
Debiasing Samples from Online Learning Using Bootstrap

Ningyuan Chen*

The Rotman School of Management
University of Toronto
ningyuan.chen@utoronto.ca

Xuefeng Gao*

Department of Systems Engineering
and Engineering Management
The Chinese University of Hong Kong
xfgao@se.cuhk.edu.hk

Yi Xiong*

Department of Systems Engineering
and Engineering Management
The Chinese University of Hong Kong
yxiong@se.cuhk.edu.hk

Abstract

It has been recently shown in the literature (Nie et al., 2018; Shin et al., 2019a,b) that the sample averages from online learning experiments are biased when used to estimate the mean reward. To correct the bias, off-policy evaluation methods, including importance sampling and doubly robust estimators, typically calculate the conditional propensity score, which is ill-defined for non-randomized policies such as UCB. This paper provides a procedure to debias the samples using bootstrap, which doesn't require the knowledge of the reward distribution and can be applied to any adaptive policies. Numerical experiments demonstrate the effective bias reduction for samples generated by popular multi-armed bandit algorithms such as Explore-Then-Commit (ETC), UCB, Thompson sampling (TS) and ϵ -greedy (EG). We analyze and provide theoretical justifications for the procedure under the ETC algorithm, including the asymptotic convergence of the bias decay rate in the real and bootstrap worlds.

1 Introduction

Online learning and specifically the multi-armed bandit problem, has seen great success in many applications, including recommender systems (Li et al., 2011), clinical trials (Villar et al., 2015), dynamic pricing (Den Boer, 2015), and A/B testing or online experiments (Burtini et al., 2015). The primary goal of online learning is to explore available options (arms) while

maximizing the payoff (or equivalently minimize the regret) at the same time. Many algorithms or policies achieve the goal remarkably well, including UCB (Upper Confidence Bound) (Lai and Robbins, 1985; Auer et al., 2002) and Thompson sampling (Thompson, 1933). After the conclusion of the experiment and the collection of the data, typically additional evaluations need to be conducted based on the collected data without running new experiments. This is usually referred to as “off-policy evaluation” (Precup et al., 2000; Li et al., 2011, 2012; Jiang and Li, 2016).

Motivating Example. To illustrate the motivation, consider the following hypothetical example in early-stage clinical trials. The online learning framework has been applied to Phase I/II trials to identify the optimal dose (Villar et al., 2015; Chen and Khademi, 2020; Aziz et al., 2021). Suppose a clinical trial is conducted on 100 patients sequentially to evaluate the efficacy of two drugs, A and B . This can be treated as a multi-armed bandit problem with two arms. The manager applies the UCB algorithm to allocate the drugs to the patients while observing their responses in terms of health measures. Eventually, the responses of the 100 patients are recorded, among which 90 patients are allocated drug A and the remaining 10 patients are allocated drug B . Apparently, the UCB algorithm determines that drug A is more efficacious than B . Suppose, in addition, the manager wants to evaluate the mean efficacy of the inferior drug B without running new trials. One natural step is to take the sample average of the responses from the 10 patients who have been allocated drug B . Does the quantity reflect the actual efficacy of drug B ?

The answer is no, which has been documented in a few recent papers (Nie et al., 2018; Shin et al., 2019b,a). To see the intuition, note that UCB tends to choose the arm with a higher empirical average. Therefore, if during the trial the sample average of drug B is less than its actual mean efficacy, then UCB tends to choose drug A more often and the realized negative

bias is not likely to be fully corrected. On the other hand, if the sample average happens to be higher than the actual mean, then UCB tends to choose drug B relatively more and correct the realized positive bias. Due to the asymmetry, the sample average of drug B is typically negatively biased.

To correct the bias, the literature on off-policy evaluations provides a few potential methods (Li et al., 2015; Swaminathan and Joachims, 2015; Jiang and Li, 2016; Dimakopoulou et al., 2017; Farajtabar et al., 2018; Hadad et al., 2021), including importance sampling and doubly robust estimators. Embedded in almost all the methods, there is a key concept called the propensity score, i.e., the probability that an arm is chosen in a period under the bandit algorithm. To estimate the unconditional propensity scores, the distribution of the rewards of the arms is required because of the adaptive nature of most bandit algorithms. However, the lack of prior knowledge of the rewards, not even their means, is precisely the reason to resort to online learning in the first place. For conditional propensity scores, they are readily available for policies with internal randomization such as TS and EG (Hadad et al., 2021). However, for other popular policies such as ETC and UCB which are non-randomized, the methods do not work and some adjustments have to be used (Dimakopoulou et al., 2017).

Our Contribution. To overcome the challenge, we introduce a simulation-based algorithm to debias the sample means and provide more reliable estimators for the mean rewards of the arms. At a high level, the algorithm can be described in the following steps:

1. Based on the collected data from the experiment in the real world, we construct bootstrap reward distributions, P_1^*, \dots, P_K^* , for the K arms. The bootstrap distribution could be Efron’s bootstrap (Efron and Tibshirani, 1994) or multiplier bootstrap (van Der Vaart and Wellner, 1996; Chernozhukov et al., 2013).
2. Simulate independent experiments in the “bootstrap” world using the same policy as in the real world, say, UCB. The rewards of the arms in the bootstrap world are generated using $\{P_k^*\}_{k=1}^K$.
3. For a bootstrap experiment, the difference between the sample average of the rewards of arm k and the mean reward in the “bootstrap” world $E_{P_k^*}[X]$, which is the sample mean of arm k in the real-world experiment, provides a realization of the bias in the “bootstrap” world.
4. Average the realized bias of all the bootstrap experiments and use it to debias the data in the real world.

The algorithm has the following features. First, it doesn’t require any knowledge of the probability model that generates the rewards. This is the benefit of using bootstrap to construct data-dependent distributions. Second, it is not specific to certain bandit algorithms. As long as the bandit algorithm that generates the data in the real world can be replicated in the bootstrap world, the debiasing procedure works. In Section 3, we apply the procedure to four popular bandit algorithms: Explore-Then-Commit (ETC), UCB, Thompson sampling, and ϵ -greedy. The corrected bias is always smaller, demonstrating the effectiveness of the procedure.

We also provide theoretical justifications when applying the approach to the ETC algorithm. In general, the exact analytical understanding of adaptive algorithms is difficult. We follow the literature (Boutilier et al., 2020; Perchet et al., 2016) and analyze ETC in the two-armed bandit problem. Specifically, for the ETC algorithm, we theoretically characterize the biases in the real world and the bootstrap world when there are two arms and the rewards are Gaussian. This allows us to evaluate how well the bootstrap bias approximates the real-world bias. In particular, we show that:

Informal Result: the bias of the real and bootstrap worlds decays exponentially as the length of the exploration phase grows. The ratio of their decay rates converges to one.

When the rewards are not Gaussian, we show that the ratio of the decay rate may depend on the tail behavior of the reward distribution, in particular, the Legendre-Fenchel transform of the reward distribution. We provide an asymptotic bound for the ratio.

Finally, we point out that in contrast to existing studies (e.g. Shin et al. (2019a)) which provide bounds on the bias for general bandit algorithms, we are able to give precise estimates of the bias under the ETC algorithm.

Related Work. In online learning problems, reward samples are collected in an adaptive manner. The dependence usually leads to bias in the sample average. This phenomenon is empirically documented in (Xu et al., 2013; Villar et al., 2015). Theoretically, Nie et al. (2018) give sufficient conditions under which the sample mean is negative. Shin et al. (2019b,a) provide a thorough analysis of the magnitude of the bias in multi-armed bandit problems.

How to debias the sample mean and learn the actual mean of the reward has attracted attention in recent literature. Xu et al. (2013) proposes a debiasing method by collecting another data log or a “held-out” set. Collecting additional data improves the estimation in most cases, but it is not always feasible in many applications

such as clinical trials. In comparison, our method only relies on the bandit experiment without additional data. In Nie et al. (2018), the authors provide an MLE-based approach that models the whole stochastic process and find the parameter (mean) that yields the maximum likelihood. Their approach requires the knowledge of the reward distribution to construct the likelihood function. Neel and Roth (2018) shows that the bias problem is mitigated if the data collection procedure is differentially private. Deshpande et al. (2018) introduces a ridge-type debiased estimator of the ordinary least squares estimator, which is computed adaptively to optimize bias-variance tradeoff. The debiasing procedure requires access to the data collection policy to construct a reasonable regularization parameter of the ridge estimation. Another line of literature uses the framework of causal inference. In particular, important sampling or propensity scoring has been used in Li et al. (2015); Swaminathan and Joachims (2015); Schnabel et al. (2016); Dimakopoulou et al. (2017); Vlassis et al. (2019); Hadad et al. (2021) to develop off-policy evaluation schemes. In order to evaluate the propensity score, the method typically only works for bandit algorithms with internal randomization such as Thompson sampling.

Our work builds on the bootstrap method, introduced by Efron (1979). Due to the data-dependent nature, bootstrap usually doesn't require any knowledge of the distribution itself and is well suited for multi-armed bandit problems. For example, Hao et al. (2019) replaces the confidence bound in UCB with the bootstrap version to circumvent the dependence on model parameters such as the subgaussian parameter of the reward distribution. (Eckles and Kaptein, 2014; Osband and Van Roy, 2015) use bootstrap in the posterior distribution of Thompson sampling to improve the computational efficiency. In addition, the bootstrap can be used to learn model coefficients in contextual bandits (Tang et al., 2015), achieve near-optimal regret (Vaswani et al., 2018), approximate Thompson sampling (Elmachtoub et al., 2017) and conduct a generally well-performed algorithm in different models (Kveton et al., 2019). These papers apply bootstrap as a component in the online algorithm, which is different from our work. Finally, bootstrap has also been used in the offline setting or off-policy evaluation. (Kostrikov and Nachum, 2020; Hao et al., 2021) propose a bootstrap Q-evaluation to evaluate the value of a target policy. Bibaut et al. (2019) obtains the confidence interval for its proposed doubly-robust estimator by bootstrapping. Mary et al. (2014) uses bootstrapping techniques to improve offline evaluation of contextual bandit algorithms in recommendation applications. In our paper, we focus on the bias of the sample average reward collected from a target multi-armed bandit algorithm and use

bootstrap to simulate additional bandit experiments so that the bias can be evaluated.

Notations. We use K as the number of arms and T as the number of rounds in the bandit experiment. We use \cdot^* for the corresponding quantity in the bootstrap world, conditional on the samples.

2 Problem Formulation

Suppose an agent has collected a dataset after running a bandit experiment. In particular, the agent has observed the number of rounds T , the action sequence $\{a_1, \dots, a_T\}$ where $a_t \in [K] = \{1, \dots, K\}$ is the arm pulled in round t , and the reward sequence $\{r_1, \dots, r_T\}$ where $r_t \in \mathbb{R}$ is the (noisy) reward observed in round t . In addition, the agent is also aware of the bandit algorithm that has been implemented, i.e., the distribution of a_t given $\mathcal{F}_{t-1} \triangleq \sigma(a_1, r_1, \dots, a_{t-1}, r_{t-1})$. However, the reward distributions $P = \{P_1, \dots, P_K\}$ are unknown, in particular, the mean μ_k for each arm k is unknown.

The agent is interested in the average reward of, say, arm $k \in [K]$. A straightforward approach is to take the sample average of the rewards generated from arm k : $\hat{\mu}_k = \frac{\sum_{t=1}^T r_t \mathbb{I}_{\{a_t=k\}}}{\sum_{t=1}^T \mathbb{I}_{\{a_t=k\}}}$. As documented in the literature (Nie et al., 2018; Shin et al., 2019b,a), such a statistic is usually biased, especially for suboptimal arms. That is, the bias given by $\mathbb{E}_P[\hat{\mu}_k] - \mu_k$ is often negative, where μ_k is the mean of the reward of arm k . However, the quantity depends on the algorithm and its interaction with the unknown P . To show the intractability of this quantity, consider $\mathbb{E}_P[\sum_{t=1}^T \mathbb{I}_{\{a_t=k\}}]$, the expected value of the total number of pulls of arm k under a particular algorithm. This seemingly simple quantity usually doesn't have a closed-form expression even when P is known, for most algorithms such as UCB and Thompson sampling. It is unclear if we can estimate the bias and use it to debias the sample averages.

