of **P1** with algorithm given in (Hickernell, Jiang, Liu, and Owen 2013), referred as **H1**, we simulate it for the same set up. Since kurtosis for exponential distribution is 9 and skewness is 2, we use these values in the algorithm **H1**. Results are reported in Table 1. We find that CI estimated by **H1** also always contains the mean even though we allow for an error in the probability that CI contains the mean to be δ. For $\delta = .05$ and .03, **H1** takes less sample than **P1** but as δ gets smaller than .03, asymptotic theory is confirmed and our proposed algorithm **P1** outperforms the **H1**. It is worth noting that the number of samples for **H1** depends on the bound on kurtosis. In our numerical experiments, it is assumed to be tight, if it were loose, more samples would be needed.

Lastly we plot the lower bound provided in (4) and (6) with δ as shown in Figure 1. These lower bounds are valid when δ is very small. Lower bound provided in (4) is valid for $(\varepsilon,\delta)-$coverage algorithms while lower bound provided in (6) is valid for $(\varepsilon,\delta)-$coverage algorithms which constructs symmetric CI around sample mean or any unbiased estimator of mean. We clearly see that the lower bound provided in (6) is larger than the lower bound provided in (4). It follows that **H1** can maximum hope to match the lower bound provided in (6) as $\delta \to 0$, while **P1** will match the smaller lower bound provided in (4).

Table 1: Performance of Algorithms **P1**, **H1** and **C1** for CI of width $\varepsilon = 0.1$ for exponential distribution with mean $\mu = 1$. 95 % CI for the estimate of $\mathbb{E}[\tau_\delta]$ for all the three algorithms is less than 20 samples. For **P1** and **H1**, actual probability of μ lying in the estimated CI comes out to be 1.

δ	Lower Bound (LB) given in (4)	Algorithm **P1**		Algorithm **H1**		Algorithm **C1**	
		$\mathbb{E}[\tau_\delta]$	Ratio of $\mathbb{E}[\tau_\delta]$ & LB	$\mathbb{E}[\tau_\delta]$	Ratio of $\mathbb{E}[\tau_\delta]$ & LB	$\mathbb{E}[\tau_\delta]$	Actual prob. of μ lies in CI
.05	1.7×10^3	1.60×10^4	9.41	1.10×10^4	6.47	1.5×10^3	0.899$\pm$.006
.03	2.1×10^3	1.70×10^4	8.09	1.60×10^4	7.62	1.8×10^3	0.925$\pm$.005
.01	2.9×10^3	1.91×10^4	6.59	3.29×10^4	11.34	2.5×10^3	0.950$\pm$.004
.005	3.5×10^3	2.03×10^4	5.8	5.20×10^4	14.85	3.1×10^3	0.961$\pm$.003

6 CONCLUSION

We study the problem of constructing fixed-width CI of the mean of a distribution with a coverage guarantee. We first provide the lower bound result on the expected number of samples required for any (ε,δ) coverage guarantee algorithm. Further, we propose an (ε,δ) coverage guarantee algorithm **P1**, which is asymptotically optimal as it matches the lower bound when δ diminishes to 0.

In this paper, we restrict our attention to the single parameter exponential family. Our results may be extended to the bounded random variables as well as random variables with explicit upper bounds on $1+\eta$ moment, where $\eta > 0$ (see (Agrawal, Juneja, and Glynn 2019)). The approach should also extend to multi-dimensional random variables where we look for fixed confidence regions.

