We thank the reviewers for their useful feedback. Our responses are below. We mistakenly omitted the code from supplemental material at submission, but it will be public following the new NeurIPS rules for accepted papers (code is already on GitHub in fact) to enable reproduction and extensions of our work, and to allow others to try a greater suite of experiment settings. Unfortunately, links (even anonymous) are disallowed in the response.

Reviewer 1:

8) Thanks for catching. Will fix typo.
9) We will add more detail. The figure seeks to illustrate the local and intrinsic efficiency uniquely achieved by our new estimators. It shows the ordering of asymptotic MSEs: if Q-functions are well specified, both our proposed estimators (EMP, REG) and DR achieve the same efficiency bound, which IS and SNIS do not; if Q-functions are misspecified, the efficiency bound is not achieved, yet our proposed estimators will *still* have better MSE than DR, IS, and SNIS.
10) Thank you for the great reference. We will definitely cite it.

Reviewer 2:

1) We followed Farajtabar et al. [5] when constructing evaluation and behavior policies. Our policies are the ones that they call friendly softening in their paper. We tried experiments on some other behavior and evaluation policy cases, and the overall results were the same (when policies are close, DR performs well among baselines and we match/exceed it; when policies are far, SNIS performs well among baselines and we match/exceed it). We will include the additional results in the supplement. Additionally, the availability of the code allows one to easily play with the policy parameters.
2) Firstly, the difference is small but usually statistically significant; we will add standard errors. Second, Tables 2–4 show that the our estimators (EMP, REG) are competitive with DR in the settings where it works, and beat it handily when it does not work, such as when the behavior and evaluation policies differ as in the Tables' 3rd line. We will comment. Similarly, Tables 5–7 show a variety of settings. While DR is sometimes competitive and sometimes less so, our estimators perform well throughout. We will add a comment that in some settings many estimators perform similarly well (DR, MDR, and ours) while in others we outperform. The point is we are never worse and sometimes better. This verifies our theory.
3) We will add standard errors.

Clarity:

1) Good point. We will add "More Robust Doubly Robust" at the first mention of MDR.
2) Re L60, we just meant to highlight that local efficiency is limited because correct parametric specification is a strong assumption. We will clarify.

Significance: Even if the Q-model is a neural network, our proposed method still works and our guarantees will still hold. We agree that often using such Q-models is necessary in more complicated situations. However, for the experiments studied, which are standard in the literature we build on, all of the domains are not complex and we followed the previous literatures regarding the choice of Q-model to most clearly demonstrate the contribution of our approach relative to previous work.

Reviewer 3:

- Re Sec 3,4 clarity: we will follow your suggestion and add more intuition on the construction. Sec 3 follows similar ideas to [5], whereas Sec 4 uses a new modified (normalized) empirical likelihood approach, which we will explain and connect more clearly to standard empirical likelihood.
- Re code: apologies for the omission in supplement submission; the code will be public. See above.
- Plots of the tables will be added to the supplement. Obviously, there were space constraints. See also R2(2) about additional discussion of results.
- Re improvements:
  - See above re Sec 3,4. Additionally, we will do another editing pass to improve clarity.
  - There are many results and we wanted to clearly contextualize the contribution in the literature. We tried to add clarity by first discussing the CB case even though it is just a special case of the RL setting. We will add additional intuitive explanations in Sec 3.
  - Code will definitely be provided for reproducible research and to allow others to build on our work.