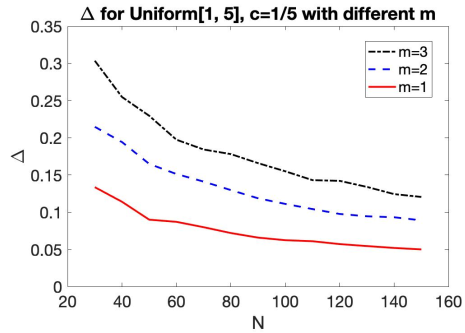

Figure 1: $\Delta^{\text{worst}}_{N,m}$

Thanks for your comments!

Review 1: Thanks for your appreciation!

Review 2:

Justification for approximate incentive compatibility. It is true that ϵ going to 0 in the limit doesn't mean the incentives are negligible in practice. But this is exactly why we study the convergence rate of ϵ.

Indeed, ϵ-BIC is fragile in a game with a few players. But in our model, there is large pool of bidders, so it is unrealistic for a bidder to find the exact best response: she has to collect a lot of information about other bidders' strategies, try many solutions, do a large amount of computation, etc. When the cost of searching for a better response exceeds ϵ, the bidder is better-off bidding truthfully. Even if the bidder suspects that some other bidders may bid non-truthfully, our incentive-awareness measure $\Delta^{\text{worst}}_{N,m}$ still guarantees that, as long as the total number of bids of all non-truthful bidders doesn't exceed m, no bidder can benefit a lot from lying. So we think ϵ-BIC is appropriate here. Also, we don't see the possibility of a (nontrivial) exact BIC mechanism when a bidder can affect her future reserve price; no BIC pricing function exists even in a two-round auction.

Revenue. Surely, the optimal revenue will be higher than Myerson if we relax BIC to ϵ-BIC. But to our knowledge, no work has studied how much higher that will be. We conjecture that it will be $O(\epsilon)$, and our $(1-\epsilon_2)$-revenue guarantee will hold against this stronger benchmark with an ϵ_2 that is larger than the current ϵ_2, but still goes to 0.

Transforming ϵ-BIC to exact BIC mechanisms. The current techniques cannot be applied here for two reasons: (1) they require some knowledge of the type distribution, either the full distribution or samples, which is unavailable in our model; (2) they are for one-shot games, but our two-phase model is a repeated game.

Asymptotic bounds. Fig. 1 shows $\Delta^{\text{worst}}_{N,m}$ for the uniform$[1,5]$ distribution with $m = 1, 2, 3$. $\Delta^{\text{worst}}_{N,1}$ goes below 0.05 already when $N = 150$, so a bidder who controls a few bids cannot reduce the price by much in a medium sized market. $\Delta^{\text{worst}}_{N,m}$ grows linearly in m roughly, which agrees with our theorem.

Review 3:

Range of m. In the most extreme case where $m = N$ (there's only one bidder), Amin et al (2013) show that it's impossible to learn a strategic bidder's value distribution to obtain near-Myerson revenue in repeated auctions unless the bidder discounts her future utility; similarly, in our model, when $m > cN$, $\Delta^{\text{worst}}_{N,m}$ always equals 1, which means the incentive property of ERM^c is very bad. So we are mainly interested in the case where m is a relatively small fraction of N, i.e., $o(\sqrt{N})$. We don't know what happens when $\Omega(\sqrt{N}) \le m \le cN$; the current framework cannot handle that. For the lower bound, thanks for suggesting considering the dependency on m! We were focused on N previously.

Rules other than ERM^c. There are indeed some algorithms that use differential privacy, e.g., Abernethy et al (NeurIPS'19). Abernethy et al's truthfulness and revenue bounds are better than ours. But their work is worse than ours in three other aspects: (1) their truthfulness notion is weaker than our PBIC notion; (2) ERM^c's running time is far less than their algorithm's; (3) ERM^c works for unbounded distributions while theirs only support bounded distributions because they need to discretize the value space; (details are in Sec. 2.4 and App. C.4.) They also need $m = o(\sqrt{N})$ to obtain approximate truthfulness and revenue optimality at the same time, as we do. We remark that our starting point is to understand ERM since it is the most fundamental learning algorithm, rather than to design good algorithms.

Exploration and exploitation at the same time. Under the BIC truthfulness assumption, it's OK to continuously update the reserve price in the second phase (rather than fixing it), which improves revenue without affecting truthfulness. But for the Perfect Bayesian Incentive-Compatibility notion (see Sec. 2.4 and App. C.4), we cannot do exploration and exploitation at the same time. Consider this example: Each auction has one bidder. In the first round bidder A submits bid 1 and 1 is set to be the reserve price in the second round. In the second round, bidder B submits 0.9 and loses, so she knows that the first round bid is greater than 0.9. Suppose bidder B will join the next round, then she can strategize her bid based on the information "the first round bid > 0.9". So the belief of bidder B on bidder A's value distribution is no longer F, and our argument fails. But doing exploration and exploitation separately is good (see App. C.4).

Review 4:

Prior-independence. In order to obtain good revenue in a second price auction, the number of bidders in the auction needs to be large since it is a $1 - 1/n$ approximation. But in our two-phase model, the number of bidders in each auction (denoted by K_1 and K_2 for the two phases, at 125) can be very small, e.g., 1 or 2. With so few bidders, prior-independent mechanisms do not have good revenue. We remark that in our model, although the total number of bidders is large, each auction may have only a few bidders because bidders do not participate in all rounds of auctions.