
On Conditional Versus Marginal Bias in Multi-Armed Bandits

Jaehyeok Shin¹ Alessandro Rinaldo¹ Aaditya Ramdas^{1,2}

Abstract

The bias of the sample means of the arms in multi-armed bandits is an important issue in adaptive data analysis that has recently received considerable attention in the literature. Existing results relate in precise ways the sign and magnitude of the bias to various sources of data adaptivity, but do not apply to the conditional inference setting in which the sample means are computed only if some specific conditions are satisfied. In this paper, we characterize the sign of the conditional bias of monotone functions of the rewards, including the sample mean. Our results hold for arbitrary conditioning events and leverage natural monotonicity properties of the data collection policy. We further demonstrate, through several examples from sequential testing and best arm identification, that the sign of the conditional and marginal bias of the sample mean of an arm can be different, depending on the conditioning event. Our analysis offers new and interesting perspectives on the subtleties of assessing the bias in data adaptive settings.

1. Introduction

In modern data analysis, it is often the case that both the data collection and the selection of the target for statistical inference are determined adaptively, i.e. in a manner that depends on the realization of the data observed so far. In these fairly common data-adaptive scenarios, the validity of traditional procedures for statistical inference (that were designed to work in nonadaptive or in i.i.d. settings) can be severely affected.

This issue is perhaps best exemplified within the stochastic multi-armed bandit (MAB) framework (Robbins, 1952).

¹Department of Statistics and Data Science, Carnegie Mellon University ²Machine Learning Department, Carnegie Mellon University. Correspondence to: Jaehyeok Shin <shinjaehyeok@cmu.edu>.

The data are collected sequentially in stages, during which the analyst draws a sample from one among finitely many unknown distributions or arms. At each stage, the determination of whether no more samples are to be acquired (*adaptive stopping*), the choice of the arm to draw from (*adaptive sampling*) and the selection of the target arm to analyze once the experiment has terminated (*adaptive choosing*) are all made based on previously observed data.

The combined effect of these sources of data adaptivity will impact the correctness of standard, nonadaptive statistical procedures in ways that are generally difficult to discern and to quantify. The recent literature on adaptive data analysis in MAB settings has focused primarily on the fundamental problem of mean estimation (Xu et al., 2013; Villar et al., 2015; Bowden & Trippa, 2017; Nie et al., 2018; Shin et al., 2019a). It is now known that the sample mean of an arm is a biased estimator of the true mean and that the sign and magnitude of the bias can be related to adaptive sampling, stopping and choosing in precise ways.

The current theoretical understanding of the bias in MAB experiments is, however limited in two aspects. First, virtually all results concern the bias in mean estimation and do not cover other functionals of the arms. Secondly, and perhaps more interestingly, the existing guarantees cover only the *marginal bias* of the sample mean, i.e., the bias obtained by accounting for all possible outcomes of the MAB experiment. However, in practice, one is often interested in the sample means of the arms only when certain outcomes have occurred. For instance, the analyst may wish to evaluate the sample mean of a given arm only when that arm was identified as the best arm or, in a sequential framework corresponding to a MAB experiment with only one arm, when the null hypothesis has been rejected or when the random criterion for determining whether enough samples have been collected has been met. In all these cases, it is of interest to compute the *conditional bias*, i.e., the bias of the sample mean given a certain *conditioning event*, such as that the arm of interest turned out to be the best arm. A priori, it is not at all clear how the sign of the conditional bias is affected by the choice of the conditioning event and by other sources of data adaptivity (e.g., sampling and stopping), or whether the signs of the marginal and conditional bias should match. See Figure 1 for illustrations of marginal and conditional biases.




Figure 1. Illustrations of marginal and conditional biases in the K prototypes example. (a) Marginal bias is obtained by accounting for all possible outcomes. (b) Conditional bias is obtained by accounting for outcomes related to a conditioning event.

As a concrete example, suppose we have K prototypes of an online service and wish to test whether the potential average revenue of each prototype (i.e., arm), μ_k , would be larger than a pre-specified threshold $\mu_0 > 0$ based on a stream of test user samples. The usual sequential testing approach will be based on an appropriate upper stopping boundary and will reject each null $\mu_k \leq \mu_0$ if the corresponding sample mean has crossed such boundary during the testing period. If the null is rejected, then we conclude the prototype as a promising one. If not, we disregard it. It is well known (see, e.g., Starr & Woodroffe, 1968; Shin et al., 2019a) that for each prototype, the sample mean based on data collected through this sequential testing procedure is positively biased: that is, for each $k \in [K] := \{1, \dots, K\}$, $\mathbb{E} [\hat{\mu}_k] \geq \mu_k$ regardless of whether the true mean μ_k is larger or smaller than the threshold μ_0 . This positive bias result can provide a useful warning signal about possible overestimation of the true revenues. At the same time, however, without careful consideration of the conditioning effect, we may end up with a false sense of comfort with low sample mean estimates from “disregarded” prototypes. Indeed, we can naively expect the true revenues of the disregarded prototypes should be even lower than the observed estimates based on the positive bias result $\mathbb{E} [\hat{\mu}_k] \geq \mu_k$. In fact, this would be a wrong conclusion: conditioned on the disregarding event C , the sample mean is negatively biased and we have $\mathbb{E} [\hat{\mu}_k \mid C] \leq \mu_k$, as demonstrated in Section 4.1.

In this paper, we make the following contributions:

- We derive in Theorem 2 sufficient conditions for de-

termining the sign of the conditional bias that hold under arbitrary conditioning events and rely on certain natural, highly interpretable monotonicity properties of the rules used for adaptive sampling, stopping and selecting. Our analysis captures in a rigorous and intuitive way the interaction between adaptivity arising from the data collection stage (sampling and stopping) with adaptivity in the choice of the target for inference.

- We characterize the sign of the conditional bias of monotone functions of the rewards of each arm, which includes the sample mean as a special case.
- In Section 4, we demonstrate with several examples in best arm identification and sequential testing problems how the marginal and conditional bias of the sample means of the arms can have opposite signs. These are, we believe, instances of a general, important phenomenon of theoretical and practical relevance.

Overall, our results advance our ability to assess the impact of the bias in adaptive data analysis problems and offer several new and interesting insights on this important issue.

