We thank the reviewers for the detailed comments and suggestions. Please find our responses below.

Response to Reviewer 1:

`Performance on real datasets:`
We are in conversation with groups having private access to wildlife poaching data where patrol scheduling to combat opportunistic crime is of major practical interest. We hope to include an extensive study on several real-world datasets as part of a journal version of the current submission.

`Number of thresholds smaller than the number of arms case:`
It is an interesting direction and will serve as a non-trivial extension of the current setting. Of course, one can naively solve it using the different-thresholds case of this paper. But we believe that a smarter, specialized algorithm whose performance will smoothly depend on how many and how *separated* the thresholds are can potentially be developed.

Response to Reviewer 2:

`Significance:`
The problem setup we considered has plenty of applications in domains such as police patrolling, poaching control, medical diagnosis, advertisement budget allocation, among many others. In this paper, we have proposed a novel framework for resource allocation (Censored Semi-Bandits (CSB)), which directly addresses such practical use cases. From the most natural way of formulating this problem, it is not at all apparent apriori that it reduces to Multi-Play or Combinatorial Semi-Bandits setup. Only a deeper understanding of the problem makes this connection explicit, which we feel is non-trivial. Furthermore, we believe showing such a reduction will help future work in this area. It would interest other researchers to look into richer models in resource allocation with censored feedback.

`Including` K `in identifying an instance of CSB:`
Not necessary. As $K = |\mu|$, K is known (implicitly) from an instance of CSB. We will make it clear in the final version.

`Arms with identical rates:`
We only require $\mu_1 \leq \mu_2 \leq \ldots \leq \mu_{K-M-1} < \mu_{K-M} \leq \ldots \leq \mu_K$, i.e., μ_{K-M-1} and μ_{K-M} are distinct, so that the KL-divergence in Theorem 1 is well defined (as required in [22]). This assumption is equivalent to saying that the set of arms under optimal allocation is unique. Note that CSB-DT does not need such an assumption. We will update this.

`Definition 1:` Thanks for catching this typo. '$\arg\min$' operator should be replaced by '$\min$.' We will update this.

`Assumptions on` N `and` M `in Lemma 1:`
We agree that a lower bound on N is needed to avoid the case $M = 0$. The assumption $N \geq \theta_c$ will ensure this (weaker than $N \geq 1$). We will state this. However, it is not necessary to define $\hat{\theta}_c$ as $\min\{N/M, 1\}$ to avoid it exceed 1. $\hat{\theta}_c$ can be allocation equivalent to θ_c even if $\hat{\theta}_c > 1$ and it does not disturb our analysis.

`Assumption on` $\mu_1 \geq \epsilon > 0$`:`
The assumption that $\mu_1 \geq \epsilon > 0$ is obviously restrictive. But it holds naturally when only arms with non zero mean loss are considered for resource allocation. In such setup, the minimum mean loss is at least ϵ for some $\epsilon > 0$. We will state it in the revision. It is still interesting to remove this assumption and will take it as future work.

Response to Reviewer 3:

`Asymptotic lower bound should be explicitly noted:`
Agreed (Theorem 3.1 of [21]). We will make it explicit.

`Asymptotic optimality of the algorithm:`
The asymptotic optimality indeed holds when we set $\delta = 1/T$. We indeed have $W = O(\log T)$, but that does not invalidate the optimality as can be checked by substitution.

`Joint estimation of threshold and loss:`
A natural algorithm for joint estimation is as follows: One first starts with a threshold vector, observes losses, and updates the mean loss for each arm. In the following round, the threshold is updated based on the observations in the previous rounds. The analysis of such an EM type algorithm for simultaneous estimation seems far more involved, but still doable. We will take it as an extension of this work.

`The tolerance parameter` γ `as input:`
We needed to know a lower bound on γ so that we can estimate $\boldsymbol{\theta}$ within some approximation. If we do not know this, we can use the joint estimation of threshold and loss as explained in the previous response. It will lead to an estimate of $\boldsymbol{\theta}$ that will eventually fall within the desired range of approximation provided $\gamma > 0$. However, the analysis of this is delicate, and we aim to take it as an extension of this work. We believe that the condition $\gamma > 0$ is necessary. Otherwise, sub-linear regret may not be achievable.