We would like to thank all the reviewers for their constructive feedback. In the following, we respond (**R**) to individual concerns (**C**) summarized in italic. Citations refer to references in the paper.

Reviewer 1. C: *"...how does the term γ_T change as we increase number of actions K, and number of agents N?"* **R:** The Reviewer is correct in that the term γ_T can dominate the regret bound, however for most widely used kernels γ_T grows only with the dimension d of the domain $\mathcal{D} = \Pi_{i=1}^N \mathcal{A}^i \times \mathcal{Z}$ and *not* with the number of actions $K = |\mathcal{A}^i|$ available to player i. Note that although the possible action tuples grow exponentially with the number of players N, the dimension d grows only linearly with N. **C:** *"...Why should the price of anarchy be bad if there is just one round t that gives poor choices of $\lambda(z_t)$ and $\mu(z_t)$?"* **R:** The obtained price of anarchy bound reflects the worst case in which, if the game is non-smooth even for a single round, the welfare function Γ for that round could have a much higher contribution than the others and thus highly deteriorate the performance. Nevertheless we agree that, based on the Reviewer's reasoning, tighter conditions could perhaps be found as a function of the contexts' sequence.
Finally, we will clarify the connections with respect to the mentioned previous works [Slivkins, 2011] and [Lu et al. 2010] that study similarity in contextual bandits. Lu et al. consider a stochastic setting and exploits Lipschitzness of the single reward function $r(a_t, z_t)$, while Slivkins studies similarity even in contextual adversarial bandits (i.e., where the player faces an adversarial sequence of reward functions $r_t(a_t, z_t)$). In a contextual game the rewards obtained by player i are generated by the game reward function $r(\cdot)$ which, however, also depends on the actions a_t^{-i} chosen by the other players. That is, a contextual game is an adversarial contextual bandit problem where $r_t(a_t, z_t) = r(a_t, a_t^{-i}, z_t)$. This makes our model unique and different from previous work in that we impose kernel-based regularity assumption on $r(\cdot)$ and study the similarity across different game outcomes (a_t, a_t^{-i}, z_t). We will clarify these aspects in the paper.

Reviewer 2. C: *"...it's not clear why would players agree to run Algorithm 1 while they can adopt dynamic sophisticated strategies (i.e., Folk Theorems)"* **R:** Although players are aware of playing against N-1 agents, these agents may be non-rational or even adversarial. In such a case we believe the proposed no-regret algorithms represent a simpler and more robust playing choice for the agents, which come with individual performance guarantees, as opposed to using Folk theorems. Moreover, they do not require the game to be known in advance. This certainty leads to a simpler non-dynamic equilibrium notion (c-CCE), which however can certify a certain level of game welfare as discussed in Section 4.
C: *"The paper indeed generalizes repeated games, but the equilibrium analysis definitely doesn't generalize the equilibria of repeated games"* **R:** We agree with the Reviewer, and we will make this clear. Our notion is a simpler, non-dynamic, and more tractable notion of equilibrium which takes into account the context information but does not have 'memory' of past rounds. The repeated game is essentially summarized with its time-average and, in this regard, c-CCEs generalize CCEs for non-contextual games (where the time-average coincides with the one-shot game).
C: *"...it's not like Proposition 4 holds for any T, and the approximation gets better with time in an online manner. It only holds if the parameters were tuned using T."* **R:** We remark that the horizon T does not necessarily need to be known in advance. In Theorem 1 the doubling trick can be used to tune the radius ϵ (see, e.g., [20, Remark 1]), while in Theorem 2 a time-varying learning rate η can be chosen (see, e.g., [11, Theorem 2.3]). Hence, Proposition 4 holds for any T and the c-CCE approximation gets better with time in an online manner. We will clarify this in the paper.
C: *"This is quite subjective, but I didn't really understand the motivation for the i.i.d. contexts case."* **R:** We agree with the Reviewer that the discussion in Section 3.3 diverges from the main point of the paper; we decided to include it only as a special case in which players can achieve significantly improved performance. However, note that even if players perfectly knew the contexts distribution, the pseudo-regret in Theorem 2 is a function of the *realized* contexts and therefore identifies a non trivial benchmark. Moreover, our motivation for Section 3.3 are settings in which the context is private information only relevant to player i (e.g., value for items in an auction, or production costs in a market) and therefore the other players do not have access to it (they can still base their decision on other private contextual information). Finally, we thank the Reviewer for other comments that can improve the exposition of the paper.

Reviewer 3. C: *"It would have been nice to study the equilibrium quality empirically as well, in addition to the evaluation of the algorithms."* **R:** Besides comparing the players' individual performance, Figure 1 also shows that the game welfare (i.e., the sum of players' payoffs) increases and the network congestion level decreases as the players approach a contextual CCE. Hence, it empirically demonstrates the quality of the computed c-CCE with respect to playing non-contextual CCEs (cyan line) and not learning (purple line).
C: *"I was wondering how the notion of contextual game relates to that of succinct games."* **R:** We believe the notion of a game having a succinct representation can co-exist with our contextual game setup, as a function of the contextual information. For instance, for any given context (e.g., network occupancy profile) the considered traffic routing game is a congestion game (hence a particular kind of succinct game). It is an interesting future direction to understand if the computational advantages of finding correlated equilibria in succinct games can be lifted to compute contextual correlated equilibria.

Reviewer 4. We thank the Reviewer for the positive feedback and comments. It is indeed an interesting future work to extend the present models and techniques to sequential settings.