The rest of this paper is organized as follows. In Section 2, we briefly formalize the three notions of adaptivity by introducing a stochastic MAB framework. Section 3 derives results on when the bias can be positive or negative. In Section 4, we demonstrate the correctness of our theoretical predictions through simulations in a variety of practical situations. We end with a brief summary in Section 5, and for reasons of space, we defer all proofs to the Appendix.

2. The Stochastic Multi-Armed Bandit Model

In this section we formally describe the MAB experiment setting, detailing the sources of data adaptivity. Let P_1, \dots, P_K be K distributions of K arms with finite means $\mu_k = \int x dP_k(x)$. If it is clear from context, every inequality and equality between two random variables is understood in the almost sure sense.

2.1. Three Sources of Adaptivity in MABs

A general stochastic MAB algorithm consists of the following sampling, stopping, and choosing rules:

- **Sampling:** For each time $t \geq 1$, choose a possibly randomized index of arm A_t from a multinomial distribution $\text{Multi}(1, \nu_t) \in \{0, 1, \dots, K\}$. Then, we draw a sample (also called a reward) Y_t from the chosen arm P_{A_t} . Here, $\nu_t \in [0, 1]^K$ refers to conditional probabilities of sampling each arm defined by

$$\nu_t(k) := \nu_t(k \mid \mathcal{D}_{t-1}) \in [0, 1] \quad \forall k \in [K],$$

with the constraint $\sum_{k=1}^K \nu_t(k) = 1$, and \mathcal{D}_{t-1} refers to observed data up to time $t - 1$ including all external randomness. If each ν_t does not depend on the previous data \mathcal{D}_{t-1} , we call it a nonadaptive sampling rule.

With a proper choice of ν_t , this sampling rule captures a broad class of existing methods including fixed design, random allocation, ϵ -greedy (Sutton & Barto, 1998), upper confidence bound (UCB) algorithms (Auer et al., 2002; Audibert & Bubeck, 2009; Garivier & Cappé, 2011; Kalyanakrishnan et al., 2012; Jamieson et al., 2014) and Thompson sampling (Thompson, 1933; Agrawal & Goyal, 2012; Kaufmann et al., 2012).

- **Stopping:** Given a sampling rule, let $\{\mathcal{F}_t\}$ be a filtration such that each A_t is \mathcal{F}_{t-1} -measurable and each Y_t and \mathcal{D}_t are \mathcal{F}_t -measurable. Let \mathcal{T} be a stopping time with respect to the filtration $\{\mathcal{F}_t\}$. Then, for each time t , the stopping event $\{\mathcal{T} = t\}$ is \mathcal{F}_t -measurable which can be used to characterize a stopping rule of a MAB algorithm.

For example, the following stopping time corresponds to the stopping rule in which we stop whenever the absolute difference between sample means for arm 1 and arm 2 exceeding a fixed threshold $\delta > 0$:

$$\mathcal{T} := \inf \{t \geq 1 : |\hat{\mu}_1(t) - \hat{\mu}_2(t)| > \delta\}.$$

Here, we refer $\hat{\mu}_k(t)$ to the sample mean for arm k defined as

$$\hat{\mu}_k(t) := \frac{S_k(t)}{N_k(t)},$$

where $S_k(t)$ and $N_k(t)$ are the running sum and number of draws for arm k defined respectively $S_k(t) := \sum_{s=1}^t \mathbb{1}(A_s = k)Y_s$ and $N_k(t) := \sum_{s=1}^t \mathbb{1}(A_s = k)$ for each $k \in [K]$ and $t \geq 1$.¹ Note that, by the construction, $S_k(t)$, $\hat{\mu}_k(t)$ are \mathcal{F}_t -measurable and $N_k(t)$ is \mathcal{F}_{t-1} -measurable for each $k \in [K]$ and $t \geq 1$.

If a stopping rule is nonadaptive, that is, if the corresponding stopping time is \mathcal{F}_0 -measurable, we use T instead of \mathcal{T} to refer the stopping time which includes all fixed time and, more generally, all data-independent stopping rules.

- **Choosing:** After stopping, let $\mathcal{D}_{\mathcal{T}}$ denote collected data up to the stopping time \mathcal{T} . Based on $\mathcal{D}_{\mathcal{T}}$ with possible randomization, we choose a data-dependent arm based on a choosing rule $\kappa : \mathcal{D}_{\mathcal{T}} \mapsto [K]$. In this paper, we denote the index of chosen arm $\kappa(\mathcal{D}_{\mathcal{T}})$ by κ for simplicity.

For example, in the best arm identification, we often choose the best arm as the arm with the largest sample mean at the stopping time \mathcal{T} . In this case, the

choosing rule is given as $\kappa = \arg \max_{k \in [K]} \hat{\mu}_k(\mathcal{T})$. If κ does not depend on the observed data, we call it a nonadaptive choosing rule.

2.2. The Tabular Model of Stochastic MABs

To analyze the theoretical properties of a MAB experiment, it is convenient to express it using an infinite tabular representation of the arm rewards (Section 2.2 of Shin et al. (2019a), Lattimore & Szepesvári (2019, Section 4.6)).

In the tabular model, we assume that each observation Y_t comes from an imaginary $\mathbb{N} \times K$ table or stacks of samples $\{X_{i,k}^*\}_{i \in \mathbb{N}, k \in [K]} := X_{\infty}^*$ where each $X_{i,k}^*$ is an independent draw from P_k , and the observation Y_t is equal to $X_{i,k}^*$ if $N_k(t) = i$ and $A_t = k$.

Now, given a MAB algorithm, let W_t be a random variable which accounts for all possible external randomness of the algorithm for each time t . Next, we define $\mathcal{D}_{\infty}^* = X_{\infty}^* \cup \{W_t\}_{t=0}^{\infty}$ as the collection of all possible arm rewards and external randomness. Then, given \mathcal{D}_{∞}^* , every property and outcome of a MAB experiment becomes a deterministic function of \mathcal{D}_{∞}^* . In particular \mathcal{T} , κ and $N_k(t)$, for each t and k , can be expressed as . See Section 2.2 in Shin et al. (2019a) for details.

3. Monotonicity Property and Signs of Biases

In a recent work by Shin et al. (2019a), a notion of monotonicity of the data collecting was introduced to characterize the marginal bias of a chosen sample mean at a stopping time. In this section, we summarize the previous result with some minor modifications in notations.

