We would first like to thank the reviewers for their insightful comments on our work. We appreciate that you like our paper and that you are helping us making it even stronger with your comments.

(Motivation concerns). One of the questions of Reviewer #2 and #3 was about the connections of our work to training GANs using GDA dynamics. At this point, we would like to clarify the motivation behind our work. Undoubtedly, understanding GAN training dynamics in its full complexity is a very important but also difficult problem. On the other hand, as it was noted in our related work section, the recent research line in min-max optimization has mainly focused on providing guarantees on simple bilinear games (BGs). In this work, we aim to make a step towards bridging this gap. In this effort, we identify two key components of GAN training that are not captured by BGs: GAN training is *i) a non-convex non-concave minmax problem* and more importantly *ii) it involves indirect competition between the two players (networks)*. Goodfellow recognized in [4] that *ii)* is important: "In practice, however, the updates are made in parameter space, so the convexity properties that the proof relies on do not apply." These observations led us to propose the study of hidden bilinear games which combine both those properties. In order to avoid any potential source of confusion, we plan to modify the corresponding parts of abstract, introduction and conclusion sections giving more emphasis on the aforementioned properties *i)* and *ii)* than on actual GAN architectures.

(Potential generalizations) The majority of the questions of Reviewer #1 were centered around potential generalizations. In regards to moving to non-zero sum games where there could be more than one equilibrium, we should distinguish two elements: 1) Our setting of hidden bilinear games already allows for games with many equilibria, e.g. due to stationary points of non-convex functions f, g. 2) Moving completely beyond zero-sum games creates a setting that is so general that it would be hard to identify usable structure. Nevertheless, we do agree that it is very interesting to understand how much can we generalize our current setting while still salvaging interesting provable properties. Our modular setting makes it particularly amenable to such generalizations (e.g. add more than two agents that play a hidden network-zero-sum game [2], introduce different non-linearities, etc). We believe that our theoretical techniques extend to more general settings, and although these questions are beyond our current scope, we hope that our work will enable for interesting followups.

(Technical Clarifications) The majority of the questions that Reviewer #3 posed revolve around some technical aspects of our work. We would like to point out that even both Section 4 and 5 are framed as classical games where each player submits a distribution to the underlying BG, our results extend beyond this case. For example one can analyze Section 5 without using Lagrange multipliers yielding an unconstrained minmax problem. See also the introduction section of [3], where a similar approach was followed. Regarding the question about the equivalence of Equation 6 and Definition 1, one can check the value of the objective function in Equation 6 is finite only if the sum to one constraints are satisfied. Thus if θ, ϕ has been chosen properly to satisfy them, any possible penalization by λ, μ can not affect the value of the game. Actually, the only difference in the non-convex non-concave case is that KKT conditions are no longer sufficient which does not affect the optimization dynamics. As far as the question about the limitation of constant step size η in Theorem 9 is concerned, the proof of Theorem 16 that we rely on does not require η to be the same at each iteration, as long as it remains positive. Reviewer #3 expressed also some interest in understanding the motivation behind the transformation of line 217. The major effect of this transformation is that it yields a dynamical system that is bipartite.This means that we can split the state of the system in two sets of variables so that the derivatives of the variables in one set depends only on the values of variables in the other set. Bipartite systems are known to be easier to analyze in terms of interesting properties like invariants and volume preservation [5, 1] . More explicitly, under line 225, in equation (5), RHS of the ODE describing $\dot{a}_i$ is a function that is independent of a_i. It is simpler to understand intuitively this observation in a standard zero-sum game such as Rock-Paper-Scissors when studied under GD. It means that the performance of the first player when playing a specific strategy e.g. Rock, does not depend on anything that he does, but on his opponents behavior.

References

[1] J. P. Bailey and G. Piliouras. Multi-agent learning in network zero-sum games is a hamiltonian system. In *AAMAS*, 2019.

[2] Y. Cai, O. Candogan, C. Daskalakis, and C. H. Papadimitriou. Zero-sum polymatrix games: A generalization of minmax. *Math. Oper. Res.*, 2016.

[3] C. Daskalakis and I. Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *ITCS*, 2019.

[4] I. J. Goodfellow. NIPS 2016 tutorial: Generative adversarial networks. *CoRR*, abs/1701.00160, 2017.

[5] J. Hofbauer. Evolutionary dynamics for bimatrix games: A hamiltonian system? *Journal of mathematical biology*, 1996.