We thank the reviewers for their positive feedback, and address their main concerns below. Given the opportunity, we will address other concerns in the final version. We are grateful to Reviewers #1 and #3 for seeing the potential of our work to spark future research in the intersection of differential privacy, TEEs and oblivious algorithms.

Reviewer #1: ***Q:*** *What is the overhead for setting up the secure environment, the encryption/decryption step. How does it compare to that of the LDP+shuffle [55,56] and ESA of [8]* **A:** Setting up an enclave is a one time cost and is proportional to the size of the code and data, giving a linear overhead. The overhead of encryption/decryption is also linear. It is hard to compare overhead of our framework to that of [55,56] as implementation of shuffle is largely left unexplained in [55, 56]. One natural way to implement shuffle step in [55] is indeed to use TEEs, then the overhead of these frameworks should be comparable. If implemented using mixnets [56], the overhead might even be higher. However, this is a very good point, and we will add this in the final version.

Q: *How trusting an anonymization primitive is different to trusting a secure environment?* **A:** Anonymization using mixnets [56] will require assumption on non-collusion between the servers. If anonymization is implemented via TEEs, then the trust model would be the same as ours. While LDP+shuffle idea in [55,56] is mathematically elegant, DP algorithms inside TEEs come with two major technical advantages: 1) One can use DP algorithms in the central model. Consider for example using private multiplicative weights algorithm for answering exponentially many linear or counting queries. This has no parallels in the LDP+shuffle model. Even with amplification result one can only answer polynomially many queries if we want to achieve same level of accuracy as the central model. 2) Given the growing software support for TEEs by Google, IBM, Microsoft, DP+TEEs approach, arguably, seems closer to adoption in practice.

Q: *Do the authors' conjecture that the weakening of the requirements of the algorithms will result in faster/more accurate algorithms?* **A:** We agree with your intuition. We also believe ODP definition should allow us to design faster algorithms for DP problems than simply combining with the stronger notion of oblivious algorithms. We plan to explore this direction in the future. The accuracies achieved by our algorithms for all the 3 problems considered match the trusted curator model asymptotically.

Reviewer #2: Q: *Composition of TTP and memory oblivious algorithms has been considered in many previous works* **A:** As we mention in our paper, we agree that composition of TEEs (or TTP) with oblivious algorithms, in general, is not new. If this was not clear, we will make sure that this is stated more clearly in the final version. First, no paper earlier to our work has suggested *running central DP* algorithms within a TEE (*). This idea leads to combining oblivious algorithms and memory-restricted algorithms for the design of differentially private algorithms, and is a new contribution. Even considering all the concurrent/other works, our ODP definition is new: we consider an adversary that *gains access to the output of the computation*. We then ask the question of how to securely and efficiently instantiate the global DP setting using TEEs. We observe that since DP output is revealed, the access patterns do not have to be strictly oblivious, motivating our new definition. The opposite direction, applying DP to oblivious algorithms, has been explored in independent work [11]. [8] explores using oblivious shuffle for anonymization, and is technically different from our ODP definition. We cannot prove the independence of our work to (e.g., [8,11]) or (*) to preserve anonymity.

Q: *In Histogram protocol, at one place the authors claimed that oblivious shuffle is expensive and then use oblivious shuffle in step 4.* **A:** We mention that *oblivious sort* is expensive so we use the *shuffle* in step 4.

Q:*The constructions of this paper seem to follow mostly from previous work.* **A:** While some of the individual components of our algorithms have appeared before (and we cite), combining ideas from oblivious algorithms literature to DP, and design of DP algorithms with limited memory are both new, and have not appeared before (as also noted by Reviewer #1). Further, some of the previous works is parallel. ***Q:*** *I did not also see a convincing analysis for the efficiency of the proposed constructions in a TTP architecture.* **A:** We have provided full proofs of theorems in the supplementary material. Due to page limit, we could not give full proofs in the main body or discuss all the improvements our theorems imply. For example, from Theorem 4.4 our histogram construction is more efficient than oblivious constructions for larger values of k (see also Table 1 in the Appendix).

Q: *Better analysis of the costs of the protocols and the gains compared with a trusted aggregator model, or the shuffle and compute model.* **A:** Our algorithms achieve same level of accuracy guarantees as that of trusted aggregator model, as can be seen from our theorems. The cost of our framework lies in the increased running time not accuracy. Compared to LDP+shuffle model, from theorems in [55,56], our accuracy guarantees are better.

Q: *Related Work suggestions* **A:** Thank you for related work suggestions. We will cite these papers appropriately and update Table 1. From a quick reading, it appears that privacy blanket paper still does not help us achieve accuracy guarantees of the trusted aggregator model, as it will require larger values of epsilon.

Reviewer #3: Q: *Address other side channel attacks a little more thoroughly and explain how to design algorithms.* **A:** Our Definition 3.1 extends to other side channels, and we will make it more clear in the final version. However, design of DP algorithms incorporating *all* side channels is challenging and is an interesting research direction on its own (even when DP is not considered). We will discuss our ideas for preventing timing attacks in the camera-ready version if given an opportunity. We will also add more details about the heavy hitters in the final version.