Definition 1. For each $k \in [K]$, we say a MAB algorithm satisfies the “monotonic increasing (or decreasing) property for arm k ” if for any $i \in \mathbb{N}$, the function $\mathcal{D}_{\infty}^* \mapsto \mathbb{1}(\kappa = k) / N_k(\mathcal{T})$, is an increasing (or decreasing) function of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_{∞}^* fixed. Further, we say that

- A sampling rule is optimistic for arm k if the function $\mathcal{D}_{\infty}^* \mapsto N_k(t)$ is an increasing function of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_{∞}^* fixed for any fixed $i \in \mathbb{N}$ and $t \geq 1$;
- A stopping rule is optimistic for arm k if the function $\mathcal{D}_{\infty}^* \mapsto \mathcal{T}$ is a decreasing function of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_{∞}^* fixed for any fixed $i \in \mathbb{N}$;
- A choosing rule is optimistic for arm k if the function $\mathcal{D}_{\infty}^* \mapsto \mathbb{1}(\kappa = k)$ is an increasing function of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_{∞}^* fixed for any fixed $i \in \mathbb{N}$.

It can be shown that many MAB algorithms, including ϵ -greedy, UCB, and Thompson sampling algorithms, have

¹Throughout this paper, we assume $N_k(t) \geq 1$ whenever we make a statement about the sample mean $\hat{\mu}_k(t)$.

optimistic sampling rules. As for optimistic stopping, a sequential testing procedure with an upper stopping boundary yields a typical example of an optimistic stopping rule based on the stopping time $\mathcal{T} := \inf \{t : \hat{\mu}(t) > U(t)\}$ where $t \mapsto U(t)$ is a deterministic function characterizing the upper stopping boundary. Finally, the index of the arm with the largest sample mean, $\kappa = \arg \max_{k \in [K]} \hat{\mu}_k(\mathcal{T})$ can be viewed as a typical optimistic choosing rule.

3.1. The Sign of the Marginal Bias

Shin et al. (2019a) demonstrated how certain natural monotonicity properties of the MAB experiment can be exploited to determine the sign of the marginal bias and of the conditional bias given conditioning events corresponding to adaptive choosing of arms. For completeness, we recall this result next.

Proposition 1 (Theorem 7 in Shin et al. (2019a)). *Given a MAB algorithm, suppose we have $\mathbb{E}[N_k(\mathcal{T})] < \infty$ and $N_k(\mathcal{T}) \geq 1$ for all k with $\mathbb{P}(\kappa = k) > 0$. If the MAB algorithm satisfies the monotonic decreasing property for each arm k with $\mathbb{P}(\kappa = k) > 0$ then we have*

$$\mathbb{E}[\hat{\mu}_k(\mathcal{T}) \mid \kappa = k] \leq \mu_k, \quad (1)$$

which also implies that

$$\mathbb{E}[\hat{\mu}_\kappa(\mathcal{T}) - \mu_\kappa] \leq 0. \quad (2)$$

Similarly if the MAB algorithm satisfies the monotonic increasing property for each arm k with $\mathbb{P}(\kappa = k) > 0$ then we have

$$\mathbb{E}[\hat{\mu}_k(\mathcal{T}) \mid \kappa = k] \geq \mu_k, \quad (3)$$

which also implies that

$$\mathbb{E}[\hat{\mu}_\kappa(\mathcal{T}) - \mu_\kappa] \geq 0. \quad (4)$$

If each arm has a bounded distribution then the condition $\mathbb{E}[N_k(\mathcal{T})] < \infty$ can be dropped.

Note that if a MAB algorithm consists of an optimistic sampling but nonadaptive stopping and choosing rules, then the algorithm satisfies the monotonic decreasing property. This is the reason why ϵ -greedy, UCB, and Thompson sampling algorithms yield a negatively biased sample mean for each arm, marginally.

Also, if a MAB algorithm consists of an optimistic stopping but nonadaptive sampling and choosing, the algorithm satisfies the monotonic increasing property, which explains why sequential testing procedures (corresponding to one-arm MBA experiments) with an upper stopping boundary produces a positively biased sample mean.

3.2. The Sign of the Conditional Bias

In this section, we generalize Proposition 1 and derive sufficient conditions for the sign of the conditional bias under

arbitrary conditioning events of monotone functions of the arms rewards.

We first introduce some notation. For each arm k , let F_k denote the corresponding cumulative distribution function (CDF): $y \in \mathbb{R} \mapsto F_k(y) = \mathbb{P}(X \leq y)$, where $X \sim P_k$. Let $\hat{F}_{k,t}$ be the empirical CDF for arm k based on samples up to time t ; that is, $\hat{F}_{k,t}$ is a random function on \mathbb{R} and taking values in $[0, 1]$ given by

$$y \in \mathbb{R} \mapsto \hat{F}_{k,t}(y) := \frac{1}{N_k(t)} \sum_{s=1}^t \mathbb{1}(A_s = k, Y_s \leq y), \quad (5)$$

which returns the fractions of samples from arm k whose values are no larger than y . For a stopping time \mathcal{T} with respect to the MAB filtration $\{\mathcal{F}_t\}$, we then define $\hat{F}_{k,\mathcal{T}}$ to be the empirical CDF of the rewards of arm k up to time \mathcal{T} ; that is, for each $y \in \mathbb{R}$,

$$\hat{F}_{k,\mathcal{T}}(y) := \lim_{t \rightarrow \infty} \hat{F}_{k,\mathcal{T} \wedge t}(y).$$

The use of the limit in the above definition is a necessary technicality in order to allow for the possibility that $\mathcal{T} = \infty$. Note that, on the event $\{N_k(\mathcal{T}) = \infty\} = \{\lim_{t \rightarrow \infty} N_k(\mathcal{T} \wedge t) = \infty\}$, we have that, for each $y \in \mathbb{R}$, $\hat{F}_{k,\mathcal{T}}(y) = F_k(y)$ almost surely, based on the strong law of large numbers; see Theorem 2.1 in Gut (2009).

