We thank the reviewers for the insightful and constructive comments. In what follows, we provide our response to the major concerns raised.

Reviewer 1: We agree with the presentation and reference issues raised, and will revise the paper as advised.

Reviewer 2, Comments 1 & 2: We agree that $\log^*(n)$ is a small number and further reducing it is not interesting. However, we note that our focus is not to reduce the round number below $\log^*(n)$, but to achieve $O(\log^*(n))$ rounds *without relying on the unrealistic assumptions made in [4].* In particular, [4] assumes that Δ_k (i.e., the difference between the means of the k^{th} and $(k+1)^{th}$ largest arms) is known, which is seldom the case in practice given that the means of all arms are unknown in advance. In contrast, our algorithms do not require any prior knowledge of Δ_k; we allow users to choose an error parameter $\epsilon \in (0, 1)$ to strike a trade-off between accuracy and efficiency. In our submission, we discuss the above issues in Lines 80-88 and 103-106. We also note that our $O(\log^*_{\frac{k}{\delta}}(n))$ result does not contradict the $\Omega(\log^*(n))$ lower bound in [4], since the latter regard k and δ as constants, whereas we consider them to be variables. If we also consider k and δ to be constants, then our result would match the lower bound in [4].

Reviewer 2, Comment 3: We will explicitly define $\log^*_b(n)$ as advised.

Reviewer 2, Comment 4: We have considered the suggested approach (which tests $\Delta = 1, \gamma, \gamma^2, \cdots$), but found that its complexity is inferior to ours, as explained in the following. Assume that the approach stops testing when $\Delta = \gamma^{2t}$. For the PAC setting (see Problem 1 in our submission), since $\gamma^{2t} < \Delta_k$, the suggest approach has to call the algorithm in [4] $O(\log \Delta_k^{-1})$ times, each of which requires $\log^*(n)$ rounds. Therefore, its round complexity is $O(\log \Delta_k^{-1} \cdot \log^* n)$, which is inferior to our $O(\log^*_{\frac{k}{\delta}}(n))$ result. For the exact top-k setting (see Problem 3 in our submission), let us consider a bandit instance where we have (i) k arms with means θ, (ii) $n-k$ arms with means $\theta - \Delta_k$, and $\gamma^t = 2\Delta_k$ (i.e., $\gamma^{2t} = 4\Delta_k^2$). If the suggested approach stops testing at $\Delta = \gamma^{2t}$, its query complexity is at least $O\left(\frac{n}{\Delta_k^4} \log \frac{k}{\delta}\right)$, which is inferior to our query complexity $O\left(\frac{n}{\Delta_k^2} \log \frac{k \cdot \log \Delta_k^{-1}}{\delta}\right)$.

Reviewer 3, Comment 1: For the proof of Lemma 1, we note that $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ holds with high probability even at the very beginning. In particular, in the first iteration (i.e., $r = 1$), S_r is exactly the same as the input arm set S (i.e., $S_1 = S$), and Algorithm 1 samples at least $\frac{32}{\epsilon^2} \log \frac{k}{\delta_1}$ times for every arm in S_1. Therefore, every arm i^* is initialized for $\hat{\theta}_{i^*}$. Based on Hoeffding bound, we have $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ with probability at least $1 - \frac{\delta_1}{k}$, for the case of $r = 1$.

The case for $r > 1$ follow from an induction on r, as shown in Lines 362-377 in our supplementary material. Specifically, suppose that $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ holds in the $(r-1)$-th iteration. If $\hat{\theta}_{i^*}$ is NOT updated in the r-th iteration, then $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ remains. Meanwhile, if $\hat{\theta}_{i^*}$ is updated in the r-th iteration, then we can apply the Hoeffding bound to show that after the update, $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ holds with at least $1 - \frac{\delta_r}{k}$ probability. By the union bound, the probability that $\hat{\theta}_{i^*} \geq \theta_i^* - \epsilon/8$ holds in all iterations is at least $1 - \frac{\delta}{2k}$ (see Eq. (7) in our supplementary material).

We will revise the proof of Lemma 1 to avoid confusions over the cases of (i) $r = 1$ and (ii) $r > 1$ and $\hat{\theta}_{i^*}$ is not updated in the r-th iteration.

Reviewer 3, Comment 2: Regarding the comparison between δE and other fixed confidence BAI algorithms: we have actually done such a comparison in our submission (see Figures 1 and 2). In particular, we compare δE and k-δE against ME [5] and ME-AS [6], both of which are state-of-the-art methods for fixed confidence instance-independent BAI. Our results demonstrate that δE (resp. k-δE) significantly outperforms ME (resp. ME-AS) in terms of query cost. In addition, in Section 4.2, we use δE as a subroutine to construct an algorithm (referred to as EG-δE) for exact (instead of PAC) top-k arm identification, and we show that it outperforms the state-of-the-art elimination-based method [8] and UCB-based method [17].

Meanwhile, for δER, we find it difficult to compare it with fixed budget BAI algorithms, due to the significant difference in the number of rounds that they require. Specifically, existing fixed budget BAI algorithms (e.g., [8]) require at least $\log(n)$ rounds, whereas δER requires *at most* $\log^*(n)$ rounds. This makes it difficult to identify a setting of round numbers to conduct a fair comparison of the algorithms.

Reviewer 3, Comments 3 & 4: We will clarify the motivation for multiple testing and detail the Exponential-Gap-Elimination algorithm as advised.