We thank all reviewers for their helpful comments and suggestions! We will address all minor issues. Please see our responses to major issues below.

R#1. Re. correlations among noise. We totally agree with the reviewer that correlated noise is an important topic for future research. Actually our results only require independent (not necessarily identical) noise, which is a classical assumption in many literatures as briefly discussed in L126–127. We will add references in the revision. In the context of social networks, our framework can model the setting where agent's preferences are mostly determined by his/her neighbors except for a small random noise.

R#1. Re. UMG(π) has cycles. We havn't thought about checking it from data, which is an interesting statistical hypothesis testing question for future research. For all neutral models such as Mallows and PL studied in this paper, we have $\pi_{\text{uni}} \in \text{CH}(\Pi)$ (L338), which means that UMG(π_{uni}), which is the empty graph, contains a weak Condorcet cycle.

R#2. Re. Theorem 1 is effectively known. We totally agree with the reviewer that the smoothed avoidance can be proved by existing techniques without using Lemma 1. Thanks for the reference and we will discuss it in the revision. We are not sure if the smoothed paradox part of Theorem 1 can be proved by techniques in the TEAC paper (though it can be proved by multi-variate Berry-Esseen because its $O(\sqrt{n})$ error is small enough), and we'd greatly appreciate it if the reviewer can shed some light.

R#2. Re. calling the proposed framework "smoothed analysis" is confusing. Thanks for the very insightful comments! The "worst-case over profiles, and then expectation over noise added" part in smoothed complexity analysis is actually a special non-parametric model, which is a special case of our framework in the following sense: the parameters are all possible ground truth (e.g. all n-profiles), and the family of distributions are distributions over n-profiles obtained from adding noise to each ground truth—the setting where each agent chooses a (possibly different) distribution from single-agent Mallows (L199) is an example. In accordance with the reviewer's comment about subjective settings, we view the loss function in our framework (the per-profile satisfaction of axioms as the reviewer pointed out) the main conceptual novelty compared to previous work in preference learning, statistical/epistemic social choice (e.g. the JMLR paper and the TEAC paper, which adopt popular statistical loss functions that measures quality of decisions w.r.t. the ground truth, such as MSE and 0-1 loss), and smoothed complexity analysis (whose loss function is the runtime that does not depend on the ground truth). We call the proposed framework "smoothed" for better presentation, because (1) the idea behind smoothed complexity analysis is well accepted in CS, and (2) the conceptual innovation of our framework is similar to that of smoothed complexity analysis: while the modeling of uncertainty and the minimax nature of inference are standard in statistical decision theory, the smoothed (complexity or social choice) analysis provide principled ways and new results to resolve criticisms on average-case analysis in disciplines beyond statistics (algorithm design and social choice, respectively). We will expand the discussions about the second point in L132–134 in the revision, and will add more detailed and technical discussions in the full version.

R#2. Re. distortion, Big-Oh. Distortion is related to the single-agent PL (L204–207) whose ground truth is cardinal utilities and whose votes are rankings. $k \leq m$ is treated as a constant because all results are proved for fixed m. We'd certainly follow the reviewer's suggestions in the revision. Thanks so much for the suggestions!

R#3. Re. "results are incremental". We agree that some high-level messages behind our results align well with people's intuition, but like R#4, we view it a strength of the paper as for smoothed complexity analysis (the simplex algorithm was widely believed to be fast before it was formally proved by Spielman and Teng). We will make it clear in the revision. On a more technical level, it is unclear whether anything (impossibility and possibility) should be expected, because both hold under certain conditions. Our technical contribution is not to prove whether the impossibilities vanish or not, but is to characterize when do they vanish and at what rate. The non-triviality of such characterizations is evident in the vast literature on this topic and many open questions such as the conjecture in Section 4 of Tsetlin et al. [49], which we gave an asymptotic answer (see L140–141).

R#3. Re. "basic results on the folklore impossibility theorem are also expected". We are puzzled by the reviewer's comment because we are not aware of any similar study even under IC, as discussed in L142–143. We'd greatly appreciate it if the reviewer could elaborate on the reasoning and references behind his/her expectation.

R#3. Re. "for Condorcet's theorem, after smoothing, things look the same". If we are not mistaken, this is not true because (1) Condorcet paradox holds in the worst case, and (2) our Theorem 1 (L280–282) states that under certain conditions, Condorcet's paradox *disappears* after smoothing, which is different from before smoothing.

R#3. Re. "the role of AI". Application of AI is used to motivate the large-scale and frequently-used features, which naturally lead to the study of asymptotic satisfaction under realistic statistical models (please see L49–51). We did not mean to hint that AI can break mathematical facts and will make it clearer in the revision.

R#4. We totally agree that smoothed Arrow and other impossibility theorems are natural and important next steps! We will follow your suggestion to move Mallows and PL to the appendix to improve the presentation in the revision.