We thank the reviewers for their careful examination of our work. All three reviewers had helpful suggestions that we wish to address.

We absolutely agree with R1 that a robust discussion about the differences between our results and Liu et al. ’18 is missing from the current draft. While a detailed comparison to Liu et al. was present in an earlier draft, it was removed in shortening the paper for the NeurIPS submission, and we acknowledge that this was an error. Here is a summary of this comparison that will be included in the final draft: Liu et al. show how to learn mechanisms with incentive guarantees, and their results do rely on differential privacy, but our results differ in two significant ways:

1. Our setting takes a Bayesian model of buyer types, where bidder values are drawn from a fixed distribution. Liu et al considered a model where buyers are adversarially selected.
2. Liu et al compare their mechanisms to the optimal single price (or optimal anonymous reserve price, with multiple bidders). In other words, their benchmark comes from a family governed by only a single parameter. The benchmark used in our results is much more challenging: we compare to the Bayes-optimal auction for the unknown distribution from which values are sampled. Generally, second-price auctions with anonymous reserve prices (as used in Liu et al) are suboptimal in our setting. The Bayesian-optimal mechanism we compare to is instead parametrized by a virtual value function for each bidder, a significant increase in complexity.

The Bayesian setting and Bayesian-optimal benchmark force us to adopt very different analysis techniques. Instead of the expert learning approach adopted by Liu et al., our approach is closer to the literature on sample complexity in single-shot mechanism design, and our main technical contribution is showing that introducing noise to the learning process for privacy can be done without significantly harming the revenue of the mechanisms we learn.

We thank R5 for pointing out other relevant work. We agree about the potential weakness of utility-approximate BIC (and that differential privacy doesn’t necessarily give something stronger), which is why we included results for bid-approximate BIC in Section 4. Xiao ‘11 and Nissim et al. ’12 rightly pointed this out, and we will gladly cite both in all future versions of our work. Similarly, Braverman et al. give results for nonmyopic buyers, and we will also include this reference in our revisions. Finally, the work relating privacy to online advertising, while sharing keywords with our work, generally is motivated by the idea of personalizing advertising to the viewer, rather than maintaining private estimates of the advertisers’ value distributions. We nonetheless appreciate the references.

Finally, we remark that our techniques do indeed apply to the alternative model suggested by R1 (as well as to several other natural modifications). We will gladly incorporate a discussion of other settings in which our results apply.