Next, for any function $f : \mathbb{R} \rightarrow \mathbb{R}$ integrable with respect to P_k , let $E_k f = \int f(x) dP_k(x)$ be the corresponding expectation. Similarly, we let

$$\hat{E}_{k,t} f = \frac{1}{N_k(t)} \sum_{s=1}^t f(Y_s) \mathbb{1}(A_s = k)$$

to be the expectation of f under the empirical measure of the k -th arm at time t . Clearly, $\hat{E}_{k,t} f$ is a random variable. If f is the identity function, then it is immediate to see that $E_k f = \mu_k$ and $\hat{E}_{k,t} f = \hat{\mu}_k(t)$. Also, for any $y \in \mathbb{R}$, setting f to be the indicator function $x \mapsto f(x) = \mathbb{1}(x \leq y)$ yields that $E_k f = F_k(y)$ and $\hat{E}_{k,t} f = \hat{F}_{k,t}(y)$. Finally, for a (possibly infinite) stopping time \mathcal{T} with respect to the filtration $\{\mathcal{F}_t\}$, we set

$$\hat{E}_{k,\mathcal{T}} f = \lim_{t \rightarrow \infty} \hat{E}_{k,\mathcal{T} \wedge t} f.$$

We are now ready to present the main result of this paper. The proof is given in Appendix A.1. For any event C , we let $\mathbb{E}[\cdot | C]$ denote the conditional expectation given C .

Theorem 2. *Let \mathcal{T} be a stopping time with respect to the MAB experiment natural filtration $\{\mathcal{F}_t\}$. For a fixed $k \in [K]$, suppose $N_k(\mathcal{T}) \geq 1$. Let $C \in \mathcal{F}_\mathcal{T}$ be any event at stopping time \mathcal{T} with $\mathbb{P}(C) > 0$. Assume that, for each i , the function $\mathcal{D}_\infty^* \mapsto \mathbb{1}(C) / N_k(\mathcal{T})$ is a decreasing function*

of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_∞^* fixed. Then, for each $y \in \mathbb{R}$,

$$\inf_{y \in \mathbb{R}} \left(\mathbb{E} \left[\widehat{F}_{k,\mathcal{T}}(y) \mid C \right] - F_k(y) \right) \geq 0, \quad (6)$$

or, equivalently, for any non-decreasing function $f : \mathbb{R} \rightarrow \mathbb{R}$ with $\mathbb{E} \left[\widehat{E}_{k,\mathcal{T}} | f \mid C \right] < \infty$ and $E_{k,\mathcal{T}} | f| < \infty$,

$$\mathbb{E} \left[\widehat{E}_{k,\mathcal{T}} f \mid C \right] \leq E_k f. \quad (7)$$

Similarly, if, for each i , the function $\mathcal{D}_\infty^* \mapsto \mathbb{1}(C) / N_k(\mathcal{T})$ is an increasing function of $X_{i,k}^*$ while keeping all other entries in \mathcal{D}_∞^* fixed, then we have

$$\sup_{y \in \mathbb{R}} \left(\mathbb{E} \left[\widehat{F}_{k,\mathcal{T}}(y) \mid C \right] - F_k(y) \right) \leq 0, \quad (8)$$

or, equivalently,

$$\mathbb{E} \left[\widehat{E}_{k,\mathcal{T}} f \mid C \right] \geq E_k f, \quad (9)$$

for any non-decreasing function $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\mathbb{E} \left[\widehat{E}_{k,\mathcal{T}} | f \mid C \right] < \infty$ and $E_{k,\mathcal{T}} | f| < \infty$.

The important conclusion from the above theorem is that the conditional expected value of the empirical CDF of the rewards of any given arm can be stochastically smaller or greater than the true CDF of the corresponding arm. And furthermore, the sign of the bias can be determined in the basis of natural and often verifiable (as shown below) monotonicity conditions that depend on (i) the specific sampling and stopping rules deployed in the MAB experiment and (ii) the choice of the conditioning event C .

Note that if in the theorem we choose C as the conditioning event $\{\kappa = k\}$ and the function f as the identity function, Theorem 2 yields Proposition 1 of Shin et al. (2019a) as a special case (indeed the statements about the marginal bias in that result follows from those on the conditional bias given the conditioning event $\{\kappa = k\}$). We emphasize that Theorem 2 is a generalization of Proposition 1 of Shin et al. (2019a) in at least two ways: (i) it provides a more general guarantee about the conditional bias of monotone functions of the empirical CDF of the arms as opposed to just the sample mean and (ii) it allows for virtually any conditioning event that depends on the outcome of the MAB experiment (formally, that is measurable with respect to $\mathcal{F}_\mathcal{T}$).

The monotonicity assumption in Theorem 2 captures in a mathematically concise yet intuitive manner how adaptivity in the data collection process combines with adaptivity in the selective data analysis, exemplified by the conditioning event, to affect the sign of the bias. The interaction between these two sources of adaptivity is, at least to us, not at all apparent and, in fact, rather subtle. To illustrate this

phenomenon, in the next section, we apply Theorem 2 in several, fairly routine, situations in adaptive data analysis to demonstrate how marginal and conditional biases may very well have opposite signs.

4. Applications

In this section, we discuss several practical examples of the conditional bias results in Theorem 2.

4.1. Conditional Versus Marginal Bias of a Stopped Sequential Test

Suppose we have a stream of samples from a single arm with a finite mean μ . For each t , let $\widehat{\mu}(t)$ be the sample mean of the rewards observed up to time t .

To test whether μ is larger than a reference value μ_0 , we may construct an upper boundary $t \mapsto U(t)$ and conclude that $\mu \geq \mu_0$ if the sample mean ever crosses the upper boundary.

Let \mathcal{T} be the first time the sample mean crosses the boundary, i.e. $\mathcal{T} := \inf \{t \geq 1 : \widehat{\mu}(t) \geq U(t)\}$. The stopping time \mathcal{T} is an example of an optimistic stopping and, from Proposition 1, we can check that the stopped sample mean $\widehat{\mu}(\mathcal{T})$ is always positively biased (Shin et al., 2019a).