Next, we propose a procedure to estimate the quantity based on bootstrap (Efron and Tibshirani, 1994) and in particular, multiplier bootstrap (van Der Vaart and Wellner, 1996; Chernozhukov et al., 2013).

2.1 Multiplier Bootstrap

In this section, we first introduce the idea of multiplier bootstrap generally.

Let $\mathcal{X} = \{z_1, \dots, z_n\}$ be i.i.d. observations from an unknown distribution F , whose mean is μ . Then the sample mean and sample variance can be calculated by $\hat{\mu} = \frac{z_1 + \dots + z_n}{n}$ and $\hat{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n (z_i - \hat{\mu})^2$. Conditioned on \mathcal{X} , the goal of bootstrap is to construct and simulate samples from F approximately, say $\mathcal{X}^* = \{z_1^*, \dots, z_m^*\}$,

without knowing F .

There are two bootstrap approaches involved in this study. Efron’s bootstrap (EB) (Efron, 1982) draws each sample z_j^* uniformly randomly from \mathcal{X} with replacement. That is $P(z_j^* = z_i) = \frac{1}{n}, \forall i = 1, \dots, n$. Equivalently, one can think of $z_j^* = \sum_{i=1}^n w_{ij} z_i$, where the weight (w_{1j}, \dots, w_{nj}) has a multinomial distribution with parameters $(1; 1/n, \dots, 1/n)$.

Another bootstrap approach, which is the focus of this paper, is referred to as multiplier bootstrap (MB) (van Der Vaart and Wellner, 1996). It has attracted the attention of many scholars due to its analytical tractability (Arlot et al., 2010; Chernozhukov et al., 2013; Hao et al., 2019). In multiplier bootstrap, we have $z_j^* = \frac{1}{\sqrt{n}} \sum_{i=1}^n (z_i - \hat{\mu}) w_{ij} + \hat{\mu}$, where the weight (w_{1j}, \dots, w_{nj}) is a random vector with $E[w_{ij}] = 0$ and $\text{Var}[w_{ij}] = 1$. It is clear that $E[z_j^*] = \hat{\mu}$ and $\text{Var}[z_j^*] = \hat{\sigma}^2$, so MB preserves the sample mean and variance.

In our problem, the data \mathcal{X} consists of the action sequence $\{a_1, \dots, a_T\}$ and the reward sequence $\{r_1, \dots, r_T\}$. We use MB to construct the distribution of the reward of each arm with Gaussian weights, *i.e.*, $w_{i,j} \sim N(0, 1)$. Therefore, conditional on \mathcal{X} , a reward from arm k can be generated by

$$P_k^* \sim N(\hat{\mu}_k, \hat{\sigma}_k^2), \quad (1)$$

where $\hat{\mu}_k$ and $\hat{\sigma}_k$ are the sample mean and variance of the observed rewards from arm k .

2.2 Use Multiplier Bootstrap for Debiasing

It is not a new idea to use bootstrap to correct the bias in the samples (Efron and Tibshirani, 1994; Steck and Jaakkola, 2003). But because of the structure of this application (the sequential nature of the observations), it is not straightforward to apply bootstrap. We next elaborate on the major steps and intuition of the algorithm. After collecting the data from the experiment, we simulate B additional experiments in the bootstrap world, using the same K , T , and the bandit algorithm. The reward of the bootstrap experiments is generated according to (1). Comparing the experiments in the bootstrap and the real world, the bootstrap reward distribution only offers an approximation. However, because the implemented bandit algorithm is known, the bootstrap experiments correctly capture the intertemporal dependence introduced by the sequential decision-making. Therefore, the bias calculated in the bootstrap world may be used to estimate the bias in the real world, thanks to the known reward distribution in the bootstrap world. Algorithm 1 demonstrates the process.

Algorithm 1 Using MB to debias the samples from a bandit experiment

Require: The horizon T , the number of arms K , the action sequence $\{a_1, \dots, a_T\}$, the reward sequence $\{r_1, \dots, r_T\}$, the implemented bandit algorithm

- 1: Parameters: bootstrap sample size B
- 2: **for** $b = 1, \dots, B$ **do**
- 3: Simulate a bandit experiment with the same K , T and bandit algorithm in the bootstrap world, where the reward of arm k is generated by (1)
- 4: Record the action sequence $\{a_{b,1}^*, \dots, a_{b,T}^*\}$ and the reward sequence $\{r_{b,1}^*, \dots, r_{b,T}^*\}$
- 5: Calculate the sample average of the reward of the arms in the experiment

$$\hat{\mu}_{b,k}^* = \frac{\sum_{t=1}^T r_{b,t}^* \mathbb{I}\{a_{b,t}^* = k\}}{\sum_{t=1}^T \mathbb{I}\{a_{b,t}^* = k\}}, \quad \forall k = 1, \dots, K$$

- 6: **end for**
 - 7: Calculate the average bias in the bootstrap world:
Estimated Bias = $\frac{1}{B} \sum_{b=1}^B \hat{\mu}_{b,k}^* - \hat{\mu}_k$
 - 8: **return** the debiased sample average as the estimator for $\mu_k, \forall k = 1, \dots, K$:
 $\hat{\mu}_k - \text{Estimated Bias} = 2\hat{\mu}_k - \frac{1}{B} \sum_{b=1}^B \hat{\mu}_{b,k}^*$
-

3 Numerical Performance

In this section, we test the performance of Algorithm 1 for bandit experiments using four popular bandit algorithms, including Explore-Then-Commit (ETC), upper confidence bound (UCB), Thompson sampling (TS) and ϵ -greedy (EG), for $T = 100$ and $K = 2$ (also $K = 4$). We consider two reward distributions: unit-variance normal distribution with $\mu_1 = 1, \mu_2 = 1.5$ and Bernoulli random variables with $\mu_1 = 0.3, \mu_2 = 0.6$. The detailed implementations of the algorithms are given below.

- ETC: pull each arm $m = 10$ times sequentially in the first mK rounds. Choose the arm with the highest average reward for the rest of the horizon.
- UCB1: we use the version in Chapter 2.7 of Sutton and Barto (2018). In particular, in round t , the arm that maximizes $\hat{\mu}_k(t) + \sqrt{\log t / N_k(t)}$ is pulled, where $\hat{\mu}_k(t)$ is the empirical average reward of arm k and $N_k(t)$ is the number of rounds arm k is pulled prior to t . If $N_k(t) = 0$, then k is considered to be a maximizing action/arm.
- TS: pull the arm according to its posterior probability of being the optimal arm (Russo et al., 2017). We set the prior distribution to be a normal distribution.

- EG: pull the arm with the highest empirical average reward so far (greedy) with probability $1 - \epsilon$ and selects a uniformly random arm with probability ϵ . We set $\epsilon = 0.05$.

We consider $B = 1,000$ bootstrap simulations in Algorithm 1 when calculating the debiased sample average in Step 8. To evaluate the bias in the real world, $\mathbb{E}[\hat{\mu}_k] - \mu_k$, we run the bandit experiment $L = 1000$ times and take the average. We also report the debiased sample average using MB or EB. The numerical experiments are conducted on a PC with 3.10 GHz Intel Processor and 16 GB of RAM.

Moreover, we implement the two propensity score based methods in Hadad et al. (2021): the inverse propensity score weighted (IPW) estimator and the augmented inverse propensity weighted (AIPW) estimator, for EG and TS. Note that both estimators are unbiased: if we run a large number of bandit experiments and take the average bias, it is going to be zero. In addition, both estimators only work for bandit algorithms with internal randomization. Therefore, we compare the mean-squared error (MSE) of our method and these two methods for EG and TS.

Table 1 demonstrates the performance of Algorithm 1. The column “ $\mathbb{E}[\hat{\mu}_k] - \mu_k$ ” reports the bias of arm k in the real world, and it is estimated by the sample average $\frac{1}{L} \sum_{l=1}^L (\hat{\mu}_{l,k} - \mu_k)$, where $\hat{\mu}_{l,k}$ denotes the sample mean of arm k in the l -th bandit experiment in the real world. The “Estimated Bias” in the bootstrap world in Table 1 is calculated as follows. For each of the l -th bandit experiment, we use $B = 1000$ bootstrap simulations to apply Algorithm 1 and calculate the bias in the bootstrap world by $\frac{1}{B} \sum_{b=1}^B \hat{\mu}_{l,b,k}^* - \hat{\mu}_{l,k}$, where $\hat{\mu}_{l,b,k}^*$ denotes the sample average reward of arm k in the b -th bootstrap simulation (see Lines 5–7 of Algorithm 1). Then the “Estimated Bias” in the bootstrap world is computed by $\frac{1}{L} \sum_{l=1}^L \left(\frac{1}{B} \sum_{b=1}^B \hat{\mu}_{l,b,k}^* - \hat{\mu}_{l,k} \right)$. Finally, the “Corrected $\mathbb{E}[\hat{\mu}_k]$ ” in the bootstrap world is estimated by $\frac{1}{L} \sum_{l=1}^L \left(2\hat{\mu}_{l,k} - \frac{1}{B} \sum_{b=1}^B \hat{\mu}_{l,b,k}^* \right)$ (c.f. Line 8 of Algorithm 1), which shows the average debiasing effect. In all the instances, Algorithm 1 is able to give a more accurate estimator for the mean reward of each arm. Although there are occasions when the bias is already quite small and Algorithm 1 over-corrects, Algorithm 1 is very effective in achieving bias reduction overall.

In Figures 1, 2, 3 and 4, we provide more information about the performance of Algorithm 1. In particular, we show the histogram of the raw sample mean as well as MB corrected sample means among the 1,000 bandit experiments. The average is illustrated by the dashed vertical lines and the actual mean reward by the solid

vertical lines. Again, after applying MB to the sample means, the bias is significantly reduced.

To compare to IPW and AIPW, Figures 5 and 6 illustrate the MSE of the corrected mean rewards among the 1000 bandit experiments. The MSE of all methods decrease as the horizon T increases although the MSE of MB appear to be constants. Actually it also decreases by around 10% when T increases from 25 to 100 in the figures. However, because the magnitude of the MSE of MB is smaller than that of IPW and AIPW, it becomes difficult to observe this decreasing pattern when the MSE of three methods are plotted in the same figure. When T is small (e.g. $T = 25$), there is some fluctuation of the MSE curves because the number of samples is too small, and in this case we can find that the MSE of IPW and AIPW is much larger compared with the MSE of MB. A large variance is a common issue for propensity score based methods. By trading off a little bias for variance reduction, Algorithm 1 seems to be more robust for small T . When $T = 100$, all methods achieve a similar level of MSE. We also remark that when T becomes very large (e.g. $T \geq 10000$), the MSE of all three estimators become very small and close to zero.

We also briefly discuss the effect of the bootstrap size B on the bias estimation and correction. For simplicity, we present results for the ETC algorithm with Gaussian rewards in the two-armed bandit problem. Below we report the corrected mean estimators for the two arms. From Table 1, the corrected means are $[0.9859, 1.5072]$ when $B = 1000$. In our experiments we find that the corrected means are $[0.9847, 1.5073]$ when $B = 50$, $[0.9853, 1.5058]$ when $B = 100$, and $[0.9854, 1.5075]$ when $B = 500$. This shows that when B is of moderate size, our algorithm already corrects the bias very well and produces results that are very close to the true means $[1, 1.5]$.

