[Rev.1: on Th. 4 and 5]. Th. 4 and 5 should be evaluated as the key building blocks of CFR-Jr, which is shown to outperform CFR and CFR-S in practice. Th. 4 is necessary to show the soundness of the reconstruction algorithm, which allows for working with compact normal-form strategies. Th. 5 shows that CFR-Jr approaches the set of CCEs, which is the goal of the algorithm. **[Rev.1:Q1]**. The reviewer is right. Indeed, almost all previous works on no-regret learning in games focus on the *easy* setting of two-player, zero-sum games, where optimality of the solution is always guaranteed. In multi-player, general-sum games, the problem of bounding the distance from the optimal equilibrium is a non-trivial open problem. As a first analysis, we provide some experimental evidence showing that we are able to approximate a nearly-socially-optimal equilibrium (e.g., Figures 3,5 and Tables 3,4). **[Rev.1:Q2]**. The time-complexity for one iteration of CFR-Jr is the same as the one for Alg. 2. CFR-Jr basically adds a traversal of the tree (requiring linear time), so the full time complexity is in the order of $T|Z|^2$. We will make this explicit in the paper. **[Rev.1:Q3]**. This can never happen, as plans $\bar{\sigma}$ built by the reconstruction procedure are always different (see the proof for Th. 4, Appendix D2, lines 632-635). **[Rev.1:Q4]**. The regret experienced by CFR-Jr is bounded by $\Delta|\mathcal{I}_i|\sqrt{|A_i|}/\sqrt{T}$, as in CFR [43], because our reconstruction procedure does not alter the way in which regret is minimized. We will make this explicit in the paper. **[Rev.1:Q5]**. We employed the optimal payoff (the maximum sum of players' utilities), which is not guaranteed to be achievable by a CCE, as an upper bound on the optimal social welfare. Employing an upper bound was necessary as CG could not scale to the larger instances of our test bed. We will clarify this in the paper. **[Rev.2:CFR-Jr]**. CFR-Jr has two fundamental components: a regret minimizer employing compactly representable strategies, and a poly-time reconstruction oracle (Alg. 2). The former comes almost directly from the available literature, but the latter was not obvious and is a key contribution of this paper. **[Rev.2: CFR-S]**. The sampling procedure has to take place at all information states because CFR-S needs a complete normal-form plan to keep track of the joint empirical frequency of play. Sampling paths of execution (from the root to a leaf) and recording them does not allow to build a well-defined correlation device, and the same holds for average marginal strategies (see Sec. 4.1 and Fig. 2). In general, CFR-S performs worse than CFR-Jr, as it needs much more iterations to converge. It has however a better per-iteration time complexity (linear in the input, as for vanilla CFR), and, thus, it could be preferred for very large game instances in which explicitly reconstructing the joint strategy (as in CFR-Jr) would be demanding. **[Rev.2: EFCEs]**. At the moment, there are no extensions of the CCE similar to the EFCE (we are aware of a team currently working on a notion of coarse EFCE). Providing a characterization of a new solution concept is beyond the scope of this paper. **[Rev.3: dominated actions]**. Properties shown in [19] should carry over to CFR-Jr as the core regret minimizer from CFR is left unchanged, and reconstruction of joint strategies does not introduce actions that were played with zero probability in the marginal strategies of the players. **[Rev.3: naïve algorithm]**. We explored the possibility of employing the reviewer's algorithm, but we decided not to include it for a number of reasons: i) We are interested in providing an explicit representation of the equilibrium (i.e., a joint probability distribution, as it is customary in the equilibrium computation community). In order to provide such a representation from T behavioral strategy profiles, a reconstruction step would still be required, limiting the possible advantages of such approach. ii) In practice, CFR-Jr allows to build dramatically smaller solutions, e.g., the figure displays the percentage difference between CFR-Jr and the reviewer's algorithm storage spaces. The figure considers G2-4 with different tie-breaking rules. iii) CFR-Jr is amenable to further improvements to reduce its storage space. E.g., we are currently working on heuristic approaches that perform the reconstruction step by prioritizing plans already selected in the past. These are part of the reasons for which we presented CFR-S as a naïve baseline (see next answer) instead of an algorithm analogous to the reviewer's one. **[Rev.3: CFR-S]**. Other than being used as a baseline, CFR-S is also instrumental in clarifying what needs to be added to vanilla CFR in order to find a CCE. Moreover, CFR-S is already faster than previous state-of-the-art approaches (based on MILPs). **[Rev.3: CG]**. See answer to Q.5 of Rev.1 for more details on how we compute the social welfare ratio. When we write >24h as the runtime of CG, we mean that the execution was killed by the time limit imposed on the system. The utility ratio is reported only for general-sum games (in zero-sum games, e.g., Kuhn, the social welfare is always zero). **[Rev.3: size of the final solutions]**. In typical scenarios, the support of the joint strategies will often be the same among different iterations, so that the final solution size will be significantly smaller than the worst case. The figure shows the ratio against $T|Z|$, the ratio against the worst case would be even smaller as the denominator would be $T|Z|^2$. **[Rev.3: size of reconstructed strategies]**. The size of a reconstructed normal-form strategy profile is upper bounded by the number of leaves reached with strictly positive probability when following the original behavioral profile, which is, in its turn, upper bounded by the size of the support of the latter strategy (i.e., the number of non-zero action probabilities). Then, the reconstructed normal-form strategy for a single player has size upper bounded by the size of the original behavioral strategy. **[Rev.3: experimental evaluation]**. The reviewer claims that the experimental comparison of running times is not surprising. However, we remark that CFR-Jr has been compared with the current state-of-the-art technique for this setting (CG), and also with a new baseline algorithm (CFR-S). Moreover, CFR-Jr was also evaluated against the de facto standard in two-player games (CFR). CFR-Jr shows better convergence and attains a higher social welfare than all the other techniques.