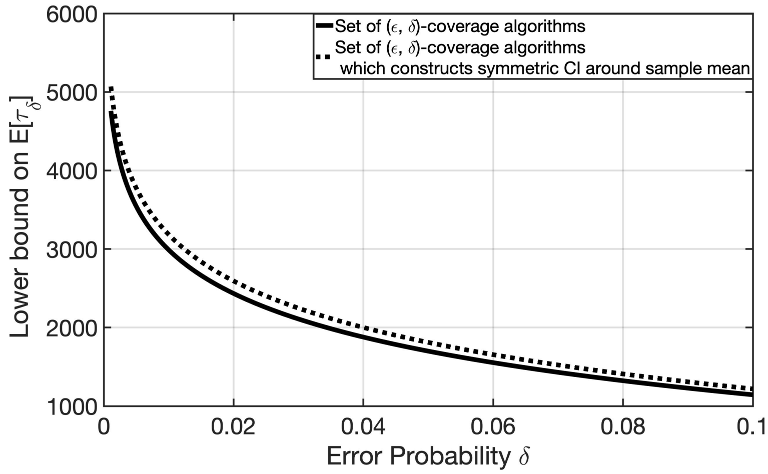

Figure 1: Lower bound provided in (4) and (6) on $\mathbb{E}[\tau_\delta]$ for small δ values for exponential distribution with mean 1 for constructing CI of width (ε) 0.1.

7 PROOFS

Proof of Proposition 1 Recall that alternate probability measure $\tilde{\mathbb{P}}$ is induced by, distribution of a random variable X, $\tilde{F} \in \mathscr{P}$ in which $\mathbb{E}[X] = \tilde{\mu} \neq \mu$. Consider following event $\mathscr{E} = \{\mu \notin [\mu^L_{\tau_\delta}, \mu^R_{\tau_\delta}]\}$. For this event from (1) we get, $\mathbb{P}(\mathscr{E}) \leq \delta$. Since $\tilde{\mathbb{P}}(\mathscr{E}) = \tilde{\mathbb{P}}\{\mu \notin [\mu^L_{\tau_\delta}, \mu^R_{\tau_\delta}]\} \geq \tilde{\mathbb{P}}\{\tilde{\mu} \in [\mu^L_{\tau_\delta}, \mu^R_{\tau_\delta}]\} \geq 1-\delta$, $\forall \tilde{\mu} \geq \mu+\varepsilon$ and $\forall \tilde{\mu} \leq \mu-\varepsilon$. If $\mathbb{P}(\mathscr{E}) \leq \delta$ and $\tilde{\mathbb{P}}(\mathscr{E}) \geq 1-\delta$, we get $\phi(\mathbb{P}(\mathscr{E}), \tilde{\mathbb{P}}(\mathscr{E})) \geq \log\left(\frac{1}{2.4\delta}\right)$. Using (2) we get, $\mathbb{E}[\tau_\delta] \geq \frac{\log\left(\frac{1}{2.4\delta}\right)}{\inf_{\tilde{\mu} \geq \mu+\varepsilon \bigcup \tilde{\mu} \leq \mu-\varepsilon} d(\mu, \tilde{\mu})}$ and result follows using uni modality of $d(\mu, \tilde{\mu})$ in $\tilde{\mu}$.

Proof of Theorem 2 Let $\mathscr{A}$ be the set of algorithms which satisfy the $(\varepsilon, \delta)-$coverage property and are stable. Let $\tilde{\mathbb{P}}$ be the alternate probability measure which is induced by, distribution of a random variable X, $\tilde{F} \in \mathscr{P}$ in which $\mathbb{E}[X] = \tilde{\mu} \neq \mu$. Now we divide the set of algorithms $\mathscr{A}$ into two parts $\mathscr{A}_1$ and $\mathscr{A}_2$. Any algorithm with stopping time τ_δ in $\mathscr{A}_1$ satisfies the following property.

$$\mu^L_{\tau_\delta} \xrightarrow{p} \mu^L. \tag{8}$$

Hence it follows that in any with stopping time τ_δ in $\mathscr{A}_2$ satisfies the following property.

$$\mu^L_{\tau_\delta} \xrightarrow{p} a, \text{ where, a } < \mu \neq \mu^L. \tag{9}$$