Finally, we also conduct additional bandit experiments with $K = 4$ arms: Gaussian reward with mean $\mu = \{2, 2.5, 3, 3.5\}$ and standard deviation $\sigma = \{2, 1, 2, 1\}$, and Bernoulli reward with mean $\mu = \{0.4, 0.5, 0.7, 0.8\}$. The four algorithms are repeated 1,000 times, each instance with $B = 1,000$ bootstrap simulations for both Efron’s bootstrap and Gaussian multiplier bootstrap. The results are presented in Appendix A. They further demonstrate the effectiveness of the proposed algorithm.

4 Theoretical Results for ETC

In this section, we analyze the performance of Algorithm 1 when the ETC algorithm (see Section 3) is used in the bandit experiment. The next lemma characterizes the sign of the bias of the ETC algorithm in

Table 1: Debiasing using Algorithm 1 for a two-armed bandit problem

Algorithm	Reward	Real World				Bootstrap World			
		$E[\hat{\mu}_k] - \mu_k$		$E[\hat{\mu}_k]$		Estimated Bias		Corrected $E[\hat{\mu}_k]$	
		Arm 1	Arm 2	Arm 1	Arm 2	Arm 1	Arm 2	Arm 1	Arm 2
ETC	Normal	-0.0518	-0.0324	0.9482	1.4676	-0.0377	-0.0396	0.9859	1.5072
	Bernoulli	-0.0184	-0.0167	0.2816	0.5833	-0.0148	-0.0156	0.2964	0.5989
UCB	Normal	-0.3407	-0.0348	0.6593	1.4652	-0.2026	-0.0361	0.8619	1.5012
	Bernoulli	-0.0384	-0.0035	0.2616	0.5965	-0.0406	-0.003	0.3022	0.5996
TS	Normal	-0.3717	-0.0504	0.6283	1.4496	-0.1804	-0.0268	0.8087	1.4764
	Bernoulli	-0.0458	-0.0098	0.2542	0.5902	-0.03	-0.0032	0.2842	0.5934
EG	Normal	-0.167	-0.109	0.833	1.391	-0.1171	-0.0958	0.9501	1.4868
	Bernoulli	-0.0508	-0.0828	0.2492	0.5172	-0.0533	-0.0377	0.3025	0.5549




Figure 1: Debiasing using bootstrap under ETC

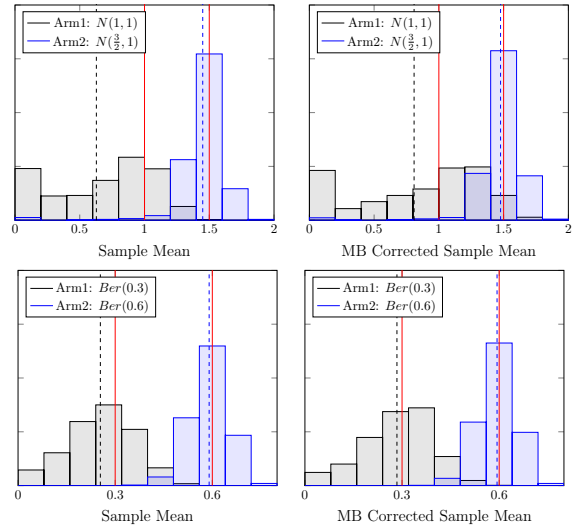


Figure 3: Debiasing using bootstrap under TS

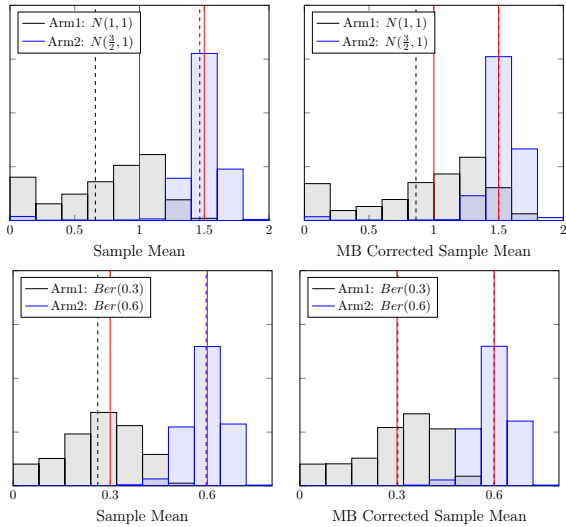


Figure 2: Debiasing using bootstrap under UCB

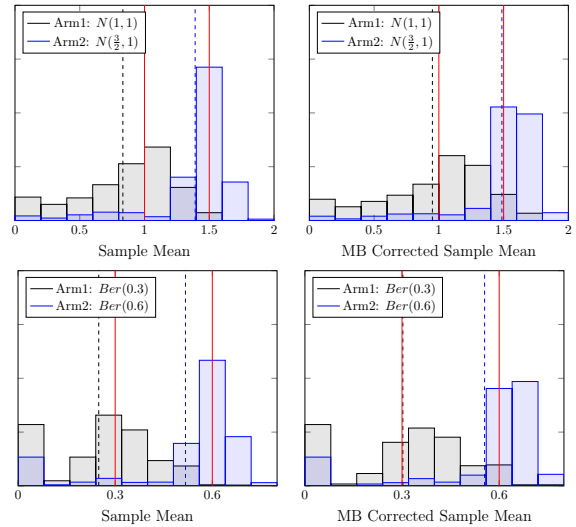


Figure 4: Debiasing using bootstrap under EG




Figure 5: MSE comparison under TS




Figure 6: MSE comparison under EG

the real world.

Lemma 1. *Bias of arm k , $\mathbb{E}_P[\hat{\mu}_k] - \mu_k$, is negative when the ETC algorithm is used for all $k \in [K]$.*

Although the result is a special case of Nie et al. (2018), in the proof (see Appendix B) we are able to explicitly characterize the bias of the sample average of arm k , in the form

$$\begin{aligned} \mathbb{E}_P[\hat{\mu}_k] - \mu_k & \\ &= \frac{T - mK}{m + T - mK} \cdot \mathbb{E}_P[(\mu_k - \hat{\mu}_k(mK))\mathbb{I}_{\{a_{mK+1}=k\}}], \end{aligned} \quad (2)$$

where $\hat{\mu}_k(mK)$ is the sample average reward of arm k in the first mK rounds. The bias is fully generated by the correlation between $\mu_k - \hat{\mu}_k(mK)$, which is the deviation of the sample average in the exploration phase from the mean, and $\mathbb{I}_{\{a_{mK+1}=k\}}$, which is the event that arm k is chosen in the exploitation phase. The ratio outside the expectation adjusts for the length of the exploration/commit phases. The expression (2)

allows us to derive the main theoretical results in the paper.

4.1 Two Arms with Gaussian Rewards

In this section, we consider the special case when $K = 2$ and the rewards of both arms are generated from the normal distribution. That is, the reward of arm i is generated from $N(\mu_i, \sigma_i^2)$ for $i = 1, 2$. We first compute the bias explicitly using the sample mean when the ETC algorithm is used.

Proposition 1. *Under the ETC algorithm, for $K = 2$ and Gaussian rewards, the bias of the sample average of arm k 's reward is*

$$\begin{aligned} \mathbb{E}_P[\hat{\mu}_k] - \mu_k & \\ &= -\frac{T - 2m}{T - m} \frac{\sigma_k^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}m} \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \mu_2)^2\right]. \end{aligned} \quad (3)$$

We have the following observations from Proposition 1:

- The bias is always negative for both arms, confirming Lemma 1.
- If $\sigma_1 = \sigma_2$, then the bias is equal for both arms. This may be counter-intuitive, as one would expect that the better arm is more likely to be chosen after the exploration phase and the sample average is thus less biased.
- Given σ_1 and σ_2 , the bias is maximized when $\mu_1 = \mu_2$, i.e., when the two arms are hard to distinguish.
- Given the reward distributions, the bias is decreasing exponentially in m .

To prove Proposition 1, we use the bias characterization (2). Under the Gaussian assumption of reward distributions, the sample average $\hat{\mu}_k(mK)$ in the exploration is still Gaussian; In addition, whether the arm k is chosen in the exploitation phase amounts to the comparison of average rewards of two arms which both follow Gaussian distributions. This allows us to obtain the explicit formula for the bias in (3) in the two-arm setting. For multiple arms, it is difficult to obtain such explicit formulas due to the comparison of multiple Gaussian distributions (Kim and Nelson, 2007).

Next, we investigate the effect of using Gaussian multiplier bootstrap to estimate (3) and correct the bias. In order to generate the reward in the bootstrap world, we need to obtain the sample mean and variance in (1). There are two outcomes in the context of ETC. If arm one has a higher average in the exploration phase,

then we have $\hat{\mu}_1 = \frac{1}{T-m} \left(\sum_{i=1}^m r_i + \sum_{i=2m+1}^T r_i \right)$, $\hat{\sigma}_1^2 = \frac{1}{T-m} \left(\sum_{i=1}^m (r_i - \hat{\mu}_1)^2 + \sum_{i=2m+1}^T (r_i - \hat{\mu}_1)^2 \right)$, $\hat{\mu}_2 = \frac{1}{m} \sum_{i=m+1}^{2m} r_i$, and $\hat{\sigma}_2^2 = \frac{1}{m} \sum_{i=m+1}^{2m} (r_i - \hat{\mu}_2)^2$. This outcome happens with probability $\mathbb{P}(N(\mu_1, \frac{\sigma_1^2}{m}) > N(\mu_2, \frac{\sigma_2^2}{m})) = \Phi \left(\frac{(\mu_1 - \mu_2)\sqrt{m}}{\sqrt{\sigma_1^2 + \sigma_2^2}} \right)$. Otherwise, arm two has a higher average in the exploration phase. We have $\hat{\mu}_1 = \frac{1}{m} \sum_{i=1}^m r_i$, $\hat{\sigma}_1^2 = \frac{1}{m} \sum_{i=1}^m (r_i - \hat{\mu}_1)^2$, $\hat{\mu}_2 = \frac{1}{T-m} \sum_{i=m+1}^T r_i$, and $\hat{\sigma}_2^2 = \frac{1}{T-m} \sum_{i=m+1}^T (r_i - \hat{\mu}_2)^2$. This outcome occurs with probability $\Phi \left(\frac{(\mu_2 - \mu_1)\sqrt{m}}{\sqrt{\sigma_1^2 + \sigma_2^2}} \right)$.

After obtaining $\hat{\mu}_k$ and $\hat{\sigma}_k$ for $k = 1, 2$, we can apply Proposition 1 to the bandit experiments in the bootstrap world, i.e., Step 7 in Algorithm 1. In particular, when $B \rightarrow \infty$, the law of large numbers implies that the bias in the bootstrap world is given by

$$\begin{aligned} \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \hat{\mu}_{b,k}^* - \hat{\mu}_k &= \mathbf{E}_{P^*}[\hat{\mu}_k^*] - \hat{\mu}_k \\ &= -\frac{T-2m}{T-m} \frac{\hat{\sigma}_k^2}{\sqrt{2\pi(\hat{\sigma}_1^2 + \hat{\sigma}_2^2)m}} \exp \left[-\frac{m}{2(\hat{\sigma}_1^2 + \hat{\sigma}_2^2)} (\hat{\mu}_1 - \hat{\mu}_2)^2 \right]. \end{aligned} \quad (4)$$

Note that the bias is conditional on the data obtained from bandit experiment in the real world, \mathcal{X} , i.e., the action sequence $\{a_1, \dots, a_T\}$, the reward sequence $\{r_1, \dots, r_T\}$.

To compare the estimated bias in the bootstrap world, (4), to the actual bias in the real world, (3), note that in the bootstrap world, we simply replace μ_k and σ_k by the sample version $\hat{\mu}_k$ and $\hat{\sigma}_k$. Such analytical tractability is one of the major benefits of Gaussian multiplier bootstrap. We define following quantity:

$$\begin{aligned} g_k(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) \\ = \log \left(\frac{T-2m}{T-m} \frac{\sigma_k^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)m}} \exp \left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)} (\mu_1 - \mu_2)^2 \right] \right), \end{aligned}$$

which is the logarithm of the absolute value of bias for arm k in the real world. The corresponding quantity in the bootstrap world is thus $g_k(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$. We focus on the logarithm because, as shown in Proposition 1, the bias decays exponentially fast as m increases. Therefore, we focus on the ratio of the decay rate. Our next result states that the ratio converges to one asymptotically.

