We thank the reviewers for their insightful comments. The three reviewers agree that the paper provides a novel approach to two-player zero-sum games with non-convex losses. In particular, **R1** states that the paper is clearly of interest for the Neurips community and could lead to interesting results in the future. We address their helpful comments below.

**Reviewer 1.** *1)* The mixed strategy approach involves computing distributions over the parameter spaces of the generator and the discriminator, and the resulting game on mixed strategies is indeed convex. The claim that we address non-convex games is to be understood in the sense that the original losses are non-convex and we study algorithms to solve the lifted convex problem with better performance than mirror descent (i.e. fixing parameters and updating weights), through the use of transport. s *Interpretation of mixed equilibria:* The resulting "mixed generator" is a mixture of distributions, each of them defined by a single generator.

*2)* As pointed out, Alg. 2 can be seen as performing gradient descent on the x parameters and multiplicative weights on the w parameters. At the level of measures, the multiplicative weights algorithm is the Fisher-Rao gradient flow, which is the gradient flow on the space of measures endowed with the Fisher-Rao (or Hellinger-Kakutani) metric. Analogously, gradient descent corresponds to the Wasserstein gradient flow on measures. When combining both algorithms, the dynamics can be seen as a gradient flow in the Wasserstein-Fisher-Rao metric, which is, loosely speaking, computed as the sum of the W an F-R metrics. See Preliminaries of Gallouët and Monsaingeon [2016] for more details (on the optimization case, analogous for games).

*3 and prior work comment)* Balandat et al. [2016] put forward an alternative way to tackle the problem of finding mixed Nash equilibria in compact strategy spaces which avoids dealing with dynamics in measure spaces. They use dual averaging (related to mirror descent) and show regret bounds. In particular, they show that the average regret tends to zero as $t \to \infty$ (Hannan consistency), and Hannan consistency implies that the empirical measures of each player converge to a mixed Nash equilibrium. However, this approach does not yield rates (contrary to us) and lies far from the gradient-based approaches frequently used in ML, which are more closely related to the measure-theoretic approach we take. We will include a more thorough comparison with this alternative work.

*4)* As pointed out, the paper is on multi-agent optimization (two agents) and not multi-objective. This will be corrected.

**Reviewer 3**. *Sample complexity for measure approximation:* We prove that the particle dynamics converge to the measure dynamics as the number of particles goes to infinity using a propagation of chaos argument. Although we do not provide quantitative rates, this convergence is in general dimensionally cursed and exponential in time. These are common drawbacks of the mean-field approach which were also encountered in the mean field analysis of neural networks literature (Mei et al. [2018], Rotskoff & Vanden-Eijnden [2018], Chizat & Bach [2018]). In practice, the number of particles needed to obtain good performance is much lower than the theoretical bound, and lower than the number needed for mirror descent ascent (Figure 1).

*Boundaryless (and compactness) assumptions:* In our theoretical analysis we assume that the parameter spaces $\mathcal{X}, \mathcal{Y}$ are compact Riemannian manifolds without boundary. The compactness and boundaryless assumptions preclude direct application of the theory to typical ML settings such as GANs. While the compactness assumption is necessary for MNE to exist, we introduce the boundaryless assumption to simplify theoretical arguments involving gradient descent and Langevin dynamics (gradient descent on spaces with boundary requires projecting after each step). However, we believe that the results could be extended to manifolds with boundary using the same ideas.

*Comment on the proof of uniqueness of Thm 4:* This proof is based on the argument of Rosen, 1965, which proves uniqueness of strictly convex games. In our case, strict convexity-concavity of the losses follows from the strict convexity of the differential entropy. As long as we can ensure strict convexity, a similar argument should allow us to prove uniqueness.

**Reviewer 5.** - *the authors should apply the method on large-scale datasets*. Evaluating our approach on larger datasets than CIFAR10 would entail training generative models on e.g. ImageNet, which is known to be very costly. We unfortunately lacked resources for this at the time of submission. *Complexity of the approach:* We compare the number of generator updates in Figure 2 and 3, and show that training mixtures is not significantly slower than training a single generator, with an additional clustering effect. Using many discriminators may slow down convergence.

*Further applications.* We thank the reviewer for his suggestions. Robust training is indeed an interesting source of 2-player games, that has been studied in the light of mixed equilibria in e.g. Pinot et al. [2020]. Our transport algorithm could be used to train robust mixtures of classifiers, although we leave this for a more applied future work.