First we find the lower bound on $\mathbb{E}[\tau_\delta]$ on algorithms in set $\mathscr{A}_2$. Using (2), on any with stopping time τ_δ in $\mathscr{A}_2$, for alternate probability measure $\tilde{\mathbb{P}}$ in which $\mathbb{E}[X_1] = a-\eta$, we get,

$$\liminf_{\delta \to 0} \frac{\mathbb{E}[\tau_\delta] d(\mu, a-\eta)}{\log(1/2.4\delta)} \geq \liminf_{\delta \to 0} \frac{\sup_{\mathscr{E} \in \mathscr{F}_\tau} \phi(\mathbb{P}(\mathscr{E}), \tilde{\mathbb{P}}(\mathscr{E}))}{\log(1/2.4\delta)}, \tag{10}$$

where η is a small fixed positive number. We claim that the following holds:

$$\liminf_{\delta \to 0} \frac{\phi(\mathbb{P}(\mathscr{E}), \tilde{\mathbb{P}}(\mathscr{E}))}{\log(1/2.4\delta)} \geq 1, \tag{11}$$

where, $\mathscr{E} = \{a-\eta \notin [\mu^L_{\tau_\delta}, \mu^R_{\tau_\delta}]\}$, where, $\mu^R_{\tau_\delta} - \mu^L_{\tau_\delta} \leq \varepsilon$.

To prove the claim made in (11), using (1), we get $\tilde{\mathbb{P}}(\mathscr{E}) \leq \delta$. Now observe that, $\mathbb{P}(\mathscr{E}) = \mathbb{P}\{a-\eta \notin [\mu^L_{\tau_\delta}, \mu^R_{\tau_\delta}]\} \geq \mathbb{P}\{a-\eta < \mu^L_{\tau_\delta}\}$. Using the (9), we get, $\lim_{\delta \to 0} \mathbb{P}(\mathscr{E}) = 1$. Using the definition of $\phi(\cdot,\cdot)$, proof of the claim follows trivially. Now we come back to our original proof and using the (11), we get $\liminf_{\delta\to 0}\frac{\mathbb{E}[\tau_\delta]}{\log(1/4\delta)} \geq \frac{1}{d(\mu,a-\eta)}$. Taking $\eta \to 0$, we get that for any in set $\mathscr{A}_2$, following holds $\liminf_{\delta\to 0}\frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{d(\mu,a)}$. If we have chosen alternate probability measure $\tilde{\mathbb{P}}$ in which $\mathbb{E}[X_1] = a+\varepsilon+\eta$, where η is a fixed positive number. Following the similar steps, we would get, $\liminf_{\delta\to 0}\frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{d(\mu,a+\varepsilon)}$. It follows that for any algorithm in $\mathscr{A}_2$, $\liminf_{\delta\to 0}\frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{\min_{\tilde{\mu}\in\{a,a+\varepsilon\}} d(\mu,\tilde{\mu})}$. Following the similar steps for any in $\mathscr{A}_1$, we would get, $\liminf_{\delta\to 0}\frac{\mathbb{E}[\tau_\delta]}{\log(1/2.4\delta)} \geq \frac{1}{\min_{\tilde{\mu}\in\{\mu^L,\mu^L+\varepsilon\}} d(\mu,\tilde{\mu})}$.

Hence to get the lower bound on algorithms in $\mathscr{A}$, we combine the lower bound on $\mathbb{E}[\tau_\delta]$ for algorithms in $\mathscr{A}_2$ and $\mathscr{A}_1$. By definition of μ^L, we know that $d(\mu,\mu^L) = d(\mu,\mu^L+\varepsilon)$. Using the uni-modality of $d(\mu,\cdot)$ we get that $d(\mu,\mu^L) = \max\{\min_{\tilde{\mu}\in\{a,a+\varepsilon\}} d(\mu,\tilde{\mu}) : a<\mu,\, a\in S\}$. This completes the proof. Notice that the proof of Corollary 3 will follow if we fix $a = \mu-\varepsilon/2$ and the proof of Proposition 4 follows from the fact that lower bound can only be achieved for the algorithms in set $\mathscr{A}_1$.