Theorem 1. *Consider the ETC algorithm for $K = 2$ and Gaussian rewards. Choosing $T > 2m$, for $k = 1, 2$, we have*

$$\lim_{m \rightarrow \infty} \frac{g_k(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)}{g_k(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)} = 1, \quad \text{in probability.}$$

Theorem 1 provides a theoretical justification for the performance of Algorithm 1. At least for the ETC

algorithm, when the exploration phase lengthens, the biases from the bootstrap world and the real world converge in the logarithmic sense.

Instead of the ratio of the log-biases, one may hope to obtain a stronger result: the convergence of the ratio of the biases themselves. We caution that the claim is not true in general. This is because the exponential may distort a small error and lead to divergence.

To prove Theorem 1, we perform a Taylor expansion of $g_k(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$ at the point $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$. We then show that $\nabla g_k/g_k = O(1)$ and $\|(\hat{\mu}_1 - \mu_1, \hat{\mu}_2 - \mu_2, \hat{\sigma}_1^2 - \sigma_1^2, \hat{\sigma}_2^2 - \sigma_2^2)\|$ converges in probability to zero as $m \rightarrow \infty$. See Appendix D for details.

4.2 General Rewards

In this section, we extend the setting in Section 4.1 by considering general rewards. In particular, we consider the ETC algorithm with $K = 2$ and the subgaussian reward distributions. The reward distribution of the first arm has mean μ_1 and variance proxy parameter a^2 , i.e. $\mathbf{E}[e^{\lambda(X_1 - \mu_1)}] \leq \exp(\frac{\lambda^2 a^2}{2})$ for all $\lambda \in \mathbb{R}$. It is well known that $\sigma_1^2 \triangleq \text{Var}(X_1) \leq a^2$, where the equality holds when X is Gaussian. For simplicity, we assume the reward of the second arm is deterministic and equals μ_2 . The analysis can be extended to the general subgaussian setting.

To facilitate the presentation, we recall some definitions from the large deviations theory (Dembo and Zeitouni, 1998). Define the log moment generating function $\eta(h) = \log \mathbf{E}[e^{hX_1}]$ for the reward of arm one. Note η is convex and continuously differentiable on \mathbb{R} with $\eta'(0) = \mathbf{E}[X_1] = \mu_1$ (Lemma 2.2.5 of Dembo and Zeitouni (1998)). Set $\mathcal{H} = \{\eta'(h) : h \in \mathbb{R}\}$. The Legendre-Fenchel transform of η is given by

$$\Lambda^*(x) = \sup_{h \in \mathbb{R}} (hx - \eta(h)).$$

Define ζ by the implicit equation $\eta'(\zeta) = \mu_2$ when $\mu_2 \in \mathcal{H}$. The next result characterizes the bias in this setting in the real world.

Proposition 2. *Under the ETC algorithm, for $K = 2$, subgaussian reward for arm one and deterministic reward for arm two, the bias of the sample average of arm one's reward is*

$$\mathbf{E}_P[\hat{\mu}_1] - \mu_1 = \frac{T-2m}{T-m} \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m \eta''(\zeta)}} (-c_* + o(1)), \quad (5)$$

as $m \rightarrow \infty$, where $c_* > 0$ is a constant.

The constant c_* in the expression is semi-explicit. The value depends on whether the reward X_1 has a lattice or

non-lattice law, which is given in (20) in the appendix. Proposition 2 shows that the bias of arm one decays exponentially fast to zero with rate $\Lambda^*(\mu_2)$ as m grows. This is a similar pattern to Proposition 1. In fact, in the special case that $X_1 \sim N(\mu_1, \sigma_1^2)$, one can readily verify that $\eta(h) = \mu_1 h + \frac{1}{2}\sigma_1^2 h^2$, $\Lambda^*(\mu_2) = \frac{1}{2\sigma_1^2}(\mu_1 - \mu_2)^2$ and $c_* = \sigma_1^2$. Hence, (5) is consistent with (3) when the rewards are Gaussian.

To prove Proposition 2, we build upon the bias characterization in (2), and use tools from large deviations theory. We need the large deviations theory since one has to compare the average rewards of two arms in the exploration to decide which arm to pull in the exploitation phase for the ETC algorithm. Mathematically, to obtain (5), we build on the Bahadur-Rao theorem (Theorem 3.7.4 in Dembo and Zeitouni (1998)) and provide precise estimates for the tail probabilities and tail expectations for the average of i.i.d. random variables. The details of the proof can be found in Appendix E.

Next, we consider the bias estimate in the bootstrap world. With Gaussian bootstrap multiplier, the bootstrap distribution (1) is always Gaussian with distribution $N(\hat{\mu}_1, \hat{\sigma}_1^2)$, capturing the first- and second-order moments of the original distribution. As a result, we can infer from (4) that the bias of arm one in the bootstrap world is given by

$$\begin{aligned} \lim_{B \rightarrow \infty} \frac{1}{B} \sum_{b=1}^B \hat{\mu}_{b,1}^* - \hat{\mu}_1 &= \mathbf{E}_{P^*}[\hat{\mu}_1^*] - \hat{\mu}_1 \\ &= -\frac{T-2m}{T-m} \frac{\hat{\sigma}_1^2}{\sqrt{2\pi\hat{\sigma}_1^2 m}} \exp\left[-m\hat{\Lambda}^*(\mu_2)\right], \end{aligned}$$

where the exponential decay rate of the bias in the bootstrap world is given by

$$\hat{\Lambda}^*(\mu_2) = \frac{1}{2\hat{\sigma}_1^2}(\hat{\mu}_1 - \mu_2)^2. \quad (6)$$

Comparing $\hat{\Lambda}^*(\mu_2)$ in (6) to $\Lambda^*(\mu_2)$ in (5), we have:

Theorem 2. *Under the same assumptions of Proposition 2, the ratio of the decay rate of the biases in the bootstrap world and the real world is asymptotically bounded by*

$$\lim_{m \rightarrow \infty} \frac{\hat{\Lambda}^*(\mu_2)}{\Lambda^*(\mu_2)} \leq \frac{a^2}{\sigma_1^2},$$

where σ_1^2 and a^2 are the variance and the variance proxy of the subgaussian reward from arm one, respectively.

Theorem 2 characterizes the relationship of the decay rates of the biases in the real and bootstrap worlds. It reveals one potential drawback of bootstrap despite its good performance in Section 3. Only the first-

and second-order moments captured by the Gaussian multiplier bootstrap are not sufficient to characterize the asymptotic decay rate. The actual decay rate may depend on the tail behavior, which is captured by a and the Legendre-Fenchel transform Λ^* . This is in general hard to capture using bootstrap, including Efron’s bootstrap. Still, Theorem 2 provides a bound for the ratio, which is attained when the reward indeed has a Gaussian distribution. For the details of the proof of Theorem 2, see Appendix F.

5 Conclusions and Future Research

In this paper, we study the problem of inferring the mean reward of the arms from the data generated in a multi-armed bandit experiment. The procedure we propose is based on bootstrap and doesn’t require any knowledge of the reward distribution. It can be used for a wide range of bandit algorithms. We demonstrate strong performance in numerical examples and analyze the theoretical properties for the ETC algorithm.

Our current theoretical analysis is limited to the ETC algorithm. It remains an open direction to analyze the procedure for other algorithms such as UCB and Thompson sampling. The major difficulty is to characterize the bias in a tractable way under such algorithms, where the adaptive data collection procedure induces complex dependency among rewards. Another interesting direction is to consider the finite-sample analysis of the bias reduction via bootstrap, which is more relevant given that the bias is typically larger in small samples. This may require a different set of mathematical techniques, and we leave it for future work.

Acknowledgements

The authors thank five anonymous referees for many helpful comments and suggestions. The research of Ningyuan Chen is partially supported by NSERC Discovery Grants Program RGPIN-2020-04038. Xuefeng Gao acknowledges support from Hong Kong RGC Grants 14201520 and 14201421.

References

- S. Arlot, G. Blanchard, E. Roquain, et al. Some nonasymptotic results on resampling in high dimension, i: confidence regions. *The Annals of Statistics*, 38(1):51–82, 2010.
- P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2):235–256, 2002.
- M. Aziz, E. Kaufmann, and M.-K. Riviere. On multi-armed bandit designs for dose-finding clinical trials. *Journal of Machine Learning Research*, 22:1–38, 2021.
- A. Bibaut, I. Malenica, N. Vlassis, and M. Van Der Laan. More efficient off-policy evaluation through regularized targeted learning. In *International Conference on Machine Learning*, pages 654–663. PMLR, 2019.
- C. Boutilier, C.-W. Hsu, B. Kveton, M. Mladenov, C. Szepesvari, and M. Zaheer. Differentiable meta-learning of bandit policies. *Advances in Neural Information Processing Systems*, 33:2122–2134, 2020.
- G. Burtini, J. Loepky, and R. Lawrence. A survey of online experiment design with the stochastic multi-armed bandit. *Working Paper*, 2015.
- N. Chen and A. Khademi. Adaptive seamless dose-finding trials. *Working Paper*, 2020.
- V. Chernozhukov, D. Chetverikov, K. Kato, et al. Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *The Annals of Statistics*, 41(6):2786–2819, 2013.
- A. DasGupta. *Probability for statistics and machine learning: fundamentals and advanced topics*. Springer Science & Business Media, 2011.
- A. Dembo and O. Zeitouni. *Large deviations techniques and applications*. Springer-Verlag, New York, 1998.
- A. V. Den Boer. Dynamic pricing and learning: historical origins, current research, and new directions. *Surveys in operations research and management science*, 20(1):1–18, 2015.
- Y. Deshpande, L. Mackey, V. Syrgkanis, and M. Taddy. Accurate inference for adaptive linear models. In *International Conference on Machine Learning*, pages 1194–1203. PMLR, 2018.
- M. Dimakopoulou, Z. Zhou, S. Athey, and G. Imbens. Estimation considerations in contextual bandits. *Working Paper*, 2017.
- D. Eckles and M. Kaptein. Thompson sampling with the online bootstrap. *arXiv preprint arXiv:1410.4009*, 2014.
- B. Efron. Bootstrap methods: another look at the jackknife. *The Annals of Statistics*, 7:1–26, 1979.
- B. Efron. *The jackknife, the bootstrap and other resampling plans*. SIAM, 1982.
- B. Efron and R. J. Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.
- A. N. Elmachtoub, R. McNellis, S. Oh, and M. Petrik. A practical method for solving contextual bandit problems using decision trees. *arXiv preprint arXiv:1706.04687*, 2017.
- M. Farajtabar, Y. Chow, and M. Ghavamzadeh. More robust doubly robust off-policy evaluation. In *International Conference on Machine Learning*, pages 1447–1456. PMLR, 2018.
- W. Feller. *An introduction to probability theory and its applications, vol 2*. John Wiley & Sons, 2008.
- V. Hadad, D. A. Hirshberg, R. Zhan, S. Wager, and S. Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the National Academy of Sciences*, 118(15), 2021.
- B. Hao, Y. Abbasi-Yadkori, Z. Wen, and G. Cheng. Bootstrapping upper confidence bound. *arXiv preprint arXiv:1906.05247*, 2019.
- B. Hao, Y. Duan, H. Lu, C. Szepesvári, M. Wang, et al. Bootstrapping statistical inference for off-policy evaluation. *arXiv preprint arXiv:2102.03607*, 2021.
- N. Jiang and L. Li. Doubly robust off-policy value evaluation for reinforcement learning. In *International Conference on Machine Learning*, pages 652–661. PMLR, 2016.
- S.-H. Kim and B. L. Nelson. Recent advances in ranking and selection. In *2007 Winter Simulation Conference*, pages 162–172. IEEE, 2007.
- I. Kostrikov and O. Nachum. Statistical bootstrapping for uncertainty estimation in off-policy evaluation. *arXiv preprint arXiv:2007.13609*, 2020.
- B. Kveton, C. Szepesvari, S. Vaswani, Z. Wen, T. Lattimore, and M. Ghavamzadeh. Garbage in, reward out: Bootstrapping exploration in multi-armed bandits. In *International Conference on Machine Learning*, pages 3601–3610. PMLR, 2019.
- T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in applied mathematics*, 6(1):4–22, 1985.
- L. Li, W. Chu, J. Langford, and X. Wang. Unbiased offline evaluation of contextual-bandit-based news article recommendation algorithms. In *Proceedings of the fourth ACM international conference on Web search and data mining*, pages 297–306, 2011.
- L. Li, W. Chu, J. Langford, T. Moon, and X. Wang. An unbiased offline evaluation of contextual bandit algorithms with generalized linear models. In *Proceedings of the Workshop on On-line Trading of Exploration*