However, the stopping time \mathcal{T} can be large or even infinite with non-zero probability. Thus, in practice, we may want to allow for the possibility of stopping the sequential test before reaching the stopping time. Let M be any fixed pre-defined time at which we stop the testing procedure (if still ongoing), and let $\mathcal{T}_M := \min \{\mathcal{T}, M\}$ be the corresponding stopping time which takes account the early stopping option. Again, by Proposition 1, we know that the stopped sample mean $\widehat{\mu}(\mathcal{T}_M)$ is still positively biased, i.e. $\mathbb{E} [\widehat{\mu}(\mathcal{T}_M)] \geq \mu$, since the function $\mathcal{D}_\infty^* \mapsto 1/\mathcal{T}_M$ is an increasing function of X_i^* while keeping all other entries in \mathcal{D}_∞^* fixed. From Theorem 2, we can also check that the expected empirical CDF at stopping time \mathcal{T}_M is negatively biased:

$$\sup_{y \in \mathbb{R}} \left(\mathbb{E} \left[\widehat{F}_{k,\mathcal{T}_M}(y) \right] - F_k(y) \right) \leq 0. \quad (10)$$

Conditioned on the early stopping event $\{M < \mathcal{T}\}$, however, Theorem 2 shows that the early stopped sample mean and the empirical CDF are negatively and positively biased, respectively, since the function $\mathcal{D}_\infty^* \mapsto \frac{\mathbb{1}(M < \mathcal{T})}{\mathcal{T}_M} = \frac{\mathbb{1}(M < \mathcal{T})}{M}$ is an decreasing function of X_i^* while keeping all other entries in \mathcal{D}_∞^* fixed. That is,

$$\mathbb{E} [\widehat{\mu}(M) \mid M < \mathcal{T}] \leq \mu, \quad (11)$$

$$\inf_{y \in \mathbb{R}} \left(\mathbb{E} \left[\widehat{F}_{k,M}(y) \mid M < \mathcal{T} \right] - F_k(y) \right) \geq 0. \quad (12)$$

However, we can also check that the function $\mathcal{D}_\infty^* \mapsto \frac{\mathbb{1}(M \geq \mathcal{T})}{\mathcal{T}_M} = \frac{\mathbb{1}(M \geq \mathcal{T})}{\mathcal{T}}$ is an increasing function of X_i^* , which




Figure 2. Average of marginal (red dashed line) and conditional empirical CDFs (green dotted and blue dot-dashed lines) from repeated stopped sequential test in Section 4.1. The black solid line refers to the true CDF of the underlying arm

implies that, on the line-crossing event $\{M \geq \mathcal{T}\}$, the sample mean and empirical CDF are positively and negatively biased, respectively. Thus, depending on the conditioning event, the conditional bias can be positive or negative. Also note that, without the early stopping condition, $\hat{\mu}(M)$ is an unbiased estimator of μ .

Experiment: We verify these facts with simulations, where we repeat a stopped sequential test 10^5 times. In each test, the arm corresponds to a standard normal distribution, and the upper boundary is constructed by the point-wise upper confidence bounds $t \mapsto U(t) = \frac{z_\alpha}{\sqrt{t}}$, where z_α is the α -upper quantile of the standard normal distribution. Note that this upper boundary does not yield a valid testing procedure since it inflates the type 1 error. However, we choose this boundary with unusual parameters $\alpha = 0.2$ and $M = 10$ to manifest the difference between marginal and conditional biases.

Figure 2 shows the point-wise averages of the observed marginal and conditional CDFs from the simulated 10^5 stopped sequential tests. The black solid line refers to the true CDF of the underlying arm. The red dashed line corresponds to the marginal CDF, which lies below the true CDF, as expected, and thus the corresponding mean bias is positive. The green dotted and blue dot-dashed lines correspond to the averages of empirical CDFs conditioned on the early stopping and line-crossing events, respectively. As predicted, conditioned on the early stopping event, the empirical CDF and the sample mean are positively and negatively biased, respectively. In contrast, conditioned on the line-crossing event, the empirical CDF and the sample mean

are negatively and positively biased, respectively.

4.2. Sequential Test for Two Arms: Conditional Biases from Upper and Lower Stopping Boundaries

Suppose we have two independent arms with unknown means μ_1 and μ_2 . In this subsection, we consider the following testing problem:

$$H_0 : \mu_1 \leq \mu_2 \text{ vs } H_1 : \mu_1 > \mu_2. \quad (13)$$

To test this hypothesis, we draw a sample from arm 1 for every odd time and from arm 2 for every even time. Then, at each even time t , we check whether the difference between the sample means $\hat{\mu}_1(t), \hat{\mu}_2(t)$ from the two arms crosses predefined upper and lower stopping boundaries, $t \mapsto U(t)$ and $t \mapsto L(t)$.

To be specific, define stopping times \mathcal{T}^U and \mathcal{T}^L as follows:

$$\mathcal{T}^U := \inf \{t \in \mathbb{N}_{\text{even}} : \hat{\mu}_1(t) - \hat{\mu}_2(t) \geq U(t)\}, \quad (14)$$

$$\mathcal{T}^L := \inf \{t \in \mathbb{N}_{\text{even}} : \hat{\mu}_1(t) - \hat{\mu}_2(t) \leq L(t)\}. \quad (15)$$

Let $M > 0$ be a predetermined maximum time budget. Based on $\mathcal{T}^U, \mathcal{T}^L$ and M , we stop sampling whenever $\hat{\mu}_1(t) - \hat{\mu}_2(t)$ crosses one of the boundaries or the maximum time budget M is met. Define the corresponding stopping time as $\mathcal{T}_M := \min \{\mathcal{T}^U, \mathcal{T}^L, M\}$. At time \mathcal{T}_M , we accept H_1 if $\mathcal{T}_M = \mathcal{T}^U$ (upper-crossing event), and accept H_0 if $\mathcal{T}_M = \mathcal{T}^L$ (lower-crossing event). Otherwise, we declare that we do not have enough evidence to accept either one of two hypotheses.

In this case, we cannot apply Proposition 1 to determine the sign of the marginal bias since the stopping rule is neither optimistic nor pessimistic, marginally. However, we can determine the sign of the conditional bias based on Theorem 2 since the function $\mathcal{D}_\infty^* \mapsto \frac{\mathbb{1}(\mathcal{T}^U \leq \min\{\mathcal{T}^L, M\})}{\mathcal{T}^U}$ is an increasing (resp. decreasing) function of $X_{i,1}^*$ (resp. $X_{i,2}^*$) for each i , keeping all other entries in \mathcal{D}_∞^* fixed.

