We would like to thank the reviewers for their thoughtful comments. We address below the main questions.

Reviewer #1

1. How realistic is the oracle of deactivating experts and how interesting is this part to the community?
   Our intent in introducing the oracle is just to provide a formal way to generalize from "assuming that all experts have low variance" to "competing with the best low-variance expert" or "competing with the best expert in the period in which it has low variance". We will try to be more clear about that.
2. Why using the sleeping expert definition in the paper instead of the one in Kleinberg et al.?
   Kleinberg et al. study the regret to the best ordering of experts, which is indeed different from our regret definition. Kanade and Steinke [2014] show that achieving optimal regret bound in this setting is computationally hard. Computationally efficient no-regret algorithms (e.g., Blum and Mansour [2007]) are known in our setting. Tackling the computation hardness of the sleeping setting is not the focus of this paper. We will add the reference and more discussion on the sleeping settings.
3. The meaning of "at the beginning of each epoch, we apply L over the average primary loss of each epoch"?
   The algorithm in Section 4.2 divides T into $T^{1-\alpha}$ epochs evenly. Let $e_i = \{(i-1)T^\alpha + 1, \ldots, iT^\alpha\}$ denote the i-th epoch and $\ell^{(1)}_{e_i,h} = \sum_{t\in e_i} \ell^{(1)}_{t,h}/T^\alpha$ denote the average primary loss of the i-th epoch. We run $\mathcal{L}$ over $\{\ell^{(1)}_{e_i,h}\}_{h\in\mathcal{H}}$ for $i = 1, \ldots, T^{1-\alpha}$. We will make sure to specify the algorithm.

Reviewer #2

1. The parameter α is required as input, whose value can only be obtained after the learning process?
   It is true that our algorithms need some information of α to obtain meaningful regret bounds. However, our algorithms only uses α to determine the length of epochs. Our algorithms can run without α with a more complicated theoretical guarantee. For example, we run Algorithm 1 with $\mathcal{A}_{\mathrm{SL}}(\mathcal{L})$ running $\mathcal{L}$ over $T^{1-\beta}$ epochs instead of $T^{1-\alpha}$ epochs, where β is given as input. Then Algorithm 1 can achieve $\mathrm{SleepReg}^{(1)}(h^*) = O(\sqrt{T_{h^*}T^\beta})$ and $\mathrm{Reg}^{(2)}_c \le \delta T^\alpha(\sqrt{\log(K)T^{1-\beta}} + K) = O(\sqrt{T^{1+2\alpha-\beta}})$. To make the bounds meaningful, we need to set $2\alpha - 1 < \beta \le \alpha$ (assuming $T_{h^*} = \omega(T^\alpha)$ as mentioned in Section 5.1). If $\alpha \le 1/2$, we can set β to any value in $[0, \alpha]$. Therefore, we do not need the exact value of α as input.
2. Missing reference?
   We will add the reference.

Reviewer #3

1. The regret $o(T)$ is large?
   We agree that $o(T)$ is inevitable when α is close to 1 as shown in the lower bounds (Theorem 2 & 3) in the "good" scenario. In the "bad" scenario, when T_{h^*} is close to T^α, the sleeping regret $o(T_{h^*})$ is inevitable as we mention in the paper.

References

Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8(Jun): 1307–1324, 2007.

Varun Kanade and Thomas Steinke. Learning hurdles for sleeping experts. *ACM Transactions on Computation Theory (TOCT)*, 6(3):1–16, 2014.