We thank the reviewers for their excellent feedback. All the comments showed a strong understanding of the paper and will be useful for presenting our results in the best form. We agree with all suggested minor changes and will update the paper accordingly. The remaining comments focus mainly on an interest in a higher dimensional argument and other algorithms.

Higher Dimensions: Reviewer 2 and 3 asked about extending the proof to higher dimensions. Our proof consists of the three components listed below.

(1) the "step-size" in the dual space is bounded; i.e., $a \leq ||y_i^t - y_i^{t-1}|| \leq b$.

(2) a proof of divergence in the dual/payoff space where the divergence grows linearly when at least one agent is not playing a pure strategy and negligibly when both agents are playing a pure strategy.

(3) a proof of recurrence where the "cycle" length (in the strategy/primal space) is bounded.

The first two components immediately extend using our current analysis. In regards to the last step, recent advancements in understanding the geometry of learning dynamics in larger games (e.g., Mertikopoulos et al, 2018; Bailey and Piliouras, 2019) suggests that although non-trivial this last step can also be eventually rigorously established. New ideas, however, will most likely be needed for this last step. Given our framework, we believe that as more techniques are developed to understand these trajectories, sublinear regret will readily follow using our approach.

Other Algorithms: Reviewer 3 expresses an interest in analysis of algorithms similar to gradient descent. Reviewer 3's hunch that a similar analysis could work on other algorithms is on point. Our proof technique should extend to all FTRL algorithms via the dual space discussed in our paper. Both (1) and (3) trivially extend to this framework. As discussed in Appendix I, the proof for (2) mainly uses the strict convexity of the regularizer in FTRL. We opted to focus this paper on gradient descent because the core principles behind the proofs are most clear when working with the regularizer $h(x) = ||x||^2$. In particular, GD has two nice properties:

(a) the dual space is the affine hull of the strategy space and the mapping from the dual space to the primal space is a simple projection. This property isn't particularly useful in the proofs. Rather, the dual space is still an underutilized concept and we believe the techniques are most accessible and most likely to be reused if there is a simple connection between the two spaces.

(b) after a finite number of iterations, all primal strategies appear on the boundary. This does not hold for all FTRL algorithms (e.g., MWU always selects fully-mixed strategies). For other FTRL algorithms and for any ϵ, after a finite number of iterations all strategies will appear within ϵ of the boundary. The proof of this follows identically to the proof in Appendix G. However, to establish linear growth in divergence as in Appendix E.3., ϵ will likely have to be carefully selected for each cycle. This is because for an algorithm like MWU, the convex conjugate h^* is never linear; rather it becomes arbitrarily close to a linear function as both agents come closer to playing a pure strategy.

We are certainly open to including an additional appendix to discuss the extension to FTRL. However, we strongly believe that the paper is best presented through the lens of gradient descent. By taking advantage of (a) and (b) while introducing our techniques, we believe researchers will find our analysis more accessible and will be more likely use similar ideas to advance the understanding of online optimization.

Other comments: On a last note, we agree with reviewer 2 that experiments are only suggestive that the stronger regret bounds extend to larger zero-sum games. We will make sure to make that point more clear.