In detail, conditioned on the event of accepting H_0 , the sample mean and empirical CDF for arm 1 are negatively and positively biased, respectively. That is,

$$\begin{aligned} & \mathbb{E} [\hat{\mu}_1(\mathcal{T}_M) \mid \mathcal{T}_M = \mathcal{T}^L] \\ &= \mathbb{E} [\hat{\mu}_1(\mathcal{T}^L) \mid \mathcal{T}^L \leq \min \{\mathcal{T}^U, M\}] \leq \mu_1, \end{aligned} \quad (16)$$

$$\begin{aligned} & \inf_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k, \mathcal{T}_M}(y) \mid \mathcal{T}_M = \mathcal{T}^L] - F_k(y) \right) \\ &= \inf_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k, \mathcal{T}^L}(y) \mid \mathcal{T}^L \leq \min \{\mathcal{T}^U, M\}] - F_k(y) \right) \\ &\geq 0. \end{aligned} \quad (17)$$

Similarly, for arm 2, we have opposite signs of the sample mean and empirical CDF.

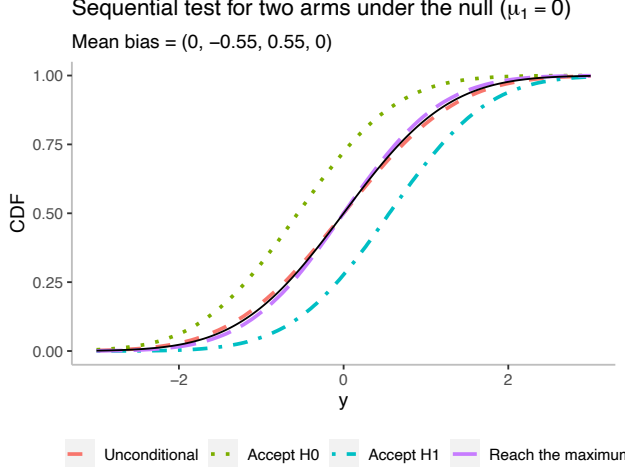


Figure 3. Average of marginal and conditional empirical CDFs of arm 1 from repeated sequential tests for two arms in Section 4.2. The black solid line refers to the true CDF of the underlying arm.

By the same reasoning, the signs of the conditional biases conditioned on the event of accepting H_1 are reversed:

$$\begin{aligned} \mathbb{E} [\hat{\mu}_1(\mathcal{T}_M) \mid \mathcal{T}_M = \mathcal{T}^U] \\ = \mathbb{E} [\hat{\mu}_1(\mathcal{T}^U) \mid \mathcal{T}^U \leq \min \{ \mathcal{T}^L, M \}] \geq \mu_1, \end{aligned} \quad (18)$$

$$\begin{aligned} \sup_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k, \mathcal{T}_M}(y) \mid \mathcal{T}_M = \mathcal{T}^U] - F_k(y) \right) \\ = \sup_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k, \mathcal{T}^U}(y) \mid \mathcal{T}^U \leq \min \{ \mathcal{T}^L, M \}] - F_k(y) \right) \\ \leq 0. \end{aligned} \quad (19)$$

Thus, we conclude that, in all cases, the expected difference between the sample means are exaggerated toward “the direction of decision”.

Remark 3. *These results hold regardless of whether the underlying distribution is under the null or alternative.*

Experiment: To demonstrate the conditional bias result, we set two standard normal arms with same means $\mu_1 = \mu_2 = 0$. In this experiment, we use upper and lower stopping boundaries based on naive point-wise confidence intervals:

$$U(t) := z_{\alpha/2} \sqrt{\frac{2}{t}}, \quad \text{and} \quad L(t) = -U(t), \quad (20)$$

where α is set to 0.2 to show the bias better. Figure 3 show the marginal and conditional biases of the empirical CDFs and sample means for arm 1 based on 10^5 repetitions of the experiment. The black solid line corresponds to the true underlying CDF. The red dashed line refers to the average

of the marginal CDFs, and the purple long-dashed line corresponds to the average of the empirical CDFs conditioned on reaching the maximal time. Note that for these two cases, the empirical CDFs are neither positively nor negatively biased across all $y \in \mathbb{R}$.

However, for the cases corresponding to accepting H_0 (green dotted line) and accepting H_1 (blue dot-dashed line), the empirical CDFs are positively and negatively biased, respectively: see inequalities (17) and (19). The conditional bias of the sample mean is also negative conditioning on the event of accepting H_0 and positive conditioning on the event of accepting H_1 : see inequalities (16) and (18).

4.3. Best-Arm Identification: (i) lil’UCB Algorithm

Suppose we have $K > 2$ sub-Gaussian arms with a common and known parameter σ^2 . In many applications, we may want to identify which of the K arms has the largest mean parameter by using as few samples as possible.

In the previous subsection, we observed that, in the two-armed bandit setting, if we use a boundary-crossing compatible with a best-arm identification style algorithm with deterministic sampling, then the optimal arm has positive bias and the sub-optimal arm has a negative bias.

Shin et al. (2019a) showed that the same phenomenon happens for the best arm as chosen by the lil’UCB algorithm, which consists of the following sampling, stopping and choosing rules (Jamieson et al., 2014):

- **Sampling:** For $t = 1, \dots, K$, draw a sample from each arm. For $t > K$, draw a sample from arm k if

$$k = \arg \max_{j \in [K]} \hat{\mu}_j(t-1) + u^{\text{lil}}(N_j(t-1)),$$

where $\beta, \epsilon, \lambda > 0$ and $\delta > 0$ are algorithm parameters and

$$\begin{aligned} u^{\text{lil}}(n) := & (1 + \beta)(1 + \sqrt{\epsilon}) \\ & \times \sqrt{2\sigma^2(1 + \epsilon) \log(\log((1 + \epsilon)n)/\delta) / n}. \end{aligned}$$

- **Stopping:** the stopping time \mathcal{T} is defined as the first time at which the following inequality holds for some $k \in [K]$:

$$N_k(t) \geq 1 + \lambda \sum_{j \neq k} N_j(t),$$

where $\lambda > 0$ is an algorithm parameter.

- **Choosing:** $\kappa = \arg \max_{k \in [K]} N_k(\mathcal{T})$.

