Thank you to all of the reviewers for their insightful comments! We respond to specific questions and comments of each reviewer below, and further provide additional discussion of the problem of developing algorithms that guarantee LSS.

Reviewer 1: We completely agree that the notation and style should be streamlined to improve readability, and have undertaken changes to address this. One such change will be removing subscripts and superscripts in cases where the distribution/generator/etc. are clear from context.

Reviewer 2: Indeed, the way Theorems 4.6. and 4.7 were stated may have made them seem weaker then they are actually are, and we thank you for pointing this out. The theorems were stated for the "interesting" values of small ϵ and δ, but also immediately hold for larger values, so long as the sample size is sufficient. Theorem 4.6 holds for *any* $\epsilon > 0$, for sufficiently large n, and Theorem 4.7 similarly holds for *any* $0 < \delta < 1$, for sufficiently large n. We have revised these two Theorem statements to reflect this.

Reviewers 2 and 3: We agree with you that developing algorithms that satisfy LSS and techniques for bounding LSS is an exciting direction for future work, and we would be glad to expand our discussion of this in the paper.

Naturally, any mechanism which guarantees Differential privacy (e.g., the Laplace and Gaussian mechanisms) will guarantee LSS as well, as a result of the DP->LMI->LSS implications. We plan to point this out more explicitly.

One can also see this, and perhaps gain additional insight, by manipulating the loss definition:

$$\sum_{x \in X_+(r)} \left(D^G_{\mathcal{X}|\mathcal{R}}(x \mid r) - D_{\mathcal{X}}(x) \right) = \sum_{x \in X_+(r)} D_{\mathcal{X}}(x) \left(\frac{D^G_{\mathcal{X}|\mathcal{R}}(x \mid r)}{D_{\mathcal{X}}(x)} - 1 \right) = \sum_{x \in X_+(r)} D_{\mathcal{X}}(x) \left(\frac{D^G_{\mathcal{R}|\mathcal{X}}(r \mid x)}{D^G_{\mathcal{R}}(r)} - 1 \right)$$

Since $D^G_{\mathcal{R}}(r) = \sum_{x' \in \mathcal{X}} D^G_{\mathcal{R}|\mathcal{X}}(r \mid x')$, it suffices to bound the quantity $\frac{D^G_{\mathcal{R}|\mathcal{X}}(r \mid x)}{D^G_{\mathcal{R}|\mathcal{X}}(r \mid x')}$ for any $x, x' \in \mathcal{X}$, which is bounded by e^{ϵ} for a Laplace mechanism with parameter $\frac{\Delta}{n\epsilon}$ in the case of a product distribution. Though this example does not provide direct improvement over DP, it may suggest a potential technique for proving LSS bounds for novel mechanisms.

Reviewer 3: There are indeed other candidates for the distance notion in Definition 2.1. We have explored some of them, but have not found another notion that we can show is both necessary and sufficient for generalization. Perhaps the most natural alternative to consider is bounded KL-Divergence, which, by Jensen's inequality, implies a bound on TV-distance. Thus, it is natural that bounded KL-Divergence would be sufficient for generalization; however, it is not clear that it is necessary. The form of the "loss assessment query" we introduce provides some intuition for the choice of the TV-distance; one cannot construct a natural analogous query for KL-Divergence, due to its unboundedness. This observation does not demonstrate that other distance measures cannot be used, but at least suggests that our proof technique may not suit them.

The fact that we handle non-iid databases is actually crucial. The reason for this is that even if the underlying data distribution were iid, the resulting *posterior* distribution given a query response might no longer be iid. Thanks for pointing out that we need to clarify this in the writeup.

The α_i values presented in Theorem 2.7 are expected losses, which might be significantly lower than ϵ_i (which can be thought of as high probability bounds on the loss). As you suggest, we will clarify the comment right after Theorem 2.7, that the Theorem is weakest when α_i is close to ϵ_i, and more meaningful when $\alpha_i \ll \epsilon_i$.