- and Exploitation 2*, pages 19–36. JMLR Workshop and Conference Proceedings, 2012.
- L. Li, R. Munos, and C. Szepesvári. Toward minimax off-policy value estimation. In *Artificial Intelligence and Statistics*, pages 608–616. PMLR, 2015.
- J. Mary, P. Preux, and O. Nicol. Improving offline evaluation of contextual bandit algorithms via bootstrapping techniques. In *International Conference on Machine Learning*, pages 172–180. PMLR, 2014.
- S. Neel and A. Roth. Mitigating bias in adaptive data gathering via differential privacy. In *International Conference on Machine Learning*, pages 3720–3729. PMLR, 2018.
- X. Nie, X. Tian, J. Taylor, and J. Zou. Why adaptively collected data have negative bias and how to correct for it. In *International Conference on Artificial Intelligence and Statistics*, pages 1261–1269. PMLR, 2018.
- I. Osband and B. Van Roy. Bootstrapped thompson sampling and deep exploration. *arXiv preprint arXiv:1507.00300*, 2015.
- V. Perchet, P. Rigollet, S. Chassang, and E. Snowberg. Batched bandit problems. *The Annals of Statistics*, 44(2):660–681, 2016.
- D. Precup, R. S. Sutton, and S. P. Singh. Eligibility traces for off-policy policy evaluation. pages 759–766. PMLR, 2000.
- D. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen. A tutorial on thompson sampling. *arXiv preprint arXiv:1707.02038*, 2017.
- T. Schnabel, A. Swaminathan, A. Singh, N. Chandak, and T. Joachims. Recommendations as treatments: Debiasing learning and evaluation. In *international conference on machine learning*, pages 1670–1679. PMLR, 2016.
- J. Shin, A. Ramdas, and A. Rinaldo. On the bias, risk and consistency of sample means in multi-armed bandits. *arXiv preprint arXiv:1902.00746*, 2019a.
- J. Shin, A. Ramdas, and A. Rinaldo. Are sample means in multi-armed bandits positively or negatively biased? In *Advances in Neural Information Processing Systems*, pages 7102–7111, 2019b.
- H. Steck and T. S. Jaakkola. Bias-corrected bootstrap and model uncertainty. In *NIPS*, pages 521–528. Citeseer, 2003.
- R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- A. Swaminathan and T. Joachims. Counterfactual risk minimization: Learning from logged bandit feedback. In *International Conference on Machine Learning*, pages 814–823. PMLR, 2015.
- L. Tang, Y. Jiang, L. Li, C. Zeng, and T. Li. Personalized recommendation via parameter-free contextual bandits. In *Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval*, pages 323–332, 2015.
- W. R. Thompson. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- A. W. Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.
- A. W. van Der Vaart and J. A. Wellner. *Weak Convergence and Empirical Processes: With Application to Statistics*. Springer, 1996.
- S. Vaswani, B. Kveton, Z. Wen, A. Rao, M. Schmidt, and Y. Abbasi-Yadkori. New insights into bootstrapping for bandits. *arXiv preprint arXiv:1805.09793*, 2018.
- S. S. Villar, J. Bowden, and J. Wason. Multi-armed bandit models for the optimal design of clinical trials: benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 30(2):199, 2015.
- N. Vlassis, A. Bibaut, M. Dimakopoulou, and T. Jebara. On the design of estimators for bandit off-policy evaluation. In *International Conference on Machine Learning*, pages 6468–6476. PMLR, 2019.
- M. Xu, T. Qin, and T.-Y. Liu. Estimation bias in multi-armed bandit algorithms for search advertising. *Advances in Neural Information Processing Systems*, 26:2400–2408, 2013.

Supplementary Material: Debiasing Samples from Online Learning Using Bootstrap

A Additional experiments

We conduct additional bandit experiments with $K = 4$ arms: Gaussian reward with mean $\mu = \{2, 2.5, 3, 3.5\}$ and standard deviation $\sigma = \{2, 1, 2, 1\}$, and Bernoulli reward with mean $\mu = \{0.4, 0.5, 0.7, 0.8\}$. The four bandit algorithms (ETC, UCB, TS, EG) with the same parameters setting as Section 3 are repeated 1,000 times, each instance with $B = 1,000$ bootstrap simulations for both Efron’s bootstrap and Gaussian multiplier bootstrap. The results are presented below. They further demonstrate the effectiveness of the proposed algorithm.

To further illustrate the effect of the bootstrap size B on bias corrections, we report the following results for the ETC algorithm with the four-armed Gaussian bandit problem. From Table 2, we observe that the corrected means are $[2.0191, 2.4927, 3.0013, 3.4921]$ when $B = 1000$. We vary B in our experiments and find that the corrected means are $[2.0221, 2.4942, 3.0012, 3.4906]$ when $B = 50$, $[2.018, 2.4926, 3.0007, 3.493]$ when $B = 100$, and $[2.0185, 2.4924, 3.0007, 3.4923]$ when $B = 500$, respectively. Similar as in the two-armed setting, we find that when B is of moderate size, our algorithm already corrects the bias very well.

Table 2: Debiasing using Bootstrap for a Four-Armed Bandit Problem.

Algorithm	Reward	Arm	True Environment		MB Environment		EB Environment	
			$E[\hat{\mu}_k] - \mu_k$	$E[\hat{\mu}_k]$	Estimated Bias	Corrected $E[\hat{\mu}_k]$	Estimated Bias	Corrected $E[\hat{\mu}_k]$
ETC	Normal	Arm1	-0.0197	1.9803	-0.0388	2.0191	-0.0369	2.0172
		Arm2	-0.0219	2.4781	-0.0146	2.4927	-0.0142	2.4924
		Arm3	-0.0347	2.9653	-0.036	3.0013	-0.0355	3.0008
		Arm4	-0.1537	3.3463	-0.1459	3.4921	-0.1462	3.4924
	Bernoulli	Arm1	-0.006	0.394	-0.0064	0.4005	-0.0092	0.4032
		Arm2	-0.0102	0.4898	-0.0127	0.5024	-0.0158	0.5055
		Arm3	-0.0366	0.6634	-0.0263	0.6897	-0.0284	0.6918
		Arm4	-0.029	0.771	-0.0259	0.797	-0.0273	0.7983
UCB	Normal	Arm1	-0.8538	1.1462	-0.2737	1.4198	-0.3998	1.546
		Arm2	-0.7435	1.7565	-0.2109	1.9674	-0.2136	1.9701
		Arm3	-0.5521	2.4479	-0.1479	2.5959	-0.1482	2.5961
		Arm4	-0.699	2.801	-0.4529	3.2539	-0.5017	3.3026
	Bernoulli	Arm1	-0.0661	0.3339	-0.0564	0.3904	-0.0503	0.3843
		Arm2	-0.0666	0.4334	-0.0580	0.4914	-0.0539	0.4873
		Arm3	-0.0393	0.6607	-0.0363	0.6969	-0.0380	0.6986
		Arm4	-0.0119	0.7881	-0.0121	0.8002	-0.0137	0.8018
TS	Normal	Arm1	-0.7917	1.2083	-0.2996	1.5079	-0.3321	1.5404
		Arm2	-0.2872	2.2128	-0.1701	2.3829	-0.1721	2.3848
		Arm3	-0.2594	2.7406	-0.1401	2.8807	-0.1398	2.8804
		Arm4	-0.6559	2.8441	-0.4913	3.3354	-0.4974	3.3415
	Bernoulli	Arm1	-0.0856	0.3144	-0.0467	0.3611	-0.0423	0.3567
		Arm2	-0.0879	0.4121	-0.0408	0.4529	-0.0391	0.4513
		Arm3	-0.0526	0.6474	-0.0283	0.6757	-0.0299	0.6773
		Arm4	-0.0346	0.7654	-0.012	0.7774	-0.0135	0.7789
EG	Normal	Arm1	-0.2252	1.7748	-0.2128	1.9876	-0.2204	1.9952
		Arm2	-0.1482	2.3518	-0.1039	2.4556	-0.1078	2.4595
		Arm3	-0.3591	2.6409	-0.2167	2.8576	-0.2235	2.8644
		Arm4	-0.959	2.541	-0.5405	3.0815	-0.5554	3.0964
	Bernoulli	Arm1	-0.1115	0.2885	-0.0652	0.3537	-0.0704	0.3590
		Arm2	-0.1310	0.3690	-0.0820	0.4510	-0.0957	0.4646
		Arm3	-0.1517	0.5483	-0.0985	0.6467	-0.1168	0.6650
		Arm4	-0.1182	0.6818	-0.0831	0.7649	-0.1083	0.7901

B Proof of Lemma 1

Proof. For arm k , the average rewards of the whole decision horizon T is

$$\hat{\mu}_k = \begin{cases} \frac{\sum_{t=(k-1)m+1}^{km} r_t + \sum_{t=Km+1}^T r_t}{m+T-mK}, & \text{if } k = a_{mK+1}, \\ \frac{\sum_{t=(k-1)m+1}^{km} r_t}{m} = \hat{\mu}_k(mK), & \text{if } k \neq a_{mK+1}. \end{cases}$$

Hence we have

$$\mathbb{E}_P[\hat{\mu}_k] = \mathbb{E}_P \left[\hat{\mu}_k(mK) \mathbb{I}_{\{k \neq a_{mK+1}\}} + \hat{\mu}_k \mathbb{I}_{\{k = a_{mK+1}\}} \right].$$

Note for each $k = 1, \dots, K$, $\mathbb{E}_P[\hat{\mu}_k(mK)] = \mu_k$. Then the bias of arm k is given by

$$\begin{aligned} \mathbb{E}_P[\hat{\mu}_k] - \mu_k &= \mathbb{E}_P \left[(-\hat{\mu}_k(mK) + \hat{\mu}_k) \mathbb{I}_{\{k = a_{mK+1}\}} \right] \\ &= \frac{T - mK}{m + T - mK} \cdot \mathbb{E}_P \left[\left(\frac{\sum_{t=Km+1}^T r_t}{T - mK} - \hat{\mu}_k(mK) \right) \mathbb{I}_{\{k = a_{mK+1}\}} \right] \\ &= \frac{T - mK}{m + T - mK} \cdot \mathbb{E}_P \left[(\mu_k - \hat{\mu}_k(mK)) \mathbb{I}_{\{k = a_{mK+1}\}} \right], \end{aligned} \quad (7)$$

where the expectation is completely determined by the joint distribution of $(\hat{\mu}_k(mK))_i$ with independent marginals. In particular, the marginal distribution of $\hat{\mu}_k(mK)$ is simply the average of m i.i.d. rewards associated with arm k . Mathematically, we have

$$\begin{aligned} \mathbb{E}_P[\hat{\mu}_k] - \mu_k &= \frac{T - mK}{m + T - mK} \cdot \int_{\mathbb{R}} (\mu_k - x) \mathbb{P}(k = a_{mK+1}) \mathbb{P}(\hat{\mu}_k(mK) \in dx) \\ &\leq \frac{T - mK}{m + T - mK} \cdot \int_{\mathbb{R}} (\mu_k - x) \mathbb{P}(\max_{k' \neq k} \hat{\mu}_{k'}(mK) \leq x) \mathbb{P}(\hat{\mu}_k(mK) \in dx). \end{aligned}$$