Proposition 4 (Corollary 8 in Shin et al. (2019a)). *In the settings of the lil’UCB algorithm, for each k with $\mathbb{P}(\kappa = k) > 0$, we have that*

$$\mathbb{E} [\hat{\mu}_k(\mathcal{T}) \mid \kappa = k] \geq \mu_k. \quad (21)$$




Figure 4. Average of conditional empirical CDFs of arm 1 from 10^5 lil'UCB algorithm runs on three unit-variance normal arms with $\mu_1 = 1$, $\mu_2 = 0.5$ and $\mu_3 = 0$, as described in Section 4.3. The black solid line refers to the true CDF of arm 1.

The following corollary to Theorem 2 shows a complementary result—the sample mean and the empirical CDF of a given arm are negatively and positively biased, respectively, conditioned on the event that the arm is not selected as the best arm. In contrast, on the same conditioning event, the empirical distribution of the selected arm is negatively biased.

Corollary 5. *In the settings of the lil'UCB algorithm, for each k with $\mathbb{P}(\kappa \neq k) > 0$, we have that*

$$\mathbb{E}[\hat{\mu}_k(\mathcal{T}) \mid \kappa \neq k] \leq \mu_k, \quad (22)$$

$$\inf_{y \in \mathbb{R}} \left(\mathbb{E} \left[\hat{F}_{k,\mathcal{T}}(y) \mid \kappa \neq k \right] - F_k(y) \right) \geq 0. \quad (23)$$

Also, for each k with $\mathbb{P}(\kappa = k) > 0$,

$$\mathbb{E}[\hat{\mu}_k(\mathcal{T}) \mid \kappa = k] \geq \mu_k, \quad (24)$$

$$\sup_{y \in \mathbb{R}} \left(\mathbb{E} \left[\hat{F}_{k,\mathcal{T}}(y) \mid \kappa = k \right] - F_k(y) \right) \leq 0. \quad (25)$$

The proof of Corollary 5 can be found in Appendix A.2.

Remark 6. *Corollary 5 remains true for any other choice of u^{lil} if the function $n \mapsto u^{\text{lil}}(n)$ is decreasing.*

Experiment: To verify the previous claims, we conducted 10^5 trials of the lil'UCB algorithm on three unit-variance normal arms with $\mu_1 = 1$, $\mu_2 = 0.5$ and $\mu_3 = 0$. It is important to note that the signs of the biases do not depend on the choice of parameters or of the underlying distributions, but the magnitudes of the biases do. To best illustrates the bias results, we use an unusual set of algorithm parameters as $\delta = 0.2$, $\epsilon = 0.1$, $\beta = 0.5$ and $\lambda = 1$ in this experiment.




Figure 5. Average of conditional empirical CDFs of arm 2 from 10^5 lil'UCB algorithm runs on three unit-variance normal arms with $\mu_1 = 1$, $\mu_2 = 0.5$ and $\mu_3 = 0$, as described in Section 4.3. The black solid line refers to the true CDF of arm 2.

Figure 4 shows the averages of the empirical CDFs of arm 1 (the arm with the largest mean) conditioned on each arm being chosen as the best arm. The black solid line corresponds to the true underlying CDF. The red dashed line, which lies below the true CDF, indicates that the empirical CDF of arm 1 conditioned on the event that the arm 1 is chosen as the best arm (i.e., $\kappa = 1$) is negatively biased; this then implies that the sample mean of the chosen arm is positively biased. In contrast, the green dotted and blue dot-dashed lines, lying above the true CDF, show that conditioned on the event the arm 1 is not chosen as the best arm (i.e., $\kappa \neq 1$), the empirical CDF is positively biased and the sample mean is negatively biased.

Figure 5 displays the averages of the empirical CDFs of arm 2. Though arm 2 is not the best arm, we can check that the signs of the conditional biases follow the same pattern as arm 1. Conditioned on the event that arm 2 is chosen as the best arm ($\kappa = 2$), the empirical CDF is negatively biased (green dotted line), but conditioned on the event that arm 2 is not chosen as the best arm ($\kappa \neq 2$), the corresponding CDFs are now positively biased (red dashed and blue dot-dashed lines), as expected.

4.4. Best-Arm Identification: (ii) Sequential Halving

The action elimination (Paulson et al., 1964; Even-Dar et al., 2002; 2006; Karnin et al., 2013) is another class of best arm identification algorithms in which each arm in an active set is sampled a fixed number of times and some of arms are eliminated from the active set at each round. In this subsection, we focus on the sequential halving algorithm

(Karnin et al., 2013), which is designed to identify the best arm with high probability and based on a fixed number of samples. To be specific, the sequential halving algorithm for a given time budget T is defined as follows :

- Initialize the set of active arms $\mathcal{A}_1 = [K]$.
- In each round $r = 1, \dots, \lceil \log_2 K \rceil$, for each arm k in \mathcal{A}_r , draw

$$m_r := \left\lfloor \frac{T}{|\mathcal{A}_r| \lceil \log_2 K \rceil} \right\rfloor$$
 i.i.d. samples from it, and let $\hat{\mu}_k^{(r)}$ be the average of the observed samples in the round.
- Let \mathcal{A}_{r+1} be the set of $\lceil |\mathcal{A}_r|/2 \rceil$ arms in \mathcal{A}_r with the largest average based on data collected in round r .
- Choose the unique arm in $\mathcal{A}_{\lceil \log_2 K \rceil + 1}$ as the best one.

Let κ be the (random) index of the chosen arm at the final time T , and for each $k \in [K]$, let $\hat{\mu}_k(T)$ be the sample mean of arm k at time T based on all previously collected $N_k(T)$ samples from that arm.

The following two corollaries show that under the sequential halving algorithm, the sample mean and empirical CDF of each arm is negatively and positively biased, respectively, marginally and conditionally on the event that the arm is not selected as the best one ($\kappa \neq k$). On the other hand, conditionally on the event that the arm is selected as the best one ($\kappa = k$), the sample mean and empirical CDF of each arm is positively and negatively biased, respectively.

