

1 We thank the reviewers for their careful reading of the manuscript and for their comments.

2 **Response to reviewer #1** Thank you for your positive reception. We will polish the writing and address the minor  
3 comments (cross-references, readability of Fig. 2,3, etc). We are happy to consider any suggestion you may have for an  
4 altered title in the final review (that may arise in the discussion): ours was meant as a rebuttal to the Nie et al. title.

5 [...*potential applications where the sign of bias can be useful...*] The sign of the bias has been used to check whether a  
6 search advertising system charged advertisers fairly [Xu et al., 2013] or to investigate the validity of MAB models for  
7 the design of clinical trials [Villar et al., 2015]. We will add these examples to the paper. As the reviewer points out,  
8 quantifying the magnitude of bias is also an important complementary problem, which we will pursue in future work.

9 [...*the term "optimism" is well-known and used in bandits...*] Our choice of the term optimism is deliberate, and is  
10 directly inspired by the principle of “optimism in the face of uncertainty” that is frequently invoked in the MAB literature.  
11 Indeed, we view our definitions of optimistic sampling, stopping and choosing as different ways of formalizing that  
12 very same principle. (For example, UCB algorithms are optimistic samplers, and our definition captures this and other  
13 methods.) We will include an explicit reference to the principle of optimism and clarify our choice of terminology. We  
14 thought hard about a different term (and looked up synonyms), but could not find a more appropriate one, and are happy  
15 to listen to a suggestion in the final review.

16 [...*Is the choosing strategy always adaptive?...*] The choosing strategy can be non-adaptive. For instance, choosing a  
17 target arm by using an independent uniform distribution on  $\{1, \dots, K\}$ , is a random but non-adaptive choice.

## 18 **Response to reviewer #2**

19 [M1. *This paper could be of interest to people in the bandit community...you have to cite them more.*] We agree with the  
20 reviewer that it would be appropriate and beneficial to the manuscript to include more references to the MAB literature  
21 and we will do so, starting from the reviewer’s suggestions.

22 [M2. *...make the goal of this analysis more explicit.*] The main objective of this paper is to provide a rigorous and  
23 comprehensive analysis of the sign of the bias in MAB settings under adaptive sampling, stopping and choosing rules.  
24 This is, we believe, a problem of practical and theoretical significance that has not been fully solved. For instance,  
25 consider the offline analysis of data that was collected by an MAB algorithm (with any aim): suppose that you want to  
26 estimate the mean reward of some of the better arms that were picked more frequently by the algorithm. The paper by  
27 Nie et al. proves that the the sample mean will be negatively biased under fairly common adaptive sampling rules: this  
28 could instill a possibly false sense of comfort with your sample mean estimate since their theory suggests that you are  
29 underestimating the effect size. However, we prove that if the algorithm was adaptively stopped and the arm index was  
30 adaptively picked, then the net bias can actually be positive. Indeed, we prove that this is the case for lil’UCB, but it is  
31 likely true more generally as captured by our main theorem. Thus, the sample mean may actually overestimate the  
32 effect size. This is an important and general phenomenon for both theoreticians (to study further and quantify) and for  
33 practitioners (to pay heed to) because if a particular arm is later deployed in practice, it may yield a lower reward than  
34 was possibly expected from the offline analysis. We can add such an explanation to the paper if the reviewer found it  
35 useful to compare messages explicitly.

36 [M3. *...relationship between the performance of a bandit algorithm and the bias?...fair comparison between UCB and*  
37 *TS in numerical experiments...*] This is a great question/suggestion, which we can study in future work. The main goal  
38 in our simulations is to visualize and corroborate our theoretical results about the sign of the bias. As a result, we did  
39 not make any attempt to optimize the parameters for UCB or TS for the purpose of minimizing the regret, since the  
40 latter was not the paper’s aim. We will include a remark in the paper to clarify this.

## 41 **Response to reviewer #3**

42 Thank you for your detailed review which captures the most important features of our work concisely.

## 43 **References**

- 44 Sofia S Villar, Jack Bowden, and James Wason. Multi-armed bandit models for the optimal design of clinical trials:  
45 benefits and challenges. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 30(2):199,  
46 2015.
- 47 Min Xu, Tao Qin, and Tie-Yan Liu. Estimation bias in multi-armed bandit algorithms for search advertising. In  
48 *Advances in Neural Information Processing Systems*, pages 2400–2408, 2013.