Reviewer 1: Thanks! We will include some intuition for Lemmas 1 and 2. If the paper is accepted, we will include additional details on the proofs with the ninth content page for the camera-ready version. An accurate theoretical analysis of the convergence rate is left as a future development of this work. The present paper presents a framework that can work with arbitrary external and internal regret minimizers. The convergence rate will definitely be impacted by the specific choice of such regret minimizers.

Reviewer 2: Thanks for your feedback!

- Re "*theoretical results are straightforward but worthwhile*". We strongly disagree that the results are straightforward. Deriving no-regret dynamics for correlated equilibria in extensive-form games has been a challenging open problem for years. Related work trying to solve the same problem date back to the early '00s, which may be an indicator that the solution to the problem was not obvious.
- Re "*finite-time analysis*". We leave giving sharp finite-time bounds as an open future direction.
- Re clarity. Thanks for your feedback! We will expand the description of the algorithm and the example in the final version, using the ninth page of content.
- Re "*set $\Sigma_i^c(I)$...is undefined*". We are not sure which symbol Σ you were referring to. The symbol $\Sigma_i^c(I)$ is defined on line 243, Σ_i is defined on line 102, and $\sigma(I)$ on line 104.
- Re "*terms of $\Sigma_i^c(I)$ in addition to $\Pi(\sigma(I))$ where sigma is again undefined*". The symbol $\Sigma_i^c(I)$ is defined on line 243. We will improve the wording and add a small example to better illustrate the intuitive meaning of $\Sigma_i^c(I)$.
- Re "*lines 252-256*". On lines 252 and 253, the instances of "all but one" should read "only one". We apologize for the typo.
- Re "*The best guess...along the line of play*". Your understanding is correct.
- Re "*The example also does not describe the regret update procedure*". Thanks for the feedback. We will fully work out the example in the final version using the ninth page of content.
- Re "*include a citation to Zinkevich (2008)...it would be worth pointing out*". Yes, we will. What you wrote about laminar regret vs CFR is correct, and we will make sure to point out the connection.
- Re "*Pure CFR*". Thanks for the pointer. The regret updates in Pure CFR seem very different from those of ICFR. From our understanding, Pure CFR instantiates a single external regret minimizer per decision point, while ICFR requires one internal and several external regret minimizers for each decision point. So, it is not immediately clear to us how ICFR could be a special case of Pure CFR.
- Re "*should (7) be*". Yes, good point, thanks!
- Re "*this would make it easier to read*". Thanks for the feedback, we will.
- Re "*Roughly five of this paper's eight pages are used as setup for the paper's contributions [sic]*". Since our paper combines many different concepts and tools (internal and external regret minimization, correlated equilibria, extensive-form games, etc), having a large preliminary section is unavoidable. Luckily, NeurIPS makes a ninth content page available so we will be able to expand the description of the algorithm and the example, and depending on the remaining space we will add a conclusion section accordingly.
- Re "*It is stated that ICFR is much more scalable than alternative algorithms, but there is no clear summary and accounting of the computational requirements for ICFR compared to its peers*". We never claimed that our algorithm outperforms the algorithm by Dudik and Gordon or the Ellipsoid-against-hope algorithm. However, we now offer some reasons why we believe it is reasonable to assume so. The Ellipsoid-against-hope is based on the ellipsoid algorithm and it is known in the community to be very impractical. On the other hand, the algorithm by Dudik and Gordon runs MCMC at every iteration (which is expensive) and is prone to numerical difficulties. To our knowledge, it was never tried beyond the original paper, and its implementation seems like a major effort. We argue that it is reasonable to expect that our algorithm will perform significantly better. Our algorithm is easy to implement, as it combines internal and external regret minimizers all of which can be developed and tested in isolation. It is decentralized, so each agent can be developed separately. And it relies on internal and external regret minimizers for simplex domains, for which strong practical algorithms have been developed in the past twenty years.
  At any rate, that is not the point of the paper. The main point of the paper is that we give the *first* (decentralized) no-regret dynamics for EFCE and EFCCE. The algorithms by Dudik and Gordon and the Ellipsoid-against-hope algorithm do not provide no-regret dynamics. We leave the comparison of those algorithm in the context of computing one correlated equilibrium for future works in this space.

Reviewer 3: Thanks! If the paper is accepted, we will include in the camera-ready version a comment on the number of normal-form plans that the algorithm needs to store. The development of a procedure to avoid recording all iterates is an interesting future development. As for the scalability of our algorithm, we hope ICFR can serve as a foundation for more scalable model-free RL methods, which we believe should be a key long term goal in the multi-agent RL research agenda. Finally, we will add more details on the experimental setting in the supplementary material, so as to improve the reproducibility, and we will include a comparison on the quality of the ICFR solution vs CFR. Exploring this aspect from a theoretical perspective is another interesting research direction.