Corollary 7. *Suppose that data are collected by the sequential halving algorithm with a time budget T . Then, for each $k \in [K]$ and for each $t \leq T$, we have that*

$$\mathbb{E} [\hat{\mu}_k(t)] \leq \mu_k, \quad (26)$$

$$\inf_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k,t}(y)] - F_k(y) \right) \geq 0. \quad (27)$$

Corollary 8. *In the settings of the sequential halving algorithm with a time budget T , for each k with $\mathbb{P}(\kappa \neq k) > 0$, we have that*

$$\mathbb{E} [\hat{\mu}_k(T) \mid \kappa \neq k] \leq \mu_k, \quad (28)$$

$$\inf_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k,T}(y) \mid \kappa \neq k] - F_k(y) \right) \geq 0. \quad (29)$$

Also, for each k with $\mathbb{P}(\kappa = k) > 0$,

$$\mathbb{E} [\hat{\mu}_k(T) \mid \kappa = k] \geq \mu_k, \quad (30)$$

$$\sup_{y \in \mathbb{R}} \left(\mathbb{E} [\hat{F}_{k,T}(y) \mid \kappa = k] - F_k(y) \right) \leq 0. \quad (31)$$

The proofs of Corollary 7 and 8 can be found in Appendix A.3. Due to limited space, an experiment about the bias of the sequential halving algorithm is presented in Appendix B.2.

5. Summary

In this paper, we have investigated the sign of the conditional bias of monotone functions of the rewards in the MAB framework under an arbitrary conditioning event, generalizing the results in Shin et al. (2019a), and complementing the results on the magnitude of the bias (Shin et al., 2019b). In our analysis, we have exploited certain natural monotonicity properties of MAB experiments and have characterized the impact on the bias of both adaptivity in the data acquisition process and in the selection of the target for inference.

Several interesting extensions of our results are worth pursuing. We emphasize two important ones:

- It is still an open problem how to characterize the bias (marginal or conditional) of other important functionals that are not necessarily monotone or linear, such as the sample variance and sample quantiles.
- Several debiasing methods have been proposed in the MAB literature: see, e.g., Xu et al. (2013); Deshpande et al. (2018); Neel & Roth (2018); Nie et al. (2018); Hadad et al. (2019). However, the existing approaches typically only adjust for adaptive sampling but ignore the other sources of adaptivity. **Current debiasing methods do not offer theoretical guarantees at stopping times, and also do not account for natural conditioning events in MABs, like selection of a promising arm.** Furthermore, they are typically not interested in characterizing the *sign* of the bias, and thus have complementary aims. It is of interest to investigate how our results and techniques can be used in order to design more general debiased estimators.

References

- Agrawal, S. and Goyal, N. Analysis of thompson sampling for the multi-armed bandit problem. In *Conference on Learning Theory*, pp. 39–1, 2012.
- Audibert, J.-Y. and Bubeck, S. Minimax policies for adversarial and stochastic bandits. In *COLT*, pp. 217–226, 2009.
- Auer, P., Cesa-Bianchi, N., and Fischer, P. Finite-time analysis of the multiarmed bandit problem. *Machine learning*, 47(2-3):235–256, 2002.
- Bowden, J. and Trippa, L. Unbiased estimation for response adaptive clinical trials. *Statistical methods in medical research*, 26(5):2376–2388, 2017.
- Deshpande, Y., Mackey, L., Syrgkanis, V., and Taddy, M. Accurate inference for adaptive linear models. In *Proceedings of the 35th International Conference on Machine Learning*, 2018.

- 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.
- Even-Dar, E., Mannor, S., and Mansour, Y. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of machine learning research*, 7(Jun):1079–1105, 2006.
- Garivier, A. and Cappé, O. The KL-UCB algorithm for bounded stochastic bandits and beyond. In *Proceedings of the 24th Annual Conference On Learning Theory*, pp. 359–376, 2011.
- Gut, A. *Stopped random walks*. Springer, 2009.
- Hadad, V., Hirshberg, D. A., Zhan, R., Wager, S., and Athey, S. Confidence intervals for policy evaluation in adaptive experiments. *arXiv preprint arXiv:1911.02768*, 2019.
- Jamieson, K., Malloy, M., Nowak, R., and Bubeck, S. lil’ UCB: An Optimal Exploration Algorithm for Multi-Armed Bandits. In *Proceedings of The 27th Conference on Learning Theory*, volume 35 of *Proceedings of Machine Learning Research*, pp. 423–439, 2014.
- Kalyanakrishnan, S., Tewari, A., Auer, P., and Stone, P. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pp. 655–662, 2012.
- Karnin, Z., Koren, T., and Somekh, O. Almost optimal exploration in multi-armed bandits. In *International Conference on Machine Learning*, pp. 1238–1246, 2013.
- Kaufmann, E., Korda, N., and Munos, R. Thompson sampling: An asymptotically optimal finite-time analysis. In *International Conference on Algorithmic Learning Theory*, pp. 199–213. Springer, 2012.
- Lattimore, T. and Szepesvári, C. *Bandit algorithms*. Cambridge University Press, 2019.
- Neel, S. and Roth, A. Mitigating bias in adaptive data gathering via differential privacy. In *International Conference on Machine Learning*, pp. 3717–3726, 2018.
- Nie, X., Tian, X., Taylor, J., and Zou, J. Why adaptively collected data have negative bias and how to correct for it. In *International Conference on Artificial Intelligence and Statistics*, pp. 1261–1269, 2018.
- Paulson, E. et al. A sequential procedure for selecting the population with the largest mean from k normal populations. *The Annals of Mathematical Statistics*, 35(1): 174–180, 1964.
- Robbins, H. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.
- Shin, J., Ramdas, A., and Rinaldo, A. Are sample means in multi-armed bandits positively or negatively biased? In *Advances in Neural Information Processing Systems*, pp. 7100–7109, 2019a.
- Shin, J., Ramdas, A., and Rinaldo, A. On the bias, risk and consistency of sample means in multi-armed bandits. *arXiv preprint arXiv:1902.00746*, 2019b.
- Starr, N. and Woodroffe, M. B. Remarks on a stopping time. *Proceedings of the National Academy of Sciences of the United States of America*, 61(4):1215, 1968.
- Sutton, R. S. and Barto, A. G. *Introduction to reinforcement learning*. MIT press Cambridge, 1998.
- Thompson, W. R. On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. *Biometrika*, 25(3/4):285–294, 1933.
- Villar, S. S., Bowden, J., and Wason, J. 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.
- Xu, M., Qin, T., and Liu, T.-Y. Estimation bias in multi-armed bandit algorithms for search advertising. In *Advances in Neural Information Processing Systems*, pp. 2400–2408, 2013.