We thank the reviewers for their detailed feedback, which will improve the presentation of our paper.

First we would like to address a high level point that was raised by the reviewers regarding Riccati perturbations. While the proof of Proposition 1 follows the argument of Konstantinov et al., Proposition 2 relies on a new elementary proof technique that is of independent interest and offers a tighter guarantee for certain classes of problems, as discussed in Lines 268-275.

Reviewer 4.

- Regarding the rich literature on certainty equivalence for tabular MDPs, we did not discuss tabular MDPs because our focus was on LQR. Nevertheless, we agree with the reviewer that the tabular MDP literature is relevant and should be included; we will introduce a few paragraphs on related works (such as Azar et al. 2013) in our revision. We would appreciate suggestions of the most relevant works studying certainty equivalence for tabular MDPs.
- It is definitely possible to verify empirically that certainty equivalence outperforms robust control methods when the model estimation error is small, while being more sensitive to the size of the error. Dean et al. [2017] observed exactly this in Figure 2; we offer a theoretical justification for their observations.

Reviewer 5.

- At this stage, our results do not offer bounds that can be used numerically in practice. Indeed, as the reviewer suggests, they would be too conservative. However, our results offer insights about the performance of certainty equivalence for LQR and LQG. For example, it explains why Dean et al. [2017] observed empirically (c.f. Figure 2 of their paper) that certainty equivalence performs poorly in the high error regime, but outperforms their robust methods in the low error regime.
- Dean et al. [2017] do not analyze certainty equivalence, a popular method used in practice. Moreover, Dean et al. [2017] study only fully observed systems, i.e. they studied LQR, but not LQG. We showed that certainty equivalence achieves a fast statistical rate for both LQR and LQG. We would like to emphasize that the partially observed case is significantly more challenging than the fully observed case.
- Theorem 2 follows from plugging in the inequality from Proposition 2 into the bound of Theorem 1. We will make this more clear in our revision.
- We will make sure to define our notation before it is used in the revision; $\mathsf{dare}(A, B, Q, R)$ is the unique positive semidefinite solution to the discrete algebraic Riccati equation associated with the parameters A, B, Q, and R.
- The parameters (ℓ, ν) quantify how controllable a system is. They are system dependent quantities. For example, if the system is controllable in the classical sense, one can choose ℓ to be equal to the state dimension and ν to be the minimal singular value of the controllability matrix. This is not the only choice, however, and this additional degree of freedom allows us to offer sharper bounds for certain systems (c.f. Lines 268-275).

Reviewer 6.

- Lower bounds would indeed be very desirable, but unfortunately so far we have not been able to derive lower bounds. We leave this for future work.
- Cohen et al. [2019] focus on the online and fully observed setting and offer an elegant method based on semidefinite programming which achieves $\sqrt{T}$ regret, where T is the horizon. However, their method requires the initialization of the system parameters to be in an error ball of radius $\mathcal{O}(1/\sqrt[4]{T})$, while our method does not have this restriction. Our method also offers computational advantages over Cohen et al. [2019], since we can take advantage of specialized DARE solvers instead of relying on general SDP solvers. We will make this comparison to the work of Cohen et al. [2019] explicit.
- It would indeed be valuable to derive bounds which rely only on observed data and not on unknown problem dependent quantities. At this time we are not aware of a way to achieve this.

References

A. Cohen, T. Koren, and Y. Mansour. Learning Linear-Quadratic Regulators Efficiently with only $\sqrt{T}$ Regret. *arXiv:1902.06223*, 2019.

S. Dean, H. Mania, N. Matni, B. Recht, and S. Tu. On the Sample Complexity of the Linear Quadratic Regulator. *arXiv:1710.01688*, 2017.