We would like to thank all the reviewers for their careful reading and constructive suggestions.

To address a common question of Reviewer #3 and #5: as currently stated in the paper, Theorem 1 applies only to the case when we have an upper bound on the target quality of distribution, and such a distribution exists. In the latest revision of the paper, we additionally describe an extension of this result which does not require this knowledge. More precisely, the algorithm we give is "semi-agnostic" – if OPT is the minimum total variation distance between the target distribution and any of the hypotheses, it outputs a hypothesis which has distance $O(\mathrm{OPT}) + \alpha$ from the target – at the cost of additive terms which are generally lower order (i.e., the sample complexity n increases by an additive $O(\log\log(1/\alpha)(1/\alpha^2 + \log^2(1/\alpha)/(\alpha\epsilon)))$. We will describe this result in the final version, together with a complete proof in the supplementary material.

We will additionally be sure to address all of the detailed feedback from Reviewers 3 and 5, which we believe will significantly improve the presentation of the paper. Regarding three specific comments:

- Reviewer #3's question about probability distributions applied to a set: the notation $P(W_1)$ stands for the probability mass which distribution P applies to the set W_1. This will be clarified in the next version.
- Reviewer #5's question about how error in p_1 and p_2 propagate in the analysis: It is sufficient to estimate these quantities up to an additive $O(\alpha)$, which can be done with $O(1/\alpha^2)$ samples to each of the candidate hypotheses. This will be rigorously argued in the next version.
- A near-linear time algorithm: While we made some attempts to derive a near-linear time algorithm, they were roadblocked by a number of known lower bounds for canonical problems (e.g., lower bounds for top-k selection). As a result, we consider this to be an interesting avenue for further study – we weakly conjecture that there may be a deeper phenomenon at play, potentially including a trade-off between time and sample complexity.