Therefore, the bias is negative since $\int_{\mathbb{R}} (\mu_k - x) \mathbb{P}(\hat{\mu}_k(mK) \in dx) = \mu_k - \mathbb{E}_P[\hat{\mu}_k(mK)] = 0$, and $\mathbb{P}(\max_{k' \neq k} \hat{\mu}_{k'}(mK) \leq x)$ puts more weights on larger values of x . \blacksquare

C Proof of Proposition 1

Proof. Recall for a general K -armed bandit problems with continuous rewards, the bias of the sample mean of arm k has the following expression:

$$\mathbb{E}_P[\hat{\mu}_k] - \mu_k = \frac{T - mK}{m + T - mK} \cdot \mathbb{E}_P \left[(\mu_k - \hat{\mu}_k(mK)) \mathbb{I}_{\{\hat{\mu}_k(mK) > \max_{k' \neq k} \hat{\mu}_{k'}(mK)\}} \right].$$

To prove Proposition 1 with $K = 2$, we first consider the special case where the arm 1 follows $N(\mu_1, \sigma_1^2)$ while the arm 2 has deterministic rewards μ_2 , i.e. $\sigma_2 = 0$. Then we have the bias of arm 2 is clearly zero, and the bias of arm 1 is given by

$$\mathbb{E}_P[\hat{\mu}_1] - \mu_1 = \frac{T - 2m}{T - m} \cdot \mathbb{E}_P \left[(\mu_1 - \hat{\mu}_1(2m)) \mathbb{I}_{\{\hat{\mu}_1(2m) \geq \mu_2\}} \right], \quad (8)$$

where $\hat{\mu}_1(2m) = \frac{1}{m} \sum_{i=1}^m r_i$ with r_i generated from $N(\mu_1, \sigma_1^2)$. Note that $\hat{\mu}_1(2m)$ is a random variable following the normal distribution $N(\mu_1, \frac{\sigma_1^2}{m})$. Hence we obtain

$$\mathbb{E}_P \left[(\mu_1 - \hat{\mu}_1(2m)) \mathbb{I}_{\{\hat{\mu}_1(2m) \geq \mu_2\}} \right] = \frac{\sigma_1}{\sqrt{m}} \cdot \mathbb{E} \left[-Z \cdot \mathbb{I}_{\{Z > \frac{\mu_2 - \mu_1}{\sigma_1 / \sqrt{m}}\}} \right],$$

where $Z \sim N(0, 1)$. Notice that $\mathbb{E} [Z \cdot \mathbb{I}_{\{Z > c\}}] = \phi(c)$ for any c , where ϕ is the density function of $N(0, 1)$. It follows that

$$\mathbb{E}_P \left[(\mu_1 - \hat{\mu}_1(2m)) \mathbb{I}_{\{\hat{\mu}_1(2m) \geq \mu_2\}} \right] = -\frac{\sigma_1}{\sqrt{m}} \phi \left(\frac{\mu_2 - \mu_1}{\sigma_1 / \sqrt{m}} \right) = -\frac{\sigma_1}{\sqrt{2\pi m}} \exp \left[-\frac{1}{2} \left(\frac{\mu_2 - \mu_1}{\sigma_1} \right)^2 m \right].$$

On combining with (8), we immediately obtain the expression for the bias of arm 1 in (3).

Next we consider the setting where the reward of arm 2 follows $N(\mu_2, \sigma_2^2)$ with $\sigma_2 > 0$. The bias of arm 1 is clearly given by

$$\mathbb{E}_P[\hat{\mu}_1] - \mu_1 = \frac{T-2m}{T-m} \cdot \mathbb{E}_P[(\mu_1 - \hat{\mu}_1(2m))\mathbb{I}_{\{\hat{\mu}_1(2m) > \hat{\mu}_2(2m)\}}], \quad (9)$$

where $\hat{\mu}_2(2m) \sim N(\mu_2, \frac{\sigma_2^2}{m})$ since $\hat{\mu}_1(2m) = \frac{1}{m} \sum_{i=m+1}^{2m} r_i$ with r_i obtained from $N(\mu_2, \sigma_2^2)$. The difference between (8) and (9) is just that we substitute a deterministic value μ_2 by a normal random variable $\hat{\mu}_2(2m)$. So it follows that

$$\begin{aligned} & \mathbb{E}_P[(\mu_1 - \hat{\mu}_1(2m))\mathbb{I}_{\{\hat{\mu}_1(2m) > \hat{\mu}_2(2m)\}}] \\ &= \int_y \left[-\frac{\sigma_1}{\sqrt{2\pi m}} \exp\left[-\frac{1}{2} \left(\frac{y - \mu_1}{\sigma_1}\right)^2 m\right] \cdot \phi\left(\frac{y - \mu_2}{\sigma_2/\sqrt{m}}\right) \cdot \frac{\sqrt{m}}{\sigma_2} \right] dy \\ &= -\frac{\sigma_1}{2\pi\sigma_2} \int_y \exp\left[-\frac{m}{2\sigma_1^2}(y - \mu_1)^2 - \frac{m}{2\sigma_2^2}(y - \mu_2)^2\right] dy. \end{aligned}$$

Note that

$$-\frac{m}{2\sigma_1^2}(y - \mu_1)^2 - \frac{m}{2\sigma_2^2}(y - \mu_2)^2 = -\frac{m}{2\sigma_1^2\sigma_2^2} \left[\left(\sqrt{\sigma_1^2 + \sigma_2^2}y - \frac{\sigma_2^2\mu_1 + \sigma_1^2\mu_2}{\sqrt{\sigma_2^2 + \sigma_1^2}} \right)^2 + \frac{\sigma_2^2\sigma_1^2}{\sigma_2^2 + \sigma_1^2}(\mu_1 - \mu_2)^2 \right].$$

Hence, we can obtain

$$\begin{aligned} & \mathbb{E}_P[(\mu_1 - \hat{\mu}_1(2m))\mathbb{I}_{\{\hat{\mu}_1(2m) > \hat{\mu}_2(2m)\}}] \\ &= -\frac{\sigma_1}{2\pi\sigma_2} \int_y \exp\left[-\frac{m}{2\sigma_1^2\sigma_2^2} \left(\sqrt{\sigma_1^2 + \sigma_2^2}y - \frac{\sigma_2^2\mu_1 + \sigma_1^2\mu_2}{\sqrt{\sigma_2^2 + \sigma_1^2}} \right)^2\right] dy \cdot \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \theta)^2\right] \\ &= -\frac{\sigma_1^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}m} \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \mu_2)^2\right]. \end{aligned}$$

Therefore, the bias of arm 1 is

$$\mathbb{E}_P[\hat{\mu}_1] - \mu_1 = -\frac{T-2m}{T-m} \frac{\sigma_1^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}m} \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \mu_2)^2\right].$$

This proves (3) for $k = 1$.

Similarly, one can compute the bias of arm 2 by

$$\mathbb{E}_P[\hat{\mu}_2] - \mu_2 = \frac{T-2m}{T-m} \cdot \mathbb{E}[(\mu_2 - \hat{\mu}_2(2m))\mathbb{I}_{\{\hat{\mu}_2(2m) > \hat{\mu}_1(2m)\}} | P],$$

where $\hat{\mu}_1(2m) = \frac{1}{m} \sum_{i=1}^m r_i \sim N(\mu_1, \frac{\sigma_1^2}{m})$ and $\hat{\mu}_2(2m) = \frac{1}{m} \sum_{i=m+1}^{2m} r_i \sim N(\mu_2, \frac{\sigma_2^2}{m})$. Using a similar argument, we can obtain that the bias of arm 2 is

$$\mathbb{E}_P[\hat{\mu}_2] - \mu_2 = -\frac{T-2m}{T-m} \frac{\sigma_2^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}m} \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \mu_2)^2\right].$$

Therefore, the proof is complete. ■

D Proof of Theorem 1

Proof. We prove the result for $k = 1$ as the proof for $k = 2$ is similar. Recall

$$g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) = \log \left(\frac{T-2m}{T-m} \frac{\sigma_1^2}{\sqrt{2\pi(\sigma_1^2 + \sigma_2^2)}m} \exp\left[-\frac{m}{2(\sigma_1^2 + \sigma_2^2)}(\mu_1 - \mu_2)^2\right] \right). \quad (10)$$

To prove the result, we will perform Taylor expansion of the (random) function $g_1(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$ at the point $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$. We can readily compute that the first order (partial) derivatives of g_1 in (10) are given by

$$\begin{aligned}\frac{\partial g_1}{\partial \mu_1}(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) &= -\frac{m(\mu_1 - \mu_2)}{\sigma_1^2 + \sigma_2^2}, \\ \frac{\partial g_1}{\partial \mu_2}(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) &= \frac{m(\mu_1 - \mu_2)}{\sigma_1^2 + \sigma_2^2}, \\ \frac{\partial g_1}{\partial \sigma_1^2}(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) &= \frac{1}{\sigma_1^2} - \frac{1}{2(\sigma_1^2 + \sigma_2^2)} + \frac{m(\mu_1 - \mu_2)^2}{2(\sigma_1^2 + \sigma_2^2)^2}, \\ \frac{\partial g_1}{\partial \sigma_2^2}(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) &= -\frac{1}{2(\sigma_1^2 + \sigma_2^2)} + \frac{m(\mu_1 - \mu_2)^2}{2(\sigma_1^2 + \sigma_2^2)^2}.\end{aligned}$$

Hence we can write $\nabla g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) = \begin{pmatrix} -d_1 m \\ d_1 m \\ d_2 m + e_1 \\ d_2 m - e_2 \end{pmatrix}$, where d_1, d_2, e_1, e_2 are nonzero constants that are independent of m .

To perform the Taylor expansion, we recall and introduce the stochastic “ σ ” symbol, see Van der Vaart (2000). For a sequence of random variables $\{X_n, n = 0, 1, 2, \dots\}$, the expression $X_n = o_P(1)$ denotes that (X_n) converges to zero in probability. Furthermore, for a given sequence of random variables $\{R_n, n = 0, 1, 2, \dots\}$, the notation $X_n = o_P(R_n)$ means $X_n = Y_n R_n$ with $Y_n = o_P(1)$. We will show later that $\|(\hat{\mu}_1 - \mu_1, \hat{\mu}_2 - \mu_2, \hat{\sigma}_1^2 - \sigma_1^2, \hat{\sigma}_2^2 - \sigma_2^2)\| = o_P(1)$ as $m \rightarrow \infty$. Then Lemma 2.12 of Van der Vaart (2000) guarantees that we can take the Taylor expansion of $g_1(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)$ at point $(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$:

$$g_1(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2) = g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2) + \nabla g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)^T \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} + o_P \left(\left\| \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right\| \right).$$

It follows that

$$\frac{g_1(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2)}{g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)} = 1 + \frac{\nabla g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)^T \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix}}{g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)} + \frac{o_P \left(\left\| \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right\| \right)}{g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)}. \quad (11)$$

To prove Theorem 1, we proceed to show the second and the third terms of (11) will both converge to zero in probability as $m \rightarrow \infty$. For notational convenience, in the following we use \hat{g}_1, g_1 and ∇g_1 to denote $g_1(\hat{\mu}_1, \hat{\mu}_2, \hat{\sigma}_1^2, \hat{\sigma}_2^2), g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$ and $\nabla g_1(\mu_1, \mu_2, \sigma_1^2, \sigma_2^2)$, respectively.

