We really appreciate the time and expertise you have invested in these reviews. We wish to express our appreciation for your in-depth comments, suggestions, and corrections, which will greatly improve the manuscript. We will reply to individual questions from reviewers respectively. Note that the numbers in the numbered lists below refer to sections of the review form as follows: 1=Contributions, 2=Detailed Comments, 5=Improvements.

Reviewer #1

1. We agree with your summary of our contributions and would like to emphasize again that we generalized the regret analysis of existing Gaussian Thompson Sampling [3] in two significant respects; (1) from the special Gaussian perturbation to general sub-Weibull or bounded perturbations, and (2) from the special Gaussian rewards distribution to general sub-Gaussian rewards. We would also like to emphasize that the lower bound in the stochastic case (Theorem 6) means that the regret analysis in important specific cases, like the Gaussian and double exponential perturbations, is tight.
2. We concur and this is an accurate summary for both settings in this work.
5. We agree and are working on the project of solving this open problem in the adversarial bandit setting. Our negative results around barriers to natural approaches to solving the open problem do have a positive aspect: they will save future researchers from spending time in these fruitless directions (as we did until we ran into these provable barriers).

Reviewer #2

1. We concur with your description of the main contributions of our paper for stochastic and adversarial bandit settings.
2. Thanks for your acknowledgement that the adversarial problem is really hard. It is this realization that makes even partial progress (in the form of barrier results) worth publishing. Thanks so much for detailed comments. We will fix them resulting in an improved manuscript.
5. Same answer as in point #5 for Reviewer #1 above.

Reviewer #3

1. We agree that this is an accurate summary of our contributions in both stochastic and adversarial bandit settings.
2. We agree. As you mention in significance part, our work paves the way for the design and analysis of efficient perturbation algorithms that enjoys both computational advantages and low regret guarantees in more complex settings such as stochastic linear bandits, combinatorial bandits and partial monitoring games.
5. In the paragraph “Failure of Bounded Perturbation” (L134-L141), we provided a counterexample that a perturbation algorithm via Uniform distribution but without log term will achieve a linear regret in two armed bandit problem. As an arm is pulled several times, the width of perturbation gets smaller because of scaling term $(1/\sqrt{T_i(t)})$, and thus the ranges of sampling distributions from two arms stop being overlapping so that the algorithm stops exploring and incurs linear regret. Therefore, we clearly need to add a term which is an *increasing* function of “global” time T so that it can compensate for narrow sampling range and restores good regret properties of the algorithm by increasing the width of perturbation. It is less easy to motivate in non-technical terms why the the term has to be *logarithmic* in time. Perhaps an analogy will help: the logarithmic term in the numerator also appears in optimistic algorithms like UCB. So it is satisfying that the randomized version randomizes within an interval of the same scaling over which UCB optimizes.