Proof of Lemma 5 Notice that μ^L is an implicit function of μ, which uniquely solves (5). Hence we represent μ^L as $\mu^L(\mu)$. It follows that $\mu^L(\mu)$ is a continuous function of μ from the definition of μ^L, which further implies that $d(\mu,\mu^L(\mu))$ is a continuous function of μ. Now using the fact that

$$\lim_{\mu\uparrow\bar{S}} d(\mu,\mu^L(\mu)) \leq \lim_{\mu\uparrow\bar{S}} d(\mu,\mu-\varepsilon), \text{ and } \lim_{\mu\uparrow\bar{S}} d(\mu,\mu^L(\mu)) \leq \lim_{\mu\downarrow\underline{S}} d(\mu,\mu+\varepsilon).$$

Using Assumption 2, we get the desired result.

Proof of Theorem 6

a): Given $\delta\in(0,1)$, we define event $\mathscr{E}_1 = \{\tau_\delta = \infty\}$. We need to show that, $\mathbb{P}(\mathscr{E}_1) = 0$. We prove it by contradiction, suppose $\mathbb{P}(\mathscr{E}_1) > 0$. Using the definition of τ_δ on any sample path in $\mathscr{E}_1$, it follows that $\forall\, n\in\mathbb{Z}^+$, $\mu_n^R - \mu_n^L > \varepsilon$. Since $\lim_{n\to\infty}\beta(n,\delta) = 0$, it follows that $\lim_{n\to\infty} d(\hat{\mu}_n,\mu_n^L) = \lim_{n\to\infty} d(\hat{\mu}_n,\mu_n^R) = 0$. From the definition of μ_n^L, we know that $\hat{\mu}_n - \varepsilon \leq \mu_n^L \leq \hat{\mu}_n$. We also know that for a given $\delta\in(0,1)$, $\hat{\mu}_n$ converges to μ on each sample path. Hence it implies that μ_n^L is a bounded sequence on each sample path. Now using continuity of $d(\cdot,\cdot)$ and the facts mentioned above we get that $\lim_{n\to\infty}\mu_n^L = \mu$ almost surely. Similarly we get $\lim_{n\to\infty}\mu_n^R = \mu$ almost surely. Hence, $\lim_{n\to\infty}(\mu_n^R - \mu_n^L) = 0$ almost surely. It follows that we get the contradiction for any sample path in $\mathscr{E}_1$.

b) (ε, δ)$-$coverage guarantee: To prove the result it suffices to show, $\mathbb{P}(\mu_{\tau_\delta}^R < \mu) \leq \delta/2$ and $\mathbb{P}(\mu_{\tau_\delta}^L > \mu) \leq \delta/2$. Now we prove $\mathbb{P}(\mu_{\tau_\delta}^R < \mu) \leq \delta/2$. From part (a) of this theorem, we know that $\{\tau_\delta < \infty\}$ occurs with probability one for given δ. Hence it follows that,

$$\mathbb{P}(\mu_{\tau_\delta}^R < \mu) \leq \mathbb{P}(\exists\, n : \mu_n^R < \mu). \tag{12}$$

Recall the definition of μ_n^R, $\mu_n^R = \max\{q > \hat{\mu}_n : d(\hat{\mu}_n,q) \leq \beta(n,\delta)\}$. Using the uni modality of $d(\hat{\mu}_n,\cdot)$ we get the following,

$$\{\mu_n^R < \mu\} \subseteq (\{\hat{\mu}_n < \mu\} \cap \{d(\hat{\mu}_n,\mu) \geq \beta(n,\delta)\}). \tag{13}$$

Combining (13) and (12), we get,

$$\mathbb{P}(\mu_{\tau_\delta}^R < \mu) \leq \mathbb{P}(\exists\, n:\ \{\hat{\mu}_n < \mu\} \cap \{d(\hat{\mu}_n,\mu) \geq \beta(n,\delta)\}).$$

Using the one side deviation result for one arm (see Corollary 31 in (Kaufmann and Koolen 2021)) we get the desired result. Similarly we can prove that $\mathbb{P}(\mu_{\tau_\delta}^L > \mu) \leq \delta/2$. This completes the proof.

c): Stable nature of P1: Firstly we claim that,

$$\lim_{\delta\to 0}\tau_\delta = \infty \text{ almost surely.} \tag{14}$$

To prove this, we use the method of contradiction. Hence, we define an event $\mathscr{E}_2 = \{\liminf_{\delta\to 0} \tau_\delta < \infty\}$ and assume $\mathbb{P}(\mathscr{E}_2) > 0$. Recall the definition of $\mu^L_{\tau_\delta}$ and $\mu^R_{\tau_\delta}$, $d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) = d(\hat{\mu}_{\tau_\delta}, \mu^R_{\tau_\delta}) = \beta(\tau_\delta, \delta)$.

Using the fact that $\lim_{\delta\to 0} \beta(\tau_\delta, \delta) = \infty$ for any sample path in $\mathscr{E}_2$, we get,

$$\liminf_{\delta\to 0} d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) = \infty \text{ for any sample path in } \mathscr{E}_2. \tag{15}$$

Recall that $\mu^L(\hat{\mu}_{\tau_\delta})$ will satisfy the following, $d(\hat{\mu}_{\tau_\delta}, \mu^L(\hat{\mu}_{\tau_\delta}) + \varepsilon) = d(\hat{\mu}_{\tau_\delta}, \mu^L(\hat{\mu}_{\tau_\delta}))$, hence it follows from the definition of $\tau(\delta)$ and uni modality of $d(\hat{\mu}_{\tau_\delta}, \cdot)$, $d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) \leq d(\hat{\mu}_{\tau_\delta}, \mu^L(\hat{\mu}_{\tau_\delta}))$. Using Lemma 5, we get, $\liminf_{\delta\to 0} d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) \leq \liminf_{\delta\to 0} d(\hat{\mu}_{\tau_\delta}, \mu^L(\hat{\mu}_{\tau_\delta})) \leq sup_{\mu\in S} d(\mu, \mu^L) < \infty$. This contradicts with (15), which completes the proof of the claim.

Now we come back to our original proof. Using (14) and the strong law of large number we get,

$$\hat{\mu}_{\tau_\delta} \to \mu \text{ almost surely as } \delta \to 0 \tag{16}$$

It follows that, $\mu^L_{\tau_\delta}$ will satisfy the following: $\hat{\mu}_{\tau_\delta} - \varepsilon \leq \mu^L_{\tau_\delta} \leq \hat{\mu}_{\tau_\delta}$. Using (16), it follows that $\mu^L_{\tau_\delta}$ is a bounded sequence on each sample path. Similarly, $\mu^R_{\tau_\delta}$ is a bounded sequence too on each sample path.

Using the definition of τ_δ, $\mu^L_{\tau_\delta}$ and $\mu^R_{\tau_\delta}$, it follows that,

$$d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) = d(\hat{\mu}_{\tau_\delta}, \mu^R_{\tau_\delta}),\ \mu^R_{\tau_\delta} \leq \mu^L_{\tau_\delta} + \varepsilon \text{ and } \mu^R_{\tau_\delta - 1} \geq \mu^L_{\tau_\delta - 1} + \varepsilon. \tag{17}$$

Using the uni modality of $d(\hat{\mu}_{\tau_\delta}, \cdot)$ and combining it with (17),we get,

$$d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) \leq d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta} + \varepsilon) \text{ and } d(\hat{\mu}_{\tau_\delta - 1}, \mu^L_{\tau_\delta - 1}) \geq d(\hat{\mu}_{\tau_\delta - 1}, \mu^L_{\tau_\delta - 1} + \varepsilon). \tag{18}$$

Now we are ready to prove that $\lim_{\delta\to 0} \mu^L_{\tau_\delta} = \mu^L$ almost surely and it will follow trivially that $\lim_{\delta\to 0} \mu^R_{\tau_\delta} = \mu^L + \varepsilon$. We prove it by contradiction, suppose $\mu^L_{\tau_\delta}$ does not converge to μ^L on a positive measure set $\mathscr{E}_3$, i.e., $\mathbb{P}(\mathscr{E}_3) > 0$. Fix any sample path in $\mathscr{E}_3$, and let $\limsup_{\mu^L_{\tau_\delta}} = \overline{K}$ and $\liminf_{\mu^L_{\tau_\delta}} = \underline{K}$, where $\overline{K} \neq \mu^L$ and $\underline{K} \neq \mu^L$. Since $\mu^L_{\tau_\delta}$ is a bounded sequence, hence there will exist a sub-sequence of $\{\delta_k, k \in \mathbb{Z}^+\}$ and $\delta_k \in (0,1)$ such that $\lim_{k\to\infty} \delta_k = 0$ and $\lim_{k\to\infty} \mu^L_{\tau(\delta_k)} = \overline{K}$. Using (18) on the sub-sequence defined above and continuity of $d(\cdot,\cdot)$, we get that,

$$\lim_{k\to\infty} d(\hat{\mu}_{\tau(\delta_k)}, \mu^L_{\tau(\delta_k)}) = \lim_{k\to\infty} d(\hat{\mu}_{\tau(\delta_k)}, \mu^L_{\tau(\delta_k)} + \varepsilon) \implies d(\mu, \overline{K}) = d(\mu, \overline{K} + \varepsilon). \tag{19}$$

Since μ^L uniquely satisfy $d(\mu + \mu^L) = d(\mu, \mu^L + \varepsilon)$, this implies that $\overline{K} = \mu^L$. Similarly we will get, $\underline{K} = \mu^L$, which gives the contradiction. This completes the proof.

Proof of Theorem 7

a):Almost sure convergence: Observe that $\mu^L_{\tau_\delta}$ and $\mu^R_{\tau_\delta}$ satisfy the following:

$$d(\hat{\mu}_{\tau_\delta}, \mu^L_{\tau_\delta}) = d(\hat{\mu}_{\tau_\delta}, \mu^R_{\tau_\delta}) = \beta(\tau_\delta, \delta). \tag{20}$$

Since we know that $\lim_{\delta\to 0} \hat{\mu}_{\tau_\delta} = \mu$ almost surely. Using Theorem 6 **(c)**, we also know that $\lim_{\delta\to 0} \mu^L_{\tau_\delta} = \mu^L$ and $\lim_{\delta\to 0} \mu^R_{\tau_\delta} = \mu^L + \varepsilon$ almost surely. Hence using continuity of $d(\cdot,\cdot)$, (20) and definition of $\beta(n,\delta)$ (see §8), we get,

$$\lim_{\delta\to 0} \frac{3\log[1 + \log(\tau_\delta)] + \mathscr{T}(\log(2/\delta))}{\tau_\delta} = d(\mu, \mu^L) \text{ almost surely.}$$

Notice from the definition of $\mathscr{T}(x)$ (see §8), we get $\lim_{\delta\to 0} \frac{\mathscr{T}(\log(2/\delta))}{\log(1/\delta)} = 1$, hence it follows that

$$\lim_{\delta\to 0} \frac{3\log[1 + \log(\tau_\delta)] + \log(1/\delta)}{\tau_\delta} = d(\mu, \mu^L) \text{ almost surely.}$$

Using (14), we know that $\tau_\delta \to \infty$ as $\delta \to 0$ almost surely, hence we get the desired result.

b) Convergence in Expectation: To get the results of convergence in expectation from almost sure, we will show that $\frac{\tau(\delta)}{\log(1/\delta)}$ is a uniform integrable random variable, which will complete the proof. Hence it suffices to show $\sup_{\delta\in(0,\delta_1)} \frac{\mathbb{E}[\tau(\delta)]^2}{(\log(1/\delta))^2} < \infty$, where, δ_1 is any fixed number in $(0,1)$. Observe that,

$$\mathbb{E}[\tau_\delta]^2 = \sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\tau_\delta > n).$$

Let $\mathscr{B}_n = \{|\hat{\mu}_n - \mu| \le c_1\}$, where c_1 is a well chosen small positive constant such that $\mu + c_1 \in S$ as well as $\mu - c_1 \in S$.

It follows that,

$$\mathbb{E}[\tau_\delta]^2 \le \sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\{\tau_\delta > n\}\cap\mathscr{B}_n) + \sum_{t=n}^{\infty}(2n-1)\mathbb{P}(\mathscr{B}_n^c). \tag{21}$$

We will handle the two series summations given in (21) separately, then we will come back to (21).

Upper bound on $\sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\mathscr{B}_n^c)$: Using Chernoff inequality for single parameter exponential family we get, $\mathbb{P}(\hat{\mu}_n \ge \mu + c_1) \le e^{-nd(\mu+c_1,\mu)}$. Similarly we have $\mathbb{P}(\hat{\mu}_n \le \mu - c_1) \le e^{-nd(\mu-c_1,\mu)}$. Hence it follows that

$$\sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\mathscr{B}_n^c) \le \sum_{n=1}^{\infty}(2n-1)(e^{-nd(\mu-c_1,\mu)} + e^{-nd(\mu+c_1,\mu)}) = c_2, \tag{22}$$

where, c_2 is a fixed positive constant independent of δ

Upper bound on $\sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\{\tau_\delta > n\}\cap\mathscr{B}_n)$: We claim that after $N(\delta) \triangleq O(\log(1/\delta))$ terms $\mathbb{P}(\{\tau_\delta > n\}$ will be 0 under the set $\mathscr{B}_n$, i.e,

$$\mathbb{P}(\{\tau_\delta > n\}\cap\mathscr{B}_n) = 0 \;\forall n \ge N(\delta). \tag{23}$$

We prove the claim mentioned above in (23) later. Using (23), we get,

$$\sum_{n=1}^{\infty}(2n-1)\mathbb{P}(\{\tau_\delta \ge n\}\cap\mathscr{B}_n) \le \sum_{n=1}^{N(\delta)}(2n-1)\mathbb{P}(\tau_\delta \ge n) \le O((\log(1/\delta))^2).$$

Combining the above inequality with (23), and substituting them in (22), we get

$$\sup_{\delta\in(0,\delta_1)} \frac{\mathbb{E}[\tau(\delta)]^2}{(\log(1/\delta))^2} \le \sup_{\delta\in(0,\delta_1)} \frac{O((\log(1/\delta))^2) + c_2}{(\log(1/\delta))^2} < \infty.$$

To complete the proof, all we need to show is that our claim (23) holds. To prove (23), observe that, $d(\mu_n,\mu^L(\mu_n)) > 0$ for $\mu_n \in \mathscr{B}_n$. Using continuity of $d(\mu,\mu^L(\mu))$ in μ, it follows that $inf_{\mu_n\in\mathscr{B}_n} d(\mu_n,\mu^L(\mu_n) = h > 0$.

Now it follows trivially from the definition of $\beta(n,\delta)$ that $\exists N(\delta) = O(\log(1/\delta))$ such that, $\beta(n,\delta) < h\,\forall n \ge N(\delta)$, under the set B_n. Using strict quasi-convexity of $d(\mu,\cdot)$, it follows that claim mentioned in (23) holds. This completes the proof.

8 Definition of $\beta(n,\delta)$.

Recall we define (see (Kaufmann and Koolen 2021)),

$$\beta(n,\delta) = \frac{3\log[1+\log(n)] + \mathscr{T}(\log(2/\delta))}{n}.$$

To state the definition of $\mathscr{T}(x)$, we need to introduce two functions. First for $l \geq 1$ the function $\psi(l) = l - \ln l$ and its inverse $\psi^{-1}(l)$. And the other function is defined for any $y \in [1, e]$ and $x \geq 0$ and given by

$$\tilde{\psi}_y(x) = \begin{cases} e^{1/\psi^{-1}(x)}\psi^{-1}(x) & \text{if } x \geq \psi^{-1}(1/\ln y), \\ y(x - \ln\ln y) & \text{o.w.} \end{cases} \tag{24}$$

Now we define function $\mathscr{T}(x) : \mathbb{R}^+ \to \mathbb{R}^+$ as follows $\mathscr{T}(x) = 2\tilde{\psi}_{3/2}\left(\frac{x + \ln 1.64}{2}\right)$.

REFERENCES

Agrawal, S., S. Juneja, and P. Glynn. 2019. "Optimal δ-Correct Best-Arm Selection for General Distributions". *arXiv preprint arXiv:1908.09094*.

Chow, Y. S., and H. Robbins. 1965. "On the asymptotic theory of fixed-width sequential confidence intervals for the mean". *The Annals of Mathematical Statistics* 36(2):457–462.

Garivier, A., and O. Cappé. 2011. "The KL-UCB algorithm for bounded stochastic bandits and beyond". In *Proceedings of the 24th annual conference on learning theory*, 359–376. JMLR Workshop and Conference Proceedings.

Hickernell, F. J., L. Jiang, Y. Liu, and A. B. Owen. 2013. "Guaranteed conservative fixed width confidence intervals via Monte Carlo sampling". In *Monte Carlo and Quasi-Monte Carlo Methods 2012*, 105–128. Springer.

Kaufmann, E., O. Cappé, and A. Garivier. 2016. "On the complexity of best-arm identification in multi-armed bandit models". *The Journal of Machine Learning Research* 17(1):1–42.

Kaufmann, E., and W. M. Koolen. 2021. "Mixture martingales revisited with applications to sequential tests and confidence intervals". *Journal of Machine Learning Research* 22(246):1–44.

Lattimore, T., and C. Szepesvári. 2020. *Bandit algorithms*. Cambridge University Press.

Mukhopadhyay, N., and S. Datta. 1996. "On sequential fixed-width confidence intervals for the mean and second-order expansions of the associated coverage probabilities". *Annals of the Institute of Statistical Mathematics* 48(3):497–507.

Siegmund, D. 1985. *Sequential analysis: tests and confidence intervals*. Springer Science & Business Media.

Yu, K. F. 1989. "On fixed-width confidence intervals associated with maximum likelihood estimation". *Journal of Theoretical Probability* 2(2):193–199.

AUTHOR BIOGRAPHIES

VIKAS DEEP is a doctoral student in the Operations department of Kellogg School Management, Northwestern University. His e-mail address is dvikas131@gmail.com.

ACHAL BASSAMBOO is Charles E. Morrison Professor of Decision Sciences at Kellogg School of Management, Northwestern University. His email address is a-bassamboo@kellogg.northwestern.edu.

SANDEEP JUNEJA is a Senior Professor at the School of Technology and Computer Science Tata Institute of Fundamental Research. His e-mail address is juneja@tifr.res.in.

ASSAF ZEEVI is the Vice Dean for Research and Henry Kravis Professor of Business at the Graduate School of Business, Columbia University. His e-mail address is assaf@gsb.columbia.edu.