We thank all reviewers for their nice reviews.

Reviewer #1: **[Re "...important results...game theory community rather than the machine learning community" & "the machine learning community may not be interested"]** We are glad to see that the reviewer appreciates the results in the paper and thinks that they are important. We strongly disagree that this paper is not a good fit for NeurIPS, and apparently so do the other reviewers. Papers in game theory are explicitly welcomed by NeurIPS (see CFP, which contains a list of welcomed topics). Also, in recent years there has been growing synergy between the game theory community and the broader machine learning community. NeurIPS has been an important cog in this process, with several seminal papers on computational equilibrium finding (e.g., [Zinkevich et al. NeurIPS06] on counterfactual regret minimization and [Brown&Sandholm NeurIPS17] on the inner workings of the Libratus poker AI). In (at least) 2017 and 2015, NeurIPS best paper awards were given to papers in computational equilibrium finding! **[Re "main difference...from the previous work..."]** This is not accurate: the proposed algorithm uses a first-order method (unlike the linear programming-based method of von Stengel and Forges), and it is based on a very different formulation (saddle-point problem instead of linear program). It is a proof of concept of the benefits of the saddle-point formulation: please see also the great insights of Reviewer #2 on this, which further crystallize the importance of such a formulation. In any case, the first goal of this paper is *not* to propose a faster algorithm for computing a social-welfare maximizing EFCE. The main goals and contributions of the paper are as in lines 51–75 in the paper. **[Re "convergence rate of the proposed algorithm?"]** The convergence rate of the algorithm by Wang and Bertsekas is still an open problem. However, convergence is guaranteed by a supermartingale convergence argument. We'll include a discussion of what is known about the algorithm's convergence rate in the final version of the paper.

Reviewer #2: **[Re "...2-player case without chance...bit restrictive."]** The reason why we focus on the two-player case with no chance is because in this case a *social-welfare-maximizing* EFCE can be computed in polynomial time, unlike games with more than two players and/or chance nodes [von Stengel and Forges, 2008]. For the same reason, the algorithm by Dudik and Gordon does not give any polynomial run-time guarantees about social-welfare-maximizing EFCE in those more general games. Furthermore, their algorithm operates with normal-form correlation plans—an exponentially big set—and uses MCMC with a tacit assumption that sampling from the proposed correlated distribution is practical. Finally, our formulation allows for the computation of EFCEs with *convex* utilities, while their method only allows for linear (regularized) ones. We'll include this discussion in the final version of the paper. **[Re "...n-player variants..."]** Interesting question! We don't know. Unfortunately, computing social-welfare-maximizing EFCE in multiplayer games is a hard problem, so this does not seem like an easy task. **[Re "I would encourage the authors to publish source code..."]** Yes! That was already our intention (see Line 63 in our paper), and we agree that it will be a step in the direction of wider accessibility of the EFCE solution concept. **[Re "...online convex optimization..." & "...Nesterov's excessive gap..."]** Yes! Regret-based methods, as well as Nesterov's EGT algorithm, are techniques that can solve convex-concave saddle-point problems like EFCE. In a way, that was one of the major inspirations of our paper: we believe that the saddle-point formulation will be important for designing scalable first-order methods, just like the saddle-point formulation for Nash equilibrium was crucial for regret-based methods like CFR and EGT-based methods like the one by Kroer et al. that the reviewer mentioned. The hurdle for designing efficient regret-based methods is constructing specialized regret minimizers for the polytope of correlated strategies, and the hurdle for EGT would be designing a smoothing scheme that can be optimized over efficiently (as the reviewer suggests). We hope that our paper will serve as a starting point for all these interesting directions of exploration. We'll make sure to include this discussion in the final version of the paper. **[Re "...fictitious play..."]** Yes, fictitious play could be used in this case. In order to make that efficient, one would need to figure out a way to compute a best response over the polytope of correlated strategies. **[Re "end-to-end learning"]** Yes! This could be possible by specifying deviators (Line 146) and correlation plans implicitly as neural networks. We thank the reviewer again for all these incredibly stimulating ideas.

Reviewer #3: **[Re "...comparison...LP formulation...saddle-point..." & "...purely-mathematical reduction..."]** Going from a bilinear saddle-point formulation (i.e., a min-max problem) to a linear program is always possible; in our case, we show that our saddle-point formulation can be used to recover the LP proposed by von Stengel and Forges (see Appendix A). The opposite direction is significantly harder: given a minimization problem over a set of variables, it is not easy to figure out which of those come from the dualization of an internal max problem. Black-box approaches such as Lagrangian or Fenchel duality are not useful in this case, as they do not recover the original min-max structure: they simply add more "dual" variables, whose intuitive meaning for the problem at hand is often not immediate. **[Re "...provide readers more intuition why the saddle-point formulation is more efficient in practice."]** Please see the answer to Reviewer #2 re "...online convex optimization..." and "...Nesterov's excessive gap...". **[Re "...this benchmark game can hardly be scalable computationally."]** We understand where this comment is coming from. However, despite their seemingly small size, our game instances have a huge number of states and actions per player (up to millions). Going forward, an interesting challenge will be around employing abstraction and approximation techniques that will allow one to scale to larger games and construct mediators that can handle larger interactions. However, our paper already significantly pushes the boundary of what can be done and has been explored so far by orders of magnitude. We will discuss this point more in the final version of the paper.