We thank the reviewers for their valuable and constructive feedback. We will incorporate the suggestions in the next revision. Our responses below.

Reviewer 2:

- **Related works**: Both of these works seem to use SGD-DD. The second paper in particular considers Gaussian Vector Autoregressive (GVAR) model for the Markov chain and our proof techniques and results on experience replay could perhaps help in this setting as well.

- **Bounds for SGD in Table 1**: (a) Lower bounds – The information theoretic lower bounds for bias in the independent setting (Theorem 1) and variance in the agnostic setting (Theorem 2) also apply to SGD. (b) Upper bounds in AR chain – Since we want to demonstrate the benefits of SGD-ER over SGD, we focus only on proving that the upper-bounds for SGD-ER are better than the *lower bounds* for SGD.

- **Section 2.1, modeling assumptions**: The agnostic and independent noise settings are standard models to account for inexact observations in real data and in nature and have been studied in literature e.g., see [7] and references there in. The AR chain is a popular chain used for forecasting and time series prediction problems, and has nice structural properties that benefit from experience replay. We will include a detailed discussion regarding this in the manuscript.

- **Text in and around equation (3)**: This is a typo. w_1 should be w_0, which is the initial iterate.

- **Theorem 5**: We restricted our attention to constant step size SGD since this is used widely in practice for least squares regression and is optimal for i.i.d data. Analysis of SGD with decreasing step size would be a great direction for future work, but even for iid data, decreasing step size leads to sub-optimal rates for the bias term.

- **Figure 1**: Our theorems show that SGD-DD is no worse than SGD in the worst case, i.e, for certain Markov chains. For the AR chain, it seems that SGD performs better than SGD-DD due to the special structure of the AR process. There is also a $\log(d)$ difference between the lower bound for SGD (specialized to the AR chain) and upper bound for SGD-DD, as seen in Table 1.

Reviewer 3: We will include a detailed explanation of the noise model by explicitly constructing the underlying probability space. We will also include the suggested references in related work.

Reviewer 4:

- **Lower bound and optimality of SGD-DD**: We respectfully disagree with the reviewer's claim that these results are already known. We request the reviewer to please share related references that he/she thinks are relevant.

- **Gaussian AR chain is very specific and is of limited interest**: (a) As demonstrated by our lower bound results, experience replay cannot in general improve upon SGD-DD in the worst-case. Improvements can be obtained only for some well-structured Markov chains and we show that the AR chain is one such. (b) Gaussian AR chain is a popular chain used for forecasting and time series prediction.

Reviewer 5: Our experimental results (a) validate the tightness of our theoretical bounds and (b) show that SGD-ER outperforms SGD and SGD-DD even on small/moderate dimensional problems. Experience replay is widely used in practice for reinforcement learning with great results, so we expect good outcomes from SGD-ER in other regression settings as well.