We first show the second term of (11) converges to 0 in probability. Note that

$$\left| \frac{1}{g_1} \cdot \nabla g_1^T \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right| \leq \left\| \begin{pmatrix} -d_1 m / g_1 \\ d_1 m / g_1 \\ (d_2 m + e_1) / g_1 \\ (d_2 m - e_2) / g_1 \end{pmatrix} \right\| \cdot \left\| \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right\|.$$

It is easy to see from (10) that $g_1 = \Theta(m)$ as $m \rightarrow \infty$. Then there exists a constant $0 < C < \infty$ such that

$$\lim_{m \rightarrow \infty} \left\| \begin{pmatrix} -d_1 m / g_1 \\ d_1 m / g_1 \\ (d_2 m + e_1) / g_1 \\ (d_2 m - e_2) / g_1 \end{pmatrix} \right\| \leq C. \text{ Hence it suffices to show}$$

$$\lim_{m \rightarrow \infty} \mathbb{P} \left(\left\| \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right\| \geq \epsilon \right) \rightarrow 0, \quad i.e., \quad \left\| \begin{pmatrix} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{pmatrix} \right\| \xrightarrow{P} 0. \quad (12)$$

To this end, we can compute that

$$\begin{aligned}
 & \mathbb{P} \left(\left\| \begin{array}{c} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{array} \right\| \geq \epsilon \right) \\
 & \leq \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| + |\hat{\mu}_2 - \mu_2| + |\hat{\sigma}_1^2 - \sigma_1^2| + |\hat{\sigma}_2^2 - \sigma_2^2| \geq \epsilon \right) \\
 & \leq \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| \geq \frac{\epsilon}{4} \right) + \mathbb{P} \left(|\hat{\mu}_2 - \mu_2| \geq \frac{\epsilon}{4} \right) + \mathbb{P} \left(|\hat{\sigma}_1^2 - \sigma_1^2| \geq \frac{\epsilon}{4} \right) + \mathbb{P} \left(|\hat{\sigma}_2^2 - \sigma_2^2| \geq \frac{\epsilon}{4} \right).
 \end{aligned}$$

We know that

$$\begin{aligned}
 & \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| \geq \frac{\epsilon}{4} \right) \\
 & = \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| \geq \frac{\epsilon}{4} \mid \text{arm 1 chosen} \right) \mathbb{P}(\text{arm 1 chosen}) + \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| \geq \frac{\epsilon}{4} \mid \text{arm 2 chosen} \right) \mathbb{P}(\text{arm 2 chosen}) \\
 & \leq \mathbb{P} \left(\left| \frac{x_1 + \dots + x_{T-m}}{T-m} - \mu_1 \right| \geq \frac{\epsilon}{4} \right) + \mathbb{P} \left(\left| \frac{x_1 + \dots + x_m}{m} - \mu_1 \right| \geq \frac{\epsilon}{4} \right) \\
 & \leq \frac{16\sigma_1^2}{(T-m)\epsilon^2} + \frac{16\sigma_1^2}{m\epsilon^2},
 \end{aligned}$$

where the last inequality comes from Chebyshev's inequality. Therefore, we can obtain $\lim_{m \rightarrow \infty} \mathbb{P} \left(|\hat{\mu}_1 - \mu_1| \geq \frac{\epsilon}{4} \right) \rightarrow 0$, since $T - m \geq m$. Similarly, we can obtain that $\lim_{m \rightarrow \infty} \mathbb{P} \left(|\hat{\mu}_2 - \mu_2| \geq \frac{\epsilon}{4} \right) \rightarrow 0$, $\lim_{m \rightarrow \infty} \mathbb{P} \left(|\hat{\sigma}_1^2 - \sigma_1^2| \geq \frac{\epsilon}{4} \right) \rightarrow 0$ and $\lim_{m \rightarrow \infty} \mathbb{P} \left(|\hat{\sigma}_2^2 - \sigma_2^2| \geq \frac{\epsilon}{4} \right) \rightarrow 0$. This establishes (12), and it follows that the second term of (11) converges to 0 in probability.

Finally, we show that the third term of (11) also converges to 0 in probability. From (12), we have

$$o_P \left(\left\| \begin{array}{c} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{array} \right\| \right) \xrightarrow{P} 0. \text{ Moreover, } g_1 = \Theta(m) \text{ which implies } \lim_{m \rightarrow \infty} \frac{1}{g_1} = 0, \text{ then it follows that}$$

$$\frac{o_P \left(\left\| \begin{array}{c} \hat{\mu}_1 - \mu_1 \\ \hat{\mu}_2 - \mu_2 \\ \hat{\sigma}_1^2 - \sigma_1^2 \\ \hat{\sigma}_2^2 - \sigma_2^2 \end{array} \right\| \right)}{g_1} \xrightarrow{P} 0.$$

Therefore, we can deduce from (11) that the result in Theorem 1 holds. ■

E Proof of Proposition 2

Proof. Recall from (7) that the bias of arm 1 is given by

$$\mathbb{E}_P[\hat{\mu}_1(T)] - \mu_1 = \frac{T-2m}{T-m} \cdot \mathbb{E}_P \left[(\mu_1 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \geq \mu_2\}} \right]. \quad (13)$$

We focus on estimating the quantity

$$G(\bar{X}_m) := \mathbb{E}_P \left[(\mu_1 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \geq \mu_2\}} \right].$$

We first consider the case $\mu_1 < \mu_2$. Then

$$G(\bar{X}_m) = \mathbb{E} \left[(\mu_2 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \geq \mu_2\}} \right] + (\mu_1 - \mu_2) \mathbb{P}(\bar{X}_m \geq \mu_2). \quad (14)$$

Since $\mu_1 < \mu_2$, for the second term of (14), by the Bahadur-Rao theorem (Theorem 3.7.4 Dembo and Zeitouni (1998)), we have

$$\mathbb{P}(\bar{X}_m \geq \mu_2) = \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m\eta''(\zeta)}}(c_0 + o(1)), \quad (15)$$

where $\eta(h) = \log \mathbf{E}[e^{hX_1}]$, $\zeta > 0$ is defined by the implicit equation $\eta'(\zeta) = \mu_2$ for a given $\mu_2 > \mu_1$, and Λ^* is the Legendre-Fenchel transform of η :

$$\Lambda^*(x) = \sup_{h \in \mathbb{R}} (hx - \eta(h)).$$

The constant $c_0 > 0$ depends on the reward distribution X_i is lattice or non-lattice, and it is given by

$$c_0 = \begin{cases} \frac{1}{|\zeta|}, & \text{if } X_i \text{ has a non-lattice distribution,} \\ \frac{d}{1 - e^{-|\zeta|d}}, & \text{if } X_i \text{ has a lattice distribution, so that } \mathbb{P}(X_i = \mu_2) \in (0, 1), \text{ and } d \text{ is the largest number to} \\ & \text{make } (X_i - \mu_2)/d \text{ is (a.s.) an integer number.} \end{cases} \quad (16)$$

In addition, for the first term of (14), we also have the following result, the proof of which is deferred to the end of this section.

Lemma 2. *Let X_i be i.i.d. random variables with mean μ_1 and logarithmic moment generating function $\eta(h) = \log \mathbf{E}[e^{hX_i}]$ and ζ is defined by the implicit equation $\eta'(\zeta) = \mu_2$ for a given $\mu_2 > \mu_1$. Let $J_m(\mu_2) = -\zeta^2 \sqrt{2\pi m\eta''(\zeta)} m \exp(m\Lambda^*(\mu_2))$.*

(a) *If the law of X_i is non-lattice, then*

$$\lim_{m \rightarrow \infty} J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m \geq \mu_2\}}] = 1.$$

(b) *If X_i has a lattice law, so that $\mathbb{P}(X_i = \mu_2) \in (0, 1)$, $(X_i - \mu_2)/d$ is (a.s.) an integer number, and d is the largest number with this property, then*

$$\lim_{m \rightarrow \infty} J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m \geq \mu_2\}}] = \frac{\zeta d \cdot e^{-\zeta d}}{1 - e^{-\zeta d}}.$$

Therefore, we have

$$\mathbf{E}_P \left[(\mu_2 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \geq \mu_2\}} \right] = \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m^3 \eta''(\zeta)}} (c_1 + o(1)),$$

where c_1 is an explicit constant which also depends on whether the reward X_i is lattice or non-lattice:

$$c_1 = \begin{cases} -\frac{1}{\zeta^2}, & \text{if } X_i \text{ has a non-lattice distribution,} \\ -\frac{d \cdot e^{-\zeta d}}{(1 - e^{-\zeta d})\zeta}, & \text{if } X_i \text{ has a lattice distribution, so that } \mathbb{P}(X_i = \mu_2) \in (0, 1), \text{ and } d \text{ is the largest} \\ & \text{number to make } (X_i - \mu_2)/d \text{ is (a.s.) an integer number.} \end{cases}$$

Hence we obtain when $\mu_1 < \mu_2$,

$$G(\bar{X}_m) = \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m\eta''(\zeta)}} \left(c_0(\mu_1 - \mu_2) + \frac{c_1}{m} + o(1) \right) = \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m\eta''(\zeta)}} (c_0(\mu_1 - \mu_2) + o(1)), \quad (17)$$

where $c_0 > 0$ is given in (16). Hence we have completed the proof for the case $\mu_1 < \mu_2$.

We next consider the case when $\mu_1 > \mu_2$. Then $\{\bar{X}_m \geq \mu_2\}$ is not a rare event anymore when m is large, but $\{\bar{X}_m \leq \mu_2\}$ is a rare event with small probabilities. We can compute

$$\begin{aligned} G(\bar{X}_m) &= \mathbf{E}_P \left[(\mu_1 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \geq \mu_2\}} \right] \\ &= 0 - \mathbf{E}_P \left[(\mu_1 - \bar{X}_m) \mathbb{I}_{\{\bar{X}_m \leq \mu_2\}} \right] \\ &= (\mu_2 - \mu_1) \mathbb{P}(\bar{X}_m \leq \mu_2) + \mathbf{E}_P \left[(\bar{X}_m - \mu_2) \mathbb{I}_{\{\bar{X}_m \leq \mu_2\}} \right], \end{aligned} \quad (18)$$

where the second equality is due to the fact that $\mathbf{E}_P(\mu_1 - \bar{X}_m) = 0$. Applying Bahadur-Rao theorem to compute the first term in (18) and use a similar argument as in Lemma 2 to compute the second term in (18), we can similarly obtain that when $\mu_1 > \mu_2$,

$$G(\bar{X}_m) = \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m\eta''(\zeta)}} (c_0 \cdot (\mu_2 - \mu_1) + o(1)), \quad (19)$$

where ζ is the solution to the equation $\eta'(\zeta) = \mu_2$ and now $\zeta < 0$ when $\mu_1 > \mu_2$. The constant $c_0 > 0$ is given in (16).

