LEARNING LARGE-SCALE COMPETITIVE TEAM BE-HAVIORS WITH MEAN-FIELD INTERACTIONS AND ON-LINE OPPONENT MODELING

Anonymous authorsPaper under double-blind review

000

001

002

004 005 006

007

008 009 010

011 012

013

014

015

016

017

018

019

021

023

025

026

027

028

029

030

032033034

035

037

040

041

042

043

044

045

046

047

051

052

ABSTRACT

While multi-agent reinforcement learning (MARL) has been proven effective across both collaborative and competitive tasks, existing algorithms often struggle to scale to large populations of agents. Recent advancements in mean-field (MF) theory provide scalable solutions by approximating population interactions as a continuum, yet most existing frameworks focus exclusively on either fully cooperative or purely competitive settings. To bridge this gap, we introduce MF-MAPPO, a mean-field extension of PPO designed for zero-sum team games that integrate intra-team cooperation with inter-team competition. MF-MAPPO employs a shared actor and a minimally informed critic per team and is trained directly on finite-population simulators, thereby enabling deployment to realistic scenarios with thousands of agents. We further show that MF-MAPPO naturally extends to partially observable settings through a simple gradient-regularized training scheme. Our evaluation utilizes large-scale benchmark scenarios using our own testing simulation platform for MF team games (MFEnv), including offense-defense battlefield tasks as well as variants of population-based rock-paperscissors games that admit analytical solutions, for benchmarking. Across these benchmarks, MF-MAPPO outperforms existing methods and exhibits complex, heterogeneous behaviors, demonstrating the effectiveness of combining meanfield theory and MARL techniques at scale.

1 Introduction

Existing state-of-the-art MARL algorithms built upon MADDPG, MAAