

---

# Near-Optimal Confidence Sequences for Bounded Random Variables

---

Arun Kumar Kuchibhotla<sup>\*1</sup> Qinqing Zheng<sup>\*2</sup>

## Abstract

Many inference problems, such as sequential decision problems like A/B testing, adaptive sampling schemes like bandit selection, are often online in nature. The fundamental problem for online inference is to provide a sequence of confidence intervals that are valid uniformly over the growing-into-infinity sample sizes. To address this question, we provide a near-optimal confidence sequence for bounded random variables by utilizing Bentkus' concentration results. We show that it improves on the existing approaches that use the Cramér-Chernoff technique such as the Hoeffding, Bernstein, and Bennett inequalities. The resulting confidence sequence is confirmed to be favorable in synthetic coverage problems, adaptive stopping algorithms, and multi-armed bandit problems.

## 1. Introduction

The abundance of data over the decades has increased the demand for sequential algorithms and inference procedures in statistics and machine learning. For instance, when the data is too large to fit in a single machine, it is natural to split data into small batches and process one at a time. Besides, many industry or laboratory data, like user behaviors on a website, patient records, temperature histories, are naturally generated and available in a sequential order. In both scenarios, the collection or processing of new data can be costly, and practitioners often would like to stop data sampling when a required criterion is satisfied. This gives the pressing call for algorithms that minimize the number of sequential samples subject to the prescribed accuracy of the estimator is satisfied.

Many important problems fit into this framework, including

---

<sup>\*</sup>Equal contribution <sup>1</sup>Department of Statistics and Data Science, Carnegie Mellon University. <sup>2</sup>Facebook AI Research. Part of the work was done before the author joined FAIR. Correspondence to: Arun Kumar Kuchibhotla <arunku@stat.cmu.edu>, Qinqing Zheng <zhengqinqing@gmail.com>.

sequential hypothesis testing problems such as testing positiveness of the mean (Zhao et al., 2016), testing equality of distributions and testing independence (Balsubramani & Ramdas, 2016; Yang et al., 2017), A/B testing (Johari et al., 2015; 2017), sequential probability ratio test (Wald, 2004), best arm identification for multi-arm bandits (MAB) (Zhao et al., 2016; Yang et al., 2017), etc. All these applications require *confidence sequences* to determine the number of samples required for a certain guarantee.

A simple example to start from is estimating the mean of a random variable from sequentially available data. This is a classic problem in statistics and widely applied to various applications. An estimator  $\hat{\mu}$  is said to be  $(\varepsilon, \delta)$ -accurate for the mean  $\mu$  if  $\mathbb{P}(|\hat{\mu}/\mu - 1| \leq \varepsilon) \geq 1 - \delta$  (Dagum et al., 2000; Mnih et al., 2008; Huber, 2019). This means that the estimator has a relative error of at most  $\varepsilon$  with probability at least  $1 - \delta$ . In the sequential setting, one important question we would like to answer is *how many samples are needed to obtain an estimator of the mean that is  $(\varepsilon, \delta)$  accurate?* Mnih et al. (2008) shows the answer can be derived from a confidence sequence.

**Definition 1.** Let  $Y_1, Y_2, \dots$  be independent real-valued random variables, available sequentially, with mean  $\mu \in \mathbb{R}$ . Given  $\delta \in [0, 1]$ , a  $1 - \delta$  confidence sequence is a sequence of confidence intervals  $\text{ConfSeq}(\delta) = \{\text{CI}_1(\delta), \text{CI}_2(\delta), \dots\}$ , where  $\text{CI}_n(\delta)$  is constructed on-the-fly after observing data sample  $(Y_1, \dots, Y_n)$ , such that

$$\mathbb{P}(\mu \in \text{CI}_n(\delta) \text{ for all } n \geq 1) \geq 1 - \delta. \quad (1)$$

For the  $(\varepsilon, \delta)$ -mean estimation problem above, suppose one can construct a  $1 - \delta$  confidence sequence of  $\mu$ :

$$\text{ConfSeq}(\delta) = \{\text{CI}_n(\delta) = [\bar{Y}_n - Q_n, \bar{Y}_n + Q_n], \quad n \geq 1\},$$

where  $\bar{Y}_n$  is the empirical mean of the first  $n$  samples. Mnih et al. (2008) shows that with number of samples  $N = \min\{n : (1 - \varepsilon)\text{UB}_n \leq (1 + \varepsilon)\text{LB}_n\}$ , where  $\text{UB}_n$  and  $\text{LB}_n$  are two simple functions of the radius of the confidence intervals, the estimator  $\hat{\mu} = (1/2)\text{sign}(\bar{Y}_N)[(1 - \varepsilon)\text{UB}_N + (1 + \varepsilon)\text{LB}_N]$  is  $(\varepsilon, \delta)$ -accurate. See Algorithm 1 in Section 4.2 for details.

The need for sequential algorithms has triggered a surge of interest in developing sharp confidence sequences. Unlike

the traditional confidence interval in statistics, the guarantee (1) is non-asymptotic and is uniform over the sample sizes. Ideally, we want  $\text{CI}_n(\delta)$  to reduce in width as either  $n$  or  $\delta$  increase. Unfortunately, guarantee (1) is impossible to achieve non-trivially<sup>1</sup> without further assumptions (Bahadur & Savage, 1956; Singh, 1963). In this paper, we assume that the random variables are bounded: there exist known constants  $L, U \in \mathbb{R}$  such that  $\mathbb{P}(L \leq Y_i \leq U) = 1$  for all  $i \geq 1$ , which yields  $\mu \in [L, U]$ . Although boundedness can be replaced by tail assumptions such as sub-Gaussianity or polynomial tails, we will restrict our discussion to the bounded case in this paper.

In recent years, several techniques have been proposed to construct confidence sequences (Zhao et al., 2016; Mnih et al., 2008; Howard et al., 2018). These confidence sequences can be thought as a generalization of classical fixed sample size concentration inequalities including Hoeffding, Bernstein, and Bennett. Arguably the simplest construction of a confidence sequence is based on *stitching* the fixed sample size concentration inequalities. Other techniques include self-normalization, method of mixtures or pseudo-maximization (Victor et al., 2007; Howard et al., 2018). The stitching method (unlike the others) makes use of a union bound (or Bonferroni inequality) which might result in a sub-optimal confidence sequence compared to those obtained from method of mixtures.

To the best of our knowledge, all the existing confidence sequences are built upon concentration results that bound the moment generating function and follow the Cramér–Chernoff technique. The Cramér–Chernoff technique leads to conservative bounds and can be significantly improved (Philips & Nelson, 1995). In this paper, we leverage the refined concentration results introduced by Bentkus (2002). We first develop a “maximal” version of Bentkus’ concentration inequality. Based on it, we construct the confidence sequence via stitching. In honor of Bentkus, who pioneered this line of refined concentration inequalities, we call our confidence sequence as *Bentkus’ Confidence Sequence*. The fixed sample size Bentkus concentration inequality is theoretically an improvement of the best possible Cramér–Chernoff bound; see Theorem 1 and the discussion that follows. This improvement implies that stitching the Bentkus concentration inequality improves upon the stitching of the best possible Cramér–Chernoff bound. Hence, our confidence sequence is an improvement on the stitched Hoeffding, Bernstein, and Bennett confidence sequences. Although this is an obvious fact, we find in applications that our confidence sequence leads to about 50% reduction in sample complexity when compared to the classical ones. Surprisingly, we find in simulations that our confidence se-

quence also improves on the method of mixture confidence sequences that do not use a union bound like stitching.

To summarize, our major contributions are as follows.

- We provide a self-contained introduction to *near-optimal* concentration inequality based on the results of Bentkus (2002; 2004) and Pinelis (2006). Unlike the Cramér–Chernoff bounds, which can be infinitely suboptimal, our bound is optimal up to  $e^2/2$ . In other words, our tail bound is at most  $e^2/2$  times the best tail bound that can be obtained under our assumptions. We believe ours is the first application of Bentkus’ concentration inequality for confidence sequences and machine learning (ML) applications including the best-arm identification problem. All ML algorithms that use classical concentration inequalities like Hoeffding or Bernstein can be improved substantially, by simply replacing them with the concentration inequalities discussed in this paper.
- We use these results in conjunction with a “stitching” method (Zhao et al., 2016; Mnih et al., 2008) to construct non-asymptotic confidence sequences. At sample size  $n$ , for  $\bar{Y}_n = n^{-1} \sum_{i=1}^n Y_i$ , the confidence interval is  $\text{CI}_n(\delta) := [\bar{Y}_n - q_n^{\text{low}}(\delta), \bar{Y}_n + q_n^{\text{up}}(\delta)]$ , for different values  $q_n^{\text{low}}(\delta), q_n^{\text{up}}(\delta) \geq 0$  and they scale like  $\sqrt{\text{Var}(Y_1) \log \log(n)/n}$  as  $n \rightarrow \infty$ .
- Similar to the Bernstein inequality, Bentkus’ method utilizes the variance of  $Y_i$ ’s. Therefore, variance estimation is needed to make the stitched Bentkus confidence sequence actionable in practice. We propose a closed form upper bound of the unknown variance based on one-sided concentration for the non-negative variables  $(Y_i - \mu)^2$  from Pinelis (2016). This one-sided concentration bound is an improvement on the classical Cramér–Chernoff bound (Peña et al., 2008, Theorem 2.19) for non-negative random variables. Once again, this leads to a better upper bound on the unknown variance compared to the ones from Audibert et al. (2009) and Maurer & Pontil (2009).
- We derive a computable form of the Bentkus’ method based on Bentkus et al. (2006), and further provide a *constant* time algorithm to compute it efficiently (see Appendix C). In comparison, a brute-force method leads to a linear time complexity.
- We conduct numerical experiments to verify our theoretical claims. Moreover, we apply the Bentkus confidence sequence to the  $(\varepsilon, \delta)$  mean estimation problem and the best-arm identification problem. For both problems, our method significantly reduces the sample complexity by about  $\frac{1}{2}$  compared with the other methods.

The rest of this article is organized as follows. Section 2 reviews the related work. Section 3 contains our theoretical results. Section 4 presents the experiments that confirm

<sup>1</sup>Of course, if we take  $\text{CI}_n(\delta) = (-\infty, \infty)$ , then (1) is trivially satisfied.

the superiority of our method. Section 5 summarizes the contributions and discusses some future directions.

## 2. Related Work

Several confidence sequences built on classical concentration inequalities have been proposed and can be applied to bounded random variables. Zhao et al. (2016) propose confidence sequences through Hoeffding’s inequality, assuming that  $Y_i$ ’s are  $\frac{1}{2}$ -sub-Gaussian. For random variables supported on  $[L, U]$ , this assumption is satisfied after scaling by  $\frac{1}{U-L}$ . However, this confidence sequence does not scale with the true variance and hence can be conservative. Mnih et al. (2008) building on Audibert et al. (2009) construct a confidence sequence through Bernstein’s inequality. Due to the nature of Bernstein’s inequality, those intervals scale correctly with the true variance. The methods in these papers is stitching of fixed sample size concentration inequalities. As mentioned before, they make use of union bound and can have more coverage than required in practice. In probability, Darling & Robbins (1967) and Victor et al. (2007) (among others) have considered confidence sequences based on martingale techniques and method of mixtures. These methods do not require union bound and can be sharper than the stitched confidence sequences. More recently, Howard et al. (2018) have unified the techniques of obtaining confidence sequences under a variety of assumptions on random variables. This work builds on much of the existing statistics literature and we refer the reader to this paper for a detailed historical account.

All the confidence sequences in the works mentioned above depend on concentration results that bound the moment generating function and follow the Cramér-Chernoff technique. Such concentration results, and consequently the obtained confidence sequences, are conservative and can be significantly improved. To understand the deficiency of such concentration inequalities, consider for example the Bernstein’s inequality: for  $\bar{Y}_n = \sum_{i=1}^n Y_i/n$ ,

$$\mathbb{P}(\sqrt{n}(\bar{Y}_n - \mu) \geq t) \leq e^{-t^2/[2A^2+(U-L)t/(3\sqrt{n})]},$$

which scales like  $\exp(-t^2/(2A^2))$ , for “small”  $t$ . However, the central limit theorem implies

$$\mathbb{P}(\sqrt{n}(\bar{Y}_n - \mu) \geq t) \approx 1 - \Phi(t/A) \leq \frac{e^{-t^2/(2A^2)}}{\sqrt{2\pi((t/A)^2 + 1)}}.$$

See, e.g., Abramowitz & Stegun (1948, Formula 7.1.13). Therefore, Bernstein’s inequality and the true tail differ by the scaling  $\sqrt{2\pi(t^2/A^2 + 1)}$ , which can be significant for large  $t$ . This scaling difference is referred to as the missing factor in Talagrand (1995) and Fan et al. (2012). This missing factor does not exist just with Bernstein’s inequality but also with the optimal bound that could be derived from the

Cramér–Chernoff technique; see the discussion surrounding Eq. (1.4) of Talagrand (1995). This explains why a further improvement is possible and Bentkus (2002) presents such sharper concentration inequalities. Our work essentially builds on the works Bentkus (2002; 2004); Pinelis (2006; 2016), to derive a near-optimal concentration inequality, followed by an improved confidence sequence through the technique of stitching.

Given that Bentkus’ concentration inequality is an improvement on the Cramér–Chernoff inequalities and that the tightness of the stitched confidence sequence is mainly controlled by the sharpness of the fixed sample size concentration inequality used, our results are not entirely unexpected. Because the improvement we obtain over the existing confidence sequences is significant (Figs. 4-6), we believe this paper will be an important addition to the literature for practical ML applications.

## 3. Bentkus’ Confidence Sequences

For any random variable  $Y_i$  with mean  $\mu$ ,  $X_i = Y_i - \mu$  is mean zero and hence we will mostly restrict to the case of mean zero random variables. The result for general  $\mu$  will readily follow; see Theorem 4. We first discuss Bentkus’ concentration inequality for bounded mean zero random variables. Afterwards, we present a refined confidence sequence through stitching. This confidence sequence is not readily actionable because it depends on the true variance of random variables. To address this, we present a practical version where we replace the true variance by an estimated upper bound. This provides an analog of the empirical Bernstein confidence sequence, and we call our method *Empirical Bentkus Confidence Sequence*.

**Assumptions.** Suppose  $X_1, X_2, \dots$  are independent random variables satisfying

$$\mathbb{E}[X_i] = 0, \text{ Var}(X_i) \leq A_i^2, \text{ and } \mathbb{P}(X_i > B) = 0, \quad (2)$$

for all  $i \geq 1$ . We will first derive concentration inequalities under the one-sided bound assumption as in (2) which only requires  $X_i \leq B$  almost surely. To derive actionable versions of the concentration inequalities (with estimated variance), we will impose a two-sided bound assumption.

### 3.1. Bentkus’ Concentration Inequality for a fixed Sample Size

We now present a near-optimal concentration inequality for  $S_t = \sum_{i=1}^t X_i$  that holds uniformly over all sample sizes  $t \leq n$ . The main idea behind the optimality is to replace the exponential function used in the Cramér-Chernoff technique with a slowly growing function. Fix  $\alpha \in [0, \infty]$ , and set

$(a)_+ = \max\{a, 0\}$ . It is easy to verify that for all  $\nu \in \mathbb{R}$ ,

$$\mathbf{1}\{\nu \geq 0\} \leq (1 + \nu/\alpha)_+^\alpha \leq e^\nu. \quad (3)$$

Taking  $\nu = \lambda(S_n - u)$  for some  $\lambda > 0$  in inequality (3) and applying expectation, we obtain for all  $u \in \mathbb{R}$ ,

$$\mathbb{P}(S_n \geq u) \leq \inf_{\lambda \geq 0} \mathbb{E} \left[ (1 + \lambda(S_n - u)/\alpha)_+^\alpha \right]. \quad (4)$$

The second inequality in (3) readily shows that (4) is better than the Cramér-Chernoff bound. Reparameterizing  $\lambda = \alpha/(u - x)$  with  $x \leq u$  in (4), we obtain

$$\mathbb{P}(S_n \geq u) \leq \inf_{x \leq u} \frac{\mathbb{E}[(S_n - x)_+^\alpha]}{(u - x)_+^\alpha}, \quad \forall u \in \mathbb{R}. \quad (5)$$

Next, we bound  $\mathbb{E}[(S_n - x)_+^\alpha]$  for all random variables  $X_i$ 's satisfying (2). This should be done optimally in order to obtain a near-optimal concentration inequality. Surprisingly, for all  $\alpha \geq 2$ ,  $\mathbb{E}[(S_n - x)_+^\alpha]$  can be bounded in terms of a ‘‘worst case’’ two-point distribution satisfying (2).

Define independent random variables  $G_i, i \geq 1$  as

$$\begin{aligned} \mathbb{P}(G_i = -A_i^2/B) &= B^2/(A_i^2 + B^2), \\ \mathbb{P}(G_i = B) &= A_i^2/(A_i^2 + B^2). \end{aligned} \quad (6)$$

These random variables satisfy (2) and  $G_i$ 's are the worst case random variables satisfying (2), in the sense that for all  $n \geq 1, \alpha \geq 2$ , and  $x \in \mathbb{R}$ ,

$$\mathbb{E}[(\sum_{i=1}^n G_i - x)_+^\alpha] = \sup_{X_i \sim (2)} \mathbb{E}[(\sum_{i=1}^n X_i - x)_+^\alpha], \quad (7)$$

where the supremum is over all distributions of  $X_i$ 's satisfying (2). We refer the readers to [Bentkus \(2002, Eq. \(11\)\)](#) and [Pinelis \(2006, Theorem 2.1\)](#) for the proof of (7). The definition of the ‘‘worst-case’’ distribution of  $G_i$ 's follows from finding the best quadratic function that upper bounds  $t \mapsto (t - x)_+^\alpha$ ; see, e.g., [Burgess et al. \(2019, Lemma 8\)](#). [Pinelis \(2006\)](#) proves that (7) holds true when  $t \mapsto (t - x)_+^\alpha$  is replaced with any function  $t \mapsto f(t)$  that has a convex first derivative.

Inequality (5) with  $\alpha = 2$  and (7) show that<sup>2</sup>

$$\mathbb{P}(S_n \geq u) \leq \inf_{x \leq u} \frac{\mathbb{E}[(\sum_{i=1}^n G_i - x)_+^2]}{(u - x)_+^2}, \quad (8)$$

and we find  $u$  such that the right hand side of (8) is upper bounded by  $\delta$ . Set  $\mathcal{A} = \{A_1, A_2, \dots\}$  as the collection of standard deviations and for  $n \geq 1$ , define

$$\tilde{P}_{2,n}(u) := \inf_{x \leq u} \frac{\mathbb{E}[(\sum_{i=1}^n G_i - x)_+^2]}{(u - x)_+^2}. \quad (9)$$

<sup>2</sup>The function  $\alpha \mapsto (1 + \nu/\alpha)_+^\alpha$  increases as  $\alpha$  increases, so using the smallest possible  $\alpha$  leads to the best bound. Because (7) only holds for  $\alpha \geq 2$ ,  $\alpha = 2$  is optimal in this context.




Figure 1: Comparison of the concentration bounds when  $\delta = 0.05$ .  $X_i$  are centered i.i.d.  $\text{Bernoulli}(\frac{1}{4})$ . We give the true standard deviation  $A_i = \frac{\sqrt{3}}{4}$  and upper bound  $B = \frac{3}{4}$  to all the methods. The average failure frequencies across 300 trials and  $1 \leq n \leq 3000$  are: Hoeffding  $0.00205 \pm 0.00261$ , Bernstein  $0.00593 \pm 0.0044$ , Bentkus  $0.01411 \pm 0.00769$ . The Bentkus' bound is the least conservative.

For  $\delta \in [0, 1]$ , define  $q(\delta; n, \mathcal{A}, B)$  as the solution to the equation  $\tilde{P}_{2,n}(u) = \delta$ . In other words,

$$q(\delta; n, \mathcal{A}, B) = \tilde{P}_{2,n}^{-1}(\delta) \quad (10)$$

The inverse exists uniquely for  $\delta \geq \mathbb{P}(\sum_{i=1}^n G_i = nB)$  and is defined to be  $nB + 1$  if  $\delta < \mathbb{P}(\sum_{i=1}^n G_i = nB)$ . The following result provides a refined concentration inequality for  $S_n = \sum_{i=1}^n X_i$ . It is a ‘‘maximal’’ version of Theorem 2.1 of [Bentkus et al. \(2006\)](#), see Appendix D for the proof.

**Theorem 1.** Fix  $n \geq 1$ . If  $X_1, X_2, \dots, X_n$  are independent random variables satisfying (2), then

$$\mathbb{P} \left( \max_{1 \leq t \leq n} S_t \geq q(\delta; n, \mathcal{A}, B) \right) \leq \delta, \quad \forall \delta \in [0, 1]. \quad (11)$$

Further, if  $A_1 = \dots = A_n = A$  and if  $\tilde{q}(\cdot; A, B)$  is a function such that  $\mathbb{P}(\max_{1 \leq t \leq n} S_t \geq n\tilde{q}(\delta^{1/n}; A, B)) \leq \delta$  for all  $\delta \in [0, 1]$  then  $q(\delta; n, \mathcal{A}, B) \leq n\tilde{q}(\delta^{1/n}; A, B)$ .

*Remark.* The first part of Theorem 1 provides a finite sample valid estimate of the quantile. The second part implies that it is sharper than classical concentration inequalities such as Hoeffding, Bernstein, Bennett or Prokhorov inequalities. To see this fact, note that  $\mathbb{P}(\max_{1 \leq t \leq n} S_t \geq n\tilde{q}(\delta^{1/n}; A, B)) \leq \delta$  for all  $\delta \in [0, 1]$  is equivalent to the existence of a function  $H(u; A, B)$  such that  $\mathbb{P}(\max_{1 \leq k \leq n} S_t \geq nu) \leq H^n(u; A, B)$  for all  $u$ . The classical concentration inequalities mentioned above are all of this product from  $H^n(u; A, B)$  for some  $H$ , hence weaker than our bound.

### 3.2. Comparison to Classical bounds

Most of the classical concentration inequalities including Hoeffding, Bernstein, Bennett, or Prokhorov inequalities ([Bentkus, 2002](#); [Wellner, 2017](#)) are derived based on the



Cramér–Chernoff technique. The Cramér–Chernoff technique makes use of exponential moments unlike the positive part second moment used in Bentkus’ concentration inequality. We have mentioned in (7) that random variables  $G_i$ ’s defined in (6) is worst case for the positive part second moment. Interestingly, the same random variables are also worst case for exponential moments too, i.e., for all  $\lambda \geq 0$ ,

$$\mathbb{E}[\exp(\lambda \sum_{i=1}^n G_i)] = \sup_{X_i \sim (2)} \mathbb{E}[\exp(\lambda \sum_{i=1}^n X_i)].$$

See [Bennett \(1962, Page 42\)](#) for a proof. Hence, the optimal Cramér–Chernoff concentration inequality is given by

$$\mathbb{P}(\sum_{i=1}^n X_i \geq u) \leq \inf_{\lambda \geq 0} \frac{\mathbb{E}[\exp(\lambda \sum_{i=1}^n G_i)]}{\exp(\lambda u)}.$$

Furthermore, it can be proved that for all  $u \in \mathbb{R}$ ,

$$\tilde{P}_{2,n}(u) \leq \inf_{\lambda \geq 0} \frac{\mathbb{E}[\exp(\lambda \sum_{i=1}^n G_i)]}{\exp(\lambda u)}, \quad (12)$$

see Eqns (3)–(4), (9). This implies that Bentkus’ concentration inequality is sharper than the optimal Cramér–Chernoff inequality, and hence sharper than Hoeffding, Bernstein, Bennett, and Prokhorov inequalities. Inequality (12) only proves that Bentkus’ inequality is an improvement but does not show how significant the improvement is. In order to describe the improvement, let us denote the right hand side of (12) as  $\tilde{P}_{\infty,n}(u)$ . It can be proved that

$$1 \leq \limsup_{n \rightarrow \infty} \sup_{u \in \mathbb{R}} \frac{\tilde{P}_{\infty,n}(u)}{\mathbb{P}(\sum_{i=1}^n G_i \geq u)} = \infty. \quad (13)$$

See [Talagrand \(1995, Eq. \(1.4\)\)](#). Moreover,<sup>3</sup>

$$\mathbb{P}(\sum_{i=1}^n G_i \geq u) \stackrel{(i)}{\leq} \tilde{P}_{2,n}(u) \stackrel{(ii)}{\leq} \frac{e^2}{2} \mathbb{P}(\sum_{i=1}^n G_i \geq u). \quad (14)$$

Inequalities in (14) show that our concentration inequalities based on the two-point random variables  $G_i$  are sharp up to a constant factor  $e^2/2$ . Further, inequalities (13) and (14) show that there exists a distribution for which Bentkus’ inequality can be infinitely better than the optimal Cramér–Chernoff bound. See [Figure 1](#) for an illustration and [Bentkus \(2002; 2004\)](#); [Pinelis \(2006\)](#) for further discussion.

### 3.3. Computation of Bentkus’ bound

Computation of  $q(\cdot; n, A, B)$  is discussed in [Bentkus et al. \(2006, Section 9\)](#) and we provide a detailed discussion in [Appendix C](#). In this respect, the following result describes the function in (9) as a piecewise smooth function in homoscedastic case, i.e.,  $A_1 = \dots = A_n = A$ .

<sup>3</sup>Because  $G_i$ ’s satisfy assumption (2), inequality (i) is trivial using (8). Inequality (ii) holds for all  $u$  in the support of  $\sum_{i=1}^n G_i$ ; it holds for all  $u \in \mathbb{R}$  if  $\mathbb{P}(\sum_{i=1}^n G_i \geq u)$  is replaced by its logarithmic interpolation; see [Bentkus \(2002\)](#) for details.




Figure 2: Examples function  $P_2(x; Z_n)$  when  $n = 3$ ,  $A = 0.1$  and  $B = 1.0$ . We plot  $P_2(x; Z_n)$  in both linear (left) and log (right) scales on the y-axis.

**Proposition 1.** Set  $p_{AB} = A^2/(A^2 + B^2)$  and  $Z_n = \sum_{i=1}^n R_i$  where  $R_i \sim \text{Bernoulli}(p_{AB})$ . Then for  $u \in \mathbb{R}$ ,

$$\tilde{P}_{2,n}(u) = P_2(np_{AB} + u(1 - p_{AB})/B; Z_n),$$

where  $P_2(x; Z_n) = 1$  for  $x \leq np_{AB}$  and

$$P_2(x; Z_n) := \begin{cases} \frac{np_{AB}(1 - p_{AB})}{(x - np_{AB})^2 + np_{AB}(1 - p_{AB})}, & \text{if } np_{AB} < x \leq \Psi_0, \\ \frac{v_k p_k - e_k^2}{x^2 p_k - 2x e_k + v_k}, & \text{if } \Psi_{k-1} < x \leq \Psi_k, \\ \mathbb{P}(Z_n = n) = p_{AB}^n, & \text{if } x \geq \Psi_{n-1} = n. \end{cases}$$

Here  $p_k = \mathbb{P}(Z_n \geq k)$ ,  $e_k = \mathbb{E}[Z_n \mathbb{1}\{Z_n \geq k\}]$ ,  $v_k = \mathbb{E}[Z_n^2 \mathbb{1}\{Z_n \geq k\}]$ , and  $\Psi_k = (v_k - k e_k)/(e_k - k p_k)$ .

The function  $P_2(\cdot; Z_n)$  is illustrated in [Figure 2](#) for  $n = 3$  in both linear and logarithmic scale. Using [Proposition 1](#) and (10), computation of  $q(\cdot; n, A, B)$  follows. In [Appendix C.1](#), we also provide a similar piecewise description of  $q(\cdot; n, A, B)$ . It is worth pointing out that a similar expression for  $P_2(\cdot; Z_n)$  can be derived when  $A_i$ ’s are unequal. [Proposition 1](#) is stated for equal variances for simplicity and also because of the widely used i.i.d. assumption.

### 3.4. Adaptive Bentkus’ Concentration Inequality with Known Variance

Although [Theorem 1](#) leads to a uniform in sample size confidence sequence until size  $n$ , it is very wide for sample sizes much smaller than  $n$ . We now use the method of stitching to obtain a confidence sequence that is valid for all sample sizes and scales reasonably well with respect to the sample size. See [Mnih et al. \(2008, Section 3.2\)](#) and [Howard et al. \(2018, Section 3.1\)](#) for other applications. The stitching method requires two user-chosen parameters:

- a scalar  $\eta > 1$  that determines the geometric spacing.
- a function  $h : \mathbb{R}_+ \rightarrow \mathbb{R}_+$  such that  $\sum_{k=0}^{\infty} 1/h(k) \leq 1$ . Ideally,  $1/h(k)$ ,  $k \geq 0$  adds up to 1.




Figure 3: Comparison of the uniform concentration bounds when  $\delta = 0.05$ .  $X_i$  are centered i.i.d. Bernoulli( $\frac{1}{4}$ ). True standard deviation  $A_i = \sqrt{3}/4$  and upper bound  $B = 3/4$  are provided to all the methods. A-Bentkus is computed using  $\eta = 1.1$ ,  $h(k) = (k+1)^{1.1}\zeta(1.1)$ . For 3000 trials, there is zero failure for Adaptive Hoeffding and Empirical Bernstein, but 3 for A-Bentkus (15). All the bounds have failure frequency bounded above by  $\delta$  but the Bentkus' bound is the least conservative. The differences between the bounds continue to grow as  $n$  increases.

The following result (proved in Appendix E) gives a uniform over  $n$  tail inequality by splitting  $\{n \geq 1\}$  into  $\bigcup_{k \geq 1} \{\lceil \eta^k \rceil \leq n \leq \lfloor \eta^{k+1} \rfloor\}$  and then applying (11) within  $\{\lceil \eta^k \rceil \leq n \leq \lfloor \eta^{k+1} \rfloor\}$ . For each  $n \geq 1$ , set  $k_n := \min\{k \geq 0 : \lceil \eta^k \rceil \leq n \leq \lfloor \eta^{k+1} \rfloor\}$  and  $c_n := \lfloor \eta^{k_n+1} \rfloor$ .

**Theorem 2.** *If  $X_1, X_2, \dots$  are independent random variables satisfying (2), then for any  $\delta \in [0, 1]$ ,*

$$\mathbb{P}\left(\exists n \geq 1 : S_n \geq q\left(\frac{\delta}{h(k_n)}; c_n, \mathcal{A}, B\right)\right) \leq \delta. \quad (15)$$

The choice of the spacing parameter  $\eta$  and stitching function  $h(\cdot)$  determine the shape of the confidence sequence and there is no universally optimal setting. The growth rate of  $h(\cdot)$  determines how the budget of  $\delta$  is spent over sample sizes; a quickly growing  $h(\cdot)$  such as  $2^k$  yield confidence intervals of essentially 100% confidence for larger sample sizes. The choice of  $\eta$  determines how conservative the bound is for small  $n$  in  $\{\lceil \eta^k \rceil \leq n \leq \lfloor \eta^{k+1} \rfloor\}$ ; for  $\eta$  too large the bound will be conservative for  $n$  close to  $\eta^k$ . Eq. (14) shows that bound is tightest at  $n = \lfloor \eta^{k+1} \rfloor$  in each epoch. See Appendix B.1 for the graphical illustration. Throughout this paper, we use  $\eta = 1.1$  and  $h(k) = \zeta(1.1)(k+1)^{1.1}$  where  $\zeta(\cdot)$  is the Riemann zeta function.

The same stitching method used in Theorem 2 can also be used with Hoeffding and Bernstein inequalities as done in Zhao et al. (2016) and Audibert et al. (2009), respectively. However, given that inequality (11) is sharper than Hoeffding and Bernstein inequalities, our bound (15) is sharper for the same spacing parameter  $\eta$  and stitching function  $h(\cdot)$ ; see Figure 3. Stitched bounds as in Theorem 2 are always piecewise constant but the Hoeffding and Bernstein

versions from Zhao et al. (2016) and Mnih et al. (2008) are smooth because they are upper bounds of the piecewise constant boundaries (obtained using  $n \leq c_n \leq \eta n$  and  $k_n \leq \log_\eta n + 1$ ). For practical use, smoothness is immaterial and the piecewise constant versions are sharper.

### 3.5. Adaptive Bentkus Confidence Sequence with Estimated Variance

Theorem 2 is impractical in its form because it involves the unknown sequence of  $A_1, A_2, \dots$ . In the case where  $A_1 = A_2 = \dots = A$ , one needs to generate an upper bound of  $A$  (for a known  $B$ ) and obtain an actionable version of Theorem 2. Finite-sample over-estimation of  $A$  requires a two-sided bound on the  $X_i$ 's; one-sided bounds on the random variables do not suffice. This actionable version is a refined version of empirical Bernstein inequality that is uniform over the sample sizes.

We will assume that  $\mathbb{P}(\underline{B} \leq X_i \leq B) = 1, \forall i$ . It follows that  $\text{Var}(X_i) = A^2 \leq -\underline{B}\underline{B}$  (Bhatia & Davis, 2000). Because  $X_i$ 's have mean zero,  $\underline{B} \leq 0$  and  $B \geq 0$ ; this implies that  $-\underline{B}\underline{B} \geq 0$ . If one wants to avoid variance estimation, then one can use this upper bound in Theorem 2 to obtain an actionable confidence sequence. This sequence, however, will not have width scaling with the true variance.

Define  $\bar{A}_1(\delta) = (B - \underline{B})/2$  and for  $n \geq 2, \delta \in [0, 1]$

$$\begin{aligned} \hat{A}_n^2 &:= \lfloor n/2 \rfloor^{-1} \sum_{i=1}^{\lfloor n/2 \rfloor} (X_{2i} - X_{2i-1})^2/2, \\ \bar{A}_n(\delta) &:= \sqrt{\hat{A}_n^2 + g_{2,n}^2(\delta) + g_{2,n}(\delta)}, \end{aligned} \quad (16)$$

where  $g_{2,n}(\delta) := (2\sqrt{2}n)^{-1} \sqrt{\lfloor c_n/2 \rfloor} (B - \underline{B}) \times \Phi^{-1}(1 - 2\delta/(e^2 h(k_n)))$ , for the distribution function  $\Phi(\cdot)$  of a standard normal random variable. We will write  $\bar{A}_n(\delta; B, \underline{B})$ , when needed, to stress the dependence of  $\bar{A}_n(\delta)$  on  $B, \underline{B}$ . Lemma F.1 shows that  $\bar{A}_n(\delta)$  is a valid over-estimate of  $A$  uniformly over  $n$  and yields the following actionable bound. We defer the proof to Appendix F.

**Theorem 3.** *If  $X_1, X_2, \dots$  are mean-zero independent random variables satisfying  $\text{Var}(X_i) = A^2$  and  $\mathbb{P}(\underline{B} \leq X_i \leq B) = 1$  for all  $i \geq 1$ , then for any  $\delta_1, \delta_2 \in [0, 1]$ , with probability at least  $1 - \delta_1 - \delta_2$ , simultaneously for all  $n \geq 1$ ,*

$$S_n \leq q\left(\frac{\delta_1}{h(k_n)}; c_n, \bar{A}_n^*(\delta_2), B\right) \text{ and } A \leq \bar{A}_n^*(\delta_2, B, \underline{B}).$$

Similarly, with probability at least  $1 - \delta_1 - \delta_2$ , simultaneously for all  $n \geq 1$ ,

$$S_n \geq -q\left(\frac{\delta_1}{h(k_n)}; c_n, \bar{A}_n^*(\delta_2), -\underline{B}\right) \text{ and } A \leq \bar{A}_n^*(\delta_2, B, \underline{B}).$$

Here  $\bar{A}_n^*(\delta_2) := \min_{1 \leq s \leq n} \bar{A}_n(\delta_2, B, \underline{B})$ , and  $k_n, c_n$  are those defined before Theorem 2.

Theorem 3 is an analogue of the empirical Bernstein inequality Mnih et al. (2008, Eq. (5)). The over-estimate of  $A$  in (16) can be improved by using non-analytic expressions, but we present the version above for simplicity; see Appendix F for details on how to improve  $\bar{A}_n(\delta)$  in (16).

Theorem 3 can be used to construct a confidence sequence as follows. Suppose  $Y_1, Y_2, \dots$  are independent random variables with mean  $\mu$ , variance  $A^2$ , and satisfying  $\mathbb{P}(L \leq Y_i \leq U) = 1$ . Then  $X_i = Y_i - \mu$  is a zero mean random variable where  $\mathbb{P}(L - \mu \leq X_i \leq U - \mu) = 1$ , and Theorem 3 is directly applicable with  $B = -\underline{B} = U - L$ . An interesting observation is that we can refine the values of  $\underline{B}$  and  $B$  while we are updating the confidence interval for  $\mu$ . Suppose with  $n$  data points, we have:  $-q_n^{\text{low}} \leq n\bar{Y}_n - n\mu \leq q_n^{\text{up}}$ , then

$$\mu_n^{\text{low}} := \bar{Y}_n - n^{-1}q_n^{\text{up}} \leq \mu \leq \bar{Y}_n + n^{-1}q_n^{\text{low}} =: \mu_n^{\text{up}},$$

where  $\bar{Y}_n$  is the empirical mean of  $Y$ . We thus have a valid estimate  $[L - \mu_n^{\text{up}}, U - \mu_n^{\text{low}}]$  of the support of  $X$ , and when we observe  $Y_{n+1}$ , we can use  $U - \mu_n^{\text{low}}$  as  $B$  and  $L - \mu_n^{\text{up}}$  as  $\underline{B}$ . Importantly, as Theorem 3 provides a uniform concentration bound, these recursively defined upper and lower bounds hold simultaneously too. This leads to the following result, proved in Appendix G.

**Theorem 4.** *If random variables  $Y_1, Y_2, \dots$  are independent with mean  $\mu$ , variance  $A^2$  and satisfy  $\mathbb{P}(L \leq Y_i \leq U) = 1$ . Define  $\mu_0^{\text{up}} := U$ ,  $\mu_0^{\text{low}} := L$ , and for  $n \geq 1$*

$$\begin{aligned} \mu_n^{\text{up}} &= \bar{Y}_n + \frac{1}{n}q \left( \frac{\delta_1}{2h(k_n)}; c_n, \bar{A}_n^*(\delta_2, U, L), \mu_{n-1}^{\text{up}} - L \right) \\ \mu_n^{\text{low}} &= \bar{Y}_n - \frac{1}{n}q \left( \frac{\delta_1}{2h(k_n)}; c_n, \bar{A}_n^*(\delta_2, U, L), U - \mu_{n-1}^{\text{low}} \right) \end{aligned}$$


Let  $\mu_n^{\text{up}*} = \min_{0 \leq i \leq n} \mu_i^{\text{up}}$  and  $\mu_n^{\text{low}*} = \max_{0 \leq i \leq n} \mu_i^{\text{low}}$ . Then for any  $\delta_1, \delta_2 \in [0, 1]$ , with probability at least  $1 - \delta_1 - \delta_2$ , simultaneously for all  $n \geq 1$ ,

$$\mu \in [\mu_n^{\text{low}*}, \mu_n^{\text{up}*}] \quad \text{and} \quad A \leq \bar{A}_n^*(\delta_2, U, L). \quad (17)$$


Because  $\mu_0^{\text{up}} = U, \mu_0^{\text{low}} = L$ , the confidence intervals  $[\mu_n^{\text{low}*}, \mu_n^{\text{up}*}]$  is always a subset of  $[L, U]$ .

## 4. Experiments

We compare our adaptive Bentkus confidence sequence (17) with the adaptive Hoeffding (Zhao et al., 2016), empirical Bernstein (Mnih et al., 2008), and two other versions of empirical Bernstein inequality from (Howard et al., 2018): Eq. (24) and Theorem 4 with the gamma-exponential boundary. Eq. (24) of Howard et al. (2018) is a stitched confidence sequence, while Theorem 4 is a method of mixture confidence sequence. We denote these methods by A-Bentkus, A-Hoeffding, E-Bernstein, HRMS-Bernstein, and HRMS-Bernstein-GE, respectively. We use  $\delta =$



(a) One Replication



(b) Average Width over 1000 Replications

Figure 4: The 95% confidence sequences for the mean when  $Y_i \sim \text{Bernoulli}(0.1)$ . All the methods estimate the variance except A-Hoeffding. HRMS-Bernstein-GE involves a tuning parameter  $\rho$  which is chosen to optimize the sequence at  $n = 500$  as suggested in Howard et al. (2018). (a) shows the confidence sequences from a single replication. (b) shows the average widths of the confidence sequences over 1000 replications. The upper and lower bounds for all the other methods are cut at 1 and 0.

0.05 for all the experiments. For A-Bentkus, we fix the spacing parameter  $\eta = 1.1$ , the stitching function  $h(k) = (k+1)^{1.1}\zeta(1.1)$ , and  $\delta_1 = 2\delta/3, \delta_2 = \delta/3$ .

Section 4.1 examines the coverage probability and the width of the confidence intervals constructed on a synthetic data from Bernoulli(0.1); for other cases, see Appendix B. Section 4.2 and 4.3 apply the confidence sequences to an adaptive stopping algorithm for  $(\varepsilon, \delta)$ -mean estimation and the best arm identification problem.

### 4.1. Confidence Sequences for Bernoulli Variables

We generate samples  $Y_1, Y_2, \dots, Y_{20000} \stackrel{\text{i.i.d}}{\sim} \text{Bernoulli}(0.1)$  and compute the confidence sequences for  $\mu = 0.1$ . Figure 4a illustrates the confidence sequences obtained and shows the sharpness of A-Bentkus. For most of the cases ( $n \geq 20$ ), A-Bentkus dominates the other methods. For smaller sample sizes, A-Bentkus also closely traces A-Hoeffding and outperforms the others. This is because the variance estimation is likely conservative and in which case our  $\bar{A}_n^*$  ends up using the trivial upper bound

$(U - L)/2$ , which is essentially what `A-Hoeffding` is exploiting. In fact, we have provided the same upper bound for all the other Bernstein-type methods too, and `A-Bentkus` still outperforms. This phenomenon shows the intrinsic sharpness of our bound.

We repeat the above experiment 1000 times and report the average miscoverage rate:

$$\frac{1}{1000} \sum_{r=1}^{1000} \mathbb{1}\{\mu \notin \text{CI}_n^{(r)} \text{ for some } 1 \leq n \leq 20000\}.$$

where  $\text{CI}_n^{(r)}$  is the confidence interval constructed after observing  $Y_1, \dots, Y_n$  in the  $r$ -th replication. The results are 0.001 for `A-Bentkus`, 0.003 for `HRMS-Bernstein-GE`, and 0 for the others. All the methods control the miscoverage rate by  $\delta = 0.05$  but are all conservative. Recall from (14) that our failure probability bound can be conservative up to a constant of  $e^2/2$ . Furthermore, from the proofs of Theorems 2 and 4, we get that for  $\eta = 1.1$ ,  $h(k) = (k + 1)^{1.1} \zeta(1.1)$ ,

$$\begin{aligned} & \mathbb{P}(\mu \notin \text{CI}_n(\delta) \text{ for some } 1 \leq n \leq 20000) \\ & \leq \sum_{k=0}^{\log_{\eta}(20000)} \delta/h(k) \leq 0.41\delta. \end{aligned}$$

For  $\delta = 0.05$ ,  $0.41\delta = 0.0205$ . This explains why the average miscoverage rate in the experiment is small.

We also report the average width of the confidence intervals in Figure 4b. All the values are between 0 and 1 as we cut the bounds from above and below for the other methods. As mentioned above, when  $n$  is very small `A-Bentkus` closely traces `A-Hoeffding` and both have smaller width. Yet the advantage of `A-Hoeffding` disappears for  $n \geq 20$  and `A-Bentkus` enjoys smaller confidence interval width afterwards. `HRMS-Bernstein-GE` improves slightly on `A-Bentkus` after observing very large number of samples.

## 4.2. Adaptive Stopping for Mean Estimation

In this section, we apply our confidence sequence to adaptively estimate the mean of a bounded random variable  $Y$ . The goal is to obtain an estimator  $\hat{\mu}$  such that the relative error  $|\hat{\mu}/\mu - 1|$  is bounded by  $\varepsilon$ , and terminate the data sampling once such criterion is satisfied.

Given  $\bar{Y}$  the empirical mean and any confidence sequence centered at  $\bar{Y}$  satisfying (1), Algorithm 1 yields a valid stopping time and an  $(\varepsilon, \delta)$ -accurate estimator; see Mnih et al. (2008, Section 3.1) for a proof. Clearly, a tighter confidence sequence will require less data sampling and yields a smaller stopping time. We follow the setup in Mnih et al. (2008). The data samples are i.i.d generated as  $Y_i = m^{-1} \sum_{j=1}^m U_{ij}$ , where  $U_{ij}$  are i.i.d uniformly distributed in  $[0, 1]$ . This implies that  $\mu = \frac{1}{2}$  and  $A^2 = \frac{1}{12m}$ . Because Algorithm 1

---

### Algorithm 1: Adaptive Stopping Algorithm

---

**Initialization:**  $n \leftarrow 0$ ,  $\text{LB} \leftarrow 0$ ,  $\text{UB} \leftarrow \infty$   
**while**  $(1 + \varepsilon)\text{LB} < (1 - \varepsilon)\text{UB}$  **do**  
      $n \leftarrow n + 1$   
     Sample  $Y_n$  and compute the  $n$ -th CI in the sequence:  
          $[\bar{Y}_n - Q_n, \bar{Y}_n + Q_n] \leftarrow \text{ConfSeq}(n, \delta)$   
      $\text{LB} \leftarrow \max\{\text{LB}, |\bar{Y}_n| - Q_n\}$   
      $\text{UB} \leftarrow \min\{\text{UB}, |\bar{Y}_n| + Q_n\}$   
**return** stopping time  $N = n$  and estimator  
 $\hat{\mu} = (1/2)\text{sign}(\bar{Y}_N)[(1 + \varepsilon)\text{LB} + (1 - \varepsilon)\text{UB}]$

---




Figure 5: Comparison of confidence sequences for an  $(\varepsilon, \delta)$ -estimator. Here  $\varepsilon = 0.1$  and  $\delta = 0.05$ .

requires symmetric intervals, we shall symmetrize the intervals returned by `A-Bentkus` by taking the largest deviation. We consider 5 cases:  $m = 1, 10, 20, 100, 1000$  and report the average stopping time (i.e. the number of samples required to achieve  $(\varepsilon, \delta) = (0.1, 0.05)$  accuracy) based on 200 trials in Figure 5. `HRMS-Bernstein-GE` involves a tuning parameter  $\rho$ , chosen here to optimize the confidence sequence at  $n = 10$  (best out of 10, 50, 100, 200). As  $m$  increases, the variance of  $Y_i$  decreases. As expected, `A-Hoeffding` does not exploit the variance of random variables so the stopping times remains roughly the same. For others, the stopping time is decreasing. It is clear that on average, `A-Bentkus` is the best for all values of  $m$  and the ratios of our stopping time to the second best are 0.79, 0.66, 0.72, 0.86, 0.84.

## 4.3. Best Arm Identification

In this section, we study the fixed confidence best arm identification, a classic multi-arm bandit problem. An agent is presented with a set of  $K$  arms  $\mathbb{A}$ , with unknown expected rewards  $\mu_1, \dots, \mu_K$ . Sequentially, the agent pulls an arm  $\alpha \in \mathbb{A}$  of his choice and observes a reward value, until he finally claims one arm to have the largest expected reward. The goal is to correctly identify the best arm with fewer pulls  $N$ , i.e. smaller sample complexity. This problem has been extensively studied; see, e.g., Even-Dar et al. (2002); Karnin et al. (2013); Jamieson et al. (2013; 2014); Jamieson & Nowak (2014); Chen & Li (2015). Zhao et al. (2016) provided an algorithm based on `A-Hoeffding`



**Algorithm 2: Best Arm Identification****Input:** failure probability  $\delta$ , a set of arms  $\mathbb{A}$ **Initialization:**  $N \leftarrow 0$ ;  $n_\alpha \leftarrow 0$ ,  $\forall \alpha \in \mathbb{A}$ **while**  $\mathbb{A}$  has more than one arms **do**    Compute empirical mean reward  $\hat{\mu}_\alpha$ ,  $\forall \alpha \in \mathbb{A}$      $\hat{\alpha} \leftarrow \arg \max_{\alpha \in \mathbb{A}} \hat{\mu}_\alpha$     **for every arm**  $\alpha \in \mathbb{A}$  **do**

$$\delta_\alpha \leftarrow \begin{cases} \frac{\delta/2}{|\mathbb{A}|-1} & \text{if } \alpha = \hat{\alpha} \\ \frac{\delta}{2} & \text{otherwise} \end{cases}$$

 $[L_\alpha, U_\alpha] \leftarrow$  the  $n_\alpha$ -th CI of ConfSeq( $\delta_\alpha$ )         $R_\alpha \leftarrow$  radius of the  $n_\alpha$ -th CI of ConfSeq( $\delta_\alpha$ )    Sample from the arm  $\alpha$  with largest radius  $R_\alpha$      $n_\alpha \leftarrow n_\alpha + 1$ ,  $N \leftarrow N + 1$     Remove arm  $\alpha$  from  $\mathbb{A}$  if  $U_\alpha < L_{\hat{\alpha}}$ **return** the remaining arm in  $\mathbb{A}$ , number of pulls  $N$ 

that outperforms previous algorithms including LIL-UCB, LIL-LUCB. Here we present it as Algorithm 2 in a general form that utilizes any valid confidence sequences, and use A-Bentkus as well as the competing ones in it to compare their performance. Following the proof of Zhao et al. (2016, Theorem 5), one can show that Algorithm 2 outputs the best arm with probability at least  $1 - \delta$ .

The experiment setup follows Jamieson & Nowak (2014); Zhao et al. (2016). Each arm is generating random Bernoulli rewards with  $\mu_\alpha = 1 - (\frac{\alpha}{K})^{0.6}$ ,  $\alpha = 0, \dots, K - 1$ ; the first arm has highest expected reward  $\mu_0 = 1$ . The problem hardness is measured by  $\text{H1} = \sum_{\alpha \neq 0} (\mu_\alpha - \mu_0)^{-2}$  (Jamieson & Nowak, 2014), which is roughly  $0.4K^{1.2}$  in our setup.

In Algorithm 2, the sampling of an arm depends on  $R_\alpha$ , the radius of the confidence interval. In our experiments, we find that a confidence sequence for which  $R_\alpha$  stays constant for a stretch of samples yields a larger sample complexity. Our intuition is that Algorithm 2 keeps selecting the same arm when the radius is not updated, therefore it forgoes a number of samples; see Appendix B.3 for more details. This phenomenon happens for all confidence sequences when truncated to  $[0, 1]$ , where the intervals stay constant at  $[0, 1]$  for the first few samples, see Figure 4a. For A-Bentkus, the cumulative maximum/minimum ( $\mu_n^{\text{low}*}$  and  $\mu_n^{\text{up}*}$  in Theorem 4) also leads to the constant radius problem. Hence, for smaller sample complexity, we set  $L_\alpha = \mu_{n_\alpha}^{\text{low}*}$ ,  $U_\alpha = \mu_{n_\alpha}^{\text{up}*}$  and  $R_\alpha = (\mu_{n_\alpha}^{\text{up}*} - \mu_{n_\alpha}^{\text{low}*})/2$  instead of  $(\mu_{n_\alpha}^{\text{low}*} - \mu_{n_\alpha}^{\text{up}*})/2$ .

Our experiments are reported in Figure 6. A-Bentkus significantly outperforms the competing approaches, including A-Hoeffding which beats LIL-UCB, LIL-LUCB (Zhao et al., 2016). Further, A-Bentkus only requires 52% to 61% of the samples required by A-Hoeffding. Finally, we note that the Bernstein type of methods underperform because they have larger confidence intervals for small to moderate number of samples as can be seen in Figure 4a.




Figure 6: The number of pulls  $N$  versus the the number of arms  $K$ .  $\delta = 0.05$ . The results are averaged over 10 trials.

## 5. Conclusion

We proposed a confidence sequence for bounded random variables and examined its efficacy in synthetic examples and applications. Our method is favorable to methods that utilize classical concentration results. It can be applied to various problems for improved performance, including testing equality of distributions, testing independence (Balsubramani & Ramdas, 2016), etc. Our work can be extended in a few future directions. We assumed that  $X_i$ 's are independent and bounded. The generalizations for the dependent case and/or the sub-Gaussian case are of interest. The generalized results can be obtained based on the results of Pinelis (2006, Theorem 2.1) and Bentkus (2010).

## References

- Abramowitz, M. and Stegun, I. A. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*, volume 55. US Government printing office, 1948.
- Audibert, J.-Y., Munos, R., and Szepesvári, C. Exploration–exploitation tradeoff using variance estimates in multi-armed bandits. *Theoretical Computer Science*, 410(19): 1876–1902, 2009.
- Bahadur, R. R. and Savage, L. J. The nonexistence of certain statistical procedures in nonparametric problems. *The Annals of Mathematical Statistics*, 27(4):1115–1122, 1956.
- Balsubramani, A. and Ramdas, A. Sequential nonparametric testing with the law of the iterated logarithm. In *Proceedings of the Thirty-Second Conference on Uncertainty in Artificial Intelligence*, pp. 42–51, 2016.
- Bennett, G. Probability inequalities for the sum of independent random variables. *Journal of the American Statistical Association*, 57(297):33–45, 1962.
- Bentkus, V. A remark on the inequalities of Bernstein, Prokhorov, Bennett, Hoeffding, and Talagrand. *Liet. Mat. Rink*, 42(3):332–342, 2002.

- Bentkus, V. On Hoeffding's inequalities. *The Annals of Probability*, 32(2):1650–1673, 2004.
- Bentkus, V. An extension of the hoeffding inequality to unbounded random variables. *Lithuanian Mathematical Journal*, 48(2):137–157, 2008.
- Bentkus, V. Bounds for the stop loss premium for unbounded risks under the variance constraints. *Preprint. Available at <https://www.math.uni-bielefeld.de/sfb701/preprints/view/423>*, 2010.
- Bentkus, V., Kalosha, N., and Van Zuijlen, M. On domination of tail probabilities of (super) martingales: explicit bounds. *Lithuanian Mathematical Journal*, 46(1):1–43, 2006.
- Bhatia, R. and Davis, C. A better bound on the variance. *The american mathematical monthly*, 107(4):353–357, 2000.
- Burgess, M. A., Chapman, A. C., and Scott, P. An engineered empirical bernstein bound. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 86–102. Springer, 2019.
- Chen, L. and Li, J. On the optimal sample complexity for best arm identification. *arXiv preprint arXiv:1511.03774*, 2015.
- Dagum, P., Karp, R., Luby, M., and Ross, S. An optimal algorithm for Monte Carlo estimation. *SIAM Journal on computing*, 29(5):1484–1496, 2000.
- Darling, D. A. and Robbins, H. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences of the United States of America*, 58(1):66, 1967.
- Even-Dar, E., Mannor, S., and Mansour, Y. Pac bounds for multi-armed bandit and markov decision processes. In *International Conference on Computational Learning Theory*, pp. 255–270. Springer, 2002.
- Fan, X., Grama, I., and Liu, Q. The missing factor in bennett's inequality. *arXiv preprint arXiv:1206.2592*, 2012.
- Giné, E. and Nickl, R. *Mathematical foundations of infinite-dimensional statistical models*, volume 40. Cambridge University Press, 2016.
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. Uniform, nonparametric, non-asymptotic confidence sequences. *arXiv preprint arXiv:1810.08240*, 2018.
- Huber, M. An optimal  $(\epsilon, \delta)$ -randomized approximation scheme for the mean of random variables with bounded relative variance. *Random Structures & Algorithms*, 55(2):356–370, 2019.
- Jamieson, K. and Nowak, R. Best-arm identification algorithms for multi-armed bandits in the fixed confidence setting. In *2014 48th Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6. IEEE, 2014.
- Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. On finding the largest mean among many. *arXiv preprint arXiv:1306.3917*, 2013.
- Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. LIL-UCB: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pp. 423–439. PMLR, 2014.
- Johari, R., Pekelis, L., and Walsh, D. J. Always valid inference: Bringing sequential analysis to A/B testing. *arXiv preprint arXiv:1512.04922*, 2015.
- Johari, R., Koomen, P., Pekelis, L., and Walsh, D. Peeking at A/B tests: Why it matters, and what to do about it. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 1517–1525, 2017.
- Karnin, Z., Koren, T., and Somekh, O. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pp. 1238–1246. PMLR, 2013.
- Maurer, A. and Pontil, M. Empirical bernstein bounds and sample variance penalization. *arXiv preprint arXiv:0907.3740*, 2009.
- Mnih, V., Szepesvári, C., and Audibert, J.-Y. Empirical Bernstein stopping. In *Proceedings of the 25th international conference on Machine learning*, pp. 672–679, 2008.
- Peña, V. H., Lai, T. L., and Shao, Q.-M. *Self-normalized processes: Limit theory and Statistical Applications*. Springer Science & Business Media, 2008.
- Philips, T. K. and Nelson, R. The moment bound is tighter than Chernoff's bound for positive tail probabilities. *The American Statistician*, 49(2):175–178, 1995.
- Pinelis, I. Binomial upper bounds on generalized moments and tail probabilities of (super) martingales with differences bounded from above. In *High dimensional probability*, pp. 33–52. Institute of Mathematical Statistics, 2006.
- Pinelis, I. On the Bennett-Hoeffding inequality. *arXiv preprint arXiv:0902.4058*, 2009.

- Pinelis, I. Optimal binomial, poisson, and normal left-tail domination for sums of nonnegative random variables. *Electronic Journal of Probability*, 21, 2016.
- Singh, R. Existence of bounded length confidence intervals. *The Annals of Mathematical Statistics*, 34(4):1474–1485, 1963.
- Talagrand, M. The missing factor in Hoeffding’s inequalities. In *Annales de l’IHP Probabilités et statistiques*, volume 31, pp. 689–702, 1995.
- Victor, H., Klass, M. J., Lai, T. L., et al. Pseudo-maximization and self-normalized processes. *Probability Surveys*, 4:172–192, 2007.
- Wald, A. *Sequential analysis*. Courier Corporation, 2004.
- Wellner, J. A. The bennett-orlicz norm. *Sankhya A*, 79(2): 355–383, 2017.
- Yang, F., Ramdas, A., Jamieson, K. G., and Wainwright, M. J. A framework for multi-A (rmed)/B (andit) testing with online fdr control. In *Advances in Neural Information Processing Systems*, pp. 5957–5966, 2017.
- Zhao, S., Zhou, E., Sabharwal, A., and Ermon, S. Adaptive concentration inequalities for sequential decision problems. In *Advances in Neural Information Processing Systems*, pp. 1343–1351, 2016.