Finally, we can combine (13), (17) and (19) to infer that the bias of arm 1 is given by

$$\mathbf{E}_P[\hat{\mu}_1] - \mu_1 = \frac{T - 2m}{T - m} \frac{e^{-m\Lambda^*(\mu_2)}}{\sqrt{2\pi m\eta''(\zeta)}} (-c_* + o(1)),$$

as $m \rightarrow \infty$, where the constant c_* is given by

$$c_* = c_0 \cdot |\mu_1 - \mu_2| > 0, \quad (20)$$

with $c_0 > 0$ given in (16). The proof is therefore complete. \blacksquare

E.1 Proof of Lemma 2

Proof. We adapt the proof of Bahadur-Rao theorem (Theorem 3.7.4 Dembo and Zeitouni (1998)) to our setting. When $\mu_1 < \mu_2$, recall $\eta(h) = \log \mathbf{E}[e^{hX_1}]$, and $\Lambda^*(x) = \sup_{h \in \mathbb{R}} (hx - \eta(h))$. Then define $D_\eta := \{h \in \mathbb{R} : \eta(h) < \infty\}$. Then in $\text{Int}\{D_\eta\}$, we have $\eta'(h) = \mathbf{E}[X_1 \exp(hX_1 - \eta(h))]$, $\eta''(h) = \mathbf{E}[X_1^2 \exp(hX_1 - \eta(h))] - \mathbf{E}[X_1 \exp(hX_1 - \eta(h))]^2$. And ζ is the solution to the implicit equation $\eta'(\zeta) = \mu_2$ if $\mu_1 < \mu_2$, then $\Lambda^*(\mu_2) = \zeta\mu_2 - \eta(\zeta)$. We follow Dembo and Zeitouni (1998) and define a new probability measure \tilde{P} by $\frac{d\tilde{P}}{dP} = \exp(\zeta x - \eta(\zeta))$ and let $Z_i = \frac{X_i - \mu_2}{\sqrt{\eta''(\zeta)}}$. Then it follows that Z_i are i.i.d. random variables satisfying $\mathbf{E}_{\tilde{P}}[Z_i] = 0$, $\mathbf{E}_{\tilde{P}}[Z_i^2] = 1$, $\mathbf{E}_{\tilde{P}}[Z_i^3] := z_3 < \infty$. We also define $\psi_m := \zeta \sqrt{m\eta''(\zeta)}$, $W_m = \frac{\sum_{i=1}^m Z_i}{\sqrt{m}}$, F_m is the cumulative distribution function of W_n under measure \tilde{P} , then we know that $\bar{X}_m = \mu_2 + \sqrt{\eta''(\zeta)/m} W_m$. Therefore, we can calculate

$$\begin{aligned} &\mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\ &= \mathbf{E}_{\tilde{P}}[\exp(-m(\zeta\bar{X}_m - \eta(\zeta))) \cdot (-\sqrt{\eta''(\zeta)/m} W_m) \cdot \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\ &= -\exp(m(\zeta\mu_2 - \eta(\zeta))) \cdot \sqrt{\eta''(\zeta)/m} \cdot \mathbf{E}_{\tilde{P}}[\exp(-m\zeta\sqrt{\eta''(\zeta)/m} W_m) \cdot W_m \mathbf{1}_{\{W_m > 0\}}] \\ &= -\exp(-m\Lambda^*(\mu_2)) \cdot \sqrt{\eta''(\zeta)/m} \cdot \int_0^\infty \exp(-\psi_m x) x dF_m. \end{aligned}$$

Hence,

$$\begin{aligned} &J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\ &= \zeta^2 \sqrt{2\pi m\eta''(\zeta)} m \cdot \sqrt{\eta''(\zeta)/m} \cdot \int_0^\infty \exp(-\psi_m x) x dF_m \\ &= \sqrt{2\pi} \psi_m^2 \int_0^\infty \exp(-\psi_m x) x dF_m(x) \\ &= -\sqrt{2\pi} \psi_m^2 \int_0^\infty F_m(x) \exp(-\psi_m x) (-\psi_m x + 1) dx. \end{aligned}$$

Let $t = \psi_m x$, then

$$\begin{aligned} & J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\ &= -\sqrt{2\pi} \psi_m \int_0^\infty F_m(t/\psi_m) \exp(-t)(-t+1) dt \\ &= -\sqrt{2\pi} \psi_m \left[\int_0^\infty (F_m(t/\psi_m) - F_m(0)) \exp(-t)(-t+1) dx + \int_0^\infty F_m(0) \exp(-t)(-t+1) dt \right]. \end{aligned}$$

Note that $\int_0^\infty \exp(-t)(-t+1) dx = 0$, therefore, we have

$$J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] = -\sqrt{2\pi} \int_0^\infty \psi_m (F_m(t/\psi_m) - F_m(0)) \exp(-t)(-t+1) dt.$$

We now discuss the lattice and non-lattice cases separately.

(a) When X_1 has a non-lattice law, the Berry-Esseen expansion Feller (2008) of $F_m(x)$ yields

$$\lim_{m \rightarrow \infty} \left[\sqrt{m} \sup_x \left| F_m(x) - \Phi(x) - \frac{z_3}{6\sqrt{m}} (1-x^2) \phi(x) \right| \right] = 0,$$

where $z_3 = \mathbf{E}_{\tilde{P}}[Z_i^3] < \infty$.

Let $C_m := -\sqrt{2\pi} \int_0^\infty \psi_m e^{-t}(-t+1) \left[\Phi\left(\frac{t}{\psi_m}\right) + \frac{z_3}{6\sqrt{m}} \left(1 - \left(\frac{t}{\psi_m}\right)^2\right) \phi\left(\frac{t}{\psi_m}\right) - \Phi(0) - \frac{z_3}{6\sqrt{m}} \phi(0) \right] dt$. Note that $\psi_m = O(\sqrt{m})$, then we can get

$$\lim_{m \rightarrow \infty} |J_m(\mu_2) \mathbf{E}_\mu[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] - C_m| = 0.$$

Moreover,

$$\begin{aligned} \lim_{m \rightarrow \infty} C_m &= \lim_{m \rightarrow \infty} (-\sqrt{2\pi}) \int_0^\infty \psi_m e^{-t}(-t+1) \left[\Phi\left(\frac{t}{\psi_m}\right) - \Phi(0) \right] dt \\ &\quad + \lim_{m \rightarrow \infty} (-\sqrt{2\pi}) \int_0^\infty \psi_m e^{-t}(-t+1) \left[\frac{z_3}{6\sqrt{m}} \left(1 - \left(\frac{t}{\psi_m}\right)^2\right) \phi\left(\frac{t}{\psi_m}\right) - \frac{z_3}{6\sqrt{m}} \phi(0) \right] dt \\ &= \lim_{m \rightarrow \infty} (-\sqrt{2\pi}) \int_0^\infty \psi_m e^{-t}(-t+1) \left[\Phi\left(\frac{t}{\psi_m}\right) - \Phi(0) \right] dt \end{aligned}$$

Taylor expansion of $\Phi\left(\frac{t}{\psi_m}\right)$ shows that $\Phi\left(\frac{t}{\psi_m}\right) = \Phi(0) + \frac{t}{\psi_m} \phi(0) + O\left(\frac{t^2}{m}\right)$, where ϕ is the density function of $N(0, 1)$. Therefore by dominated convergence theorem we have

$$\lim_{m \rightarrow \infty} C_m = (-\sqrt{2\pi}) \int_0^\infty e^{-t}(-t+1)t\phi(0) dt = \sqrt{2\pi} \phi(0) = 1.$$

Thus, we obtain

$$\lim_{m \rightarrow \infty} J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] = 1.$$

(b) When X_1 has a lattice law, we have $Z_i = \frac{X_i - \mu_2}{\eta''(\zeta)} \in \left\{ \frac{md}{\eta''(\zeta)} : m \in \mathcal{Z} \right\}$. From the Berry-Esseen expansion Feller (2008), we have

$$\lim_{m \rightarrow \infty} \left[\sqrt{m} \sup_x \left| F_m(x) - \Phi(x) - \frac{z_3}{6\sqrt{m}} (1-x^2) \phi(x) - \phi(x) g\left(x, \frac{d}{\sqrt{\eta''(\zeta)m}}\right) \right| \right] = 0,$$

where $g(x, h) = \frac{h}{2} - (x \bmod h)$ if $(x \bmod h) \neq 0$ and $g(x, h) = -\frac{h}{2}$ if $(x \bmod h) = 0$. Then,

$$\begin{aligned} & \lim_{m \rightarrow \infty} J_m(\mu_2) \mathbf{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\ &= 1 + \lim_{m \rightarrow \infty} (-\sqrt{2\pi}) \int_0^\infty \psi_m e^{-t}(-t+1) \left[\phi\left(\frac{t}{\psi_m}\right) g\left(\frac{t}{\psi_m}, \frac{\zeta d}{\psi_m}\right) - \phi(0) g\left(0, \frac{\zeta d}{\psi_m}\right) \right] dt. \end{aligned}$$

Since $\psi_m g\left(\frac{t}{\psi_m}, \frac{\zeta d}{\psi_m}\right) = g(t, \zeta d)$, then

$$\begin{aligned}
 & \lim_{m \rightarrow \infty} J_m(\mu_2) \mathbb{E}_P[(\mu_2 - \bar{X}_m) \mathbf{1}_{\{\bar{X}_m > \mu_2\}}] \\
 &= 1 + \lim_{m \rightarrow \infty} (-\sqrt{2\pi}) \int_0^\infty e^{-t}(-t+1) \left[\phi\left(\frac{t}{\psi_m}\right) g(t, \zeta d) - \phi(0)g(0, \zeta d) \right] dt \\
 &= 1 - \sqrt{2\pi}\phi(0) \int_0^\infty e^{-t}(-t+1) [g(t, \zeta d) - g(0, \zeta d)] dt \\
 &= 1 - \left(\sum_{n=0}^\infty e^{-n\zeta d} \right) \int_0^{\zeta d} e^{-t}(-t+1)(\zeta d - t) dt \\
 &= 1 - \left[\frac{1}{1 - e^{-\zeta d}} \cdot (1 - e^{-\zeta d} - \zeta d e^{-\zeta d}) \right] \\
 &= \frac{\zeta d e^{-\zeta d}}{1 - e^{-\zeta d}}.
 \end{aligned}$$

The proof is hence complete. \blacksquare

F Proof of Theorem 2

Proof. First, in the proof of Theorem 1 (see (12)), we have established that $|\hat{\mu}_1 - \mu_1| + |\hat{\sigma}_1^2 - \sigma_1^2| \xrightarrow{P} 0$ as $m \rightarrow \infty$, where $\sigma_1^2 = \text{Var}(X_1)$, the variance of the reward from Arm 1. This result does not depend on the assumption of Gaussian rewards. It then follows from the continuous mapping theorem (Theorem 7.10 DasGupta (2011)) that

$$\hat{\Lambda}^*(\mu_2) = \frac{1}{2\hat{\sigma}_1^2} (\hat{\mu}_1 - \mu_2)^2 \xrightarrow{P} \frac{(\mu_2 - \mu_1)^2}{2\sigma_1^2}, \quad \text{as } m \rightarrow \infty.$$

Next, we show that $\Lambda^*(\mu_2) \geq \frac{(\mu_2 - \mu_1)^2}{2a^2}$, where $\Lambda^*(\mu_2)$ is given in (5). Recall that $X_1 - \mu_1$ is sub-Gaussian with variance proxy parameter a^2 , i.e. $\mathbb{E}[e^{\lambda(X_1 - \mu_1)}] \leq \exp(\frac{a^2}{2}\lambda^2)$ for all $\lambda \in \mathbb{R}$. It follows that when $\mu_1 < \mu_2$ we have the Chernoff-Hoeffding bound:

$$\mathbb{P}(\bar{X}_m \geq \mu_2) \leq e^{-\frac{m(\mu_2 - \mu_1)^2}{2a^2}},$$

where $\bar{X}_m = \frac{1}{m} \sum_{i=1}^m X_i$ is the average of m i.i.d sub-Gaussian rewards from Arm 1. Hence,

$$\lim_{m \rightarrow \infty} \frac{1}{m} \log P(\bar{X}_m \geq \mu_2) \leq -\frac{(\mu_2 - \mu_1)^2}{2a^2}.$$

On the other hand, from formula (15), we know that

$$\lim_{m \rightarrow \infty} \frac{1}{m} \log P(\bar{X}_m \geq \mu_2) = \lim_{m \rightarrow \infty} \frac{1}{m} \log \left(\frac{c_* + o(1)}{\sqrt{2\pi m \eta''(\zeta)}} \right) - \Lambda^*(\mu_2) = -\Lambda^*(\mu_2).$$

Therefore, we have

$$\Lambda^*(\mu_2) \geq \frac{(\mu_2 - \mu_1)^2}{2a^2} = \frac{(\mu_2 - \mu_1)^2}{2\sigma_1^2} \cdot \frac{\sigma_1^2}{a^2}.$$

When $\mu_1 > \mu_2$, the proof is similar by studying $P(\bar{X}_m \leq \mu_2)$, and hence we omit the details. The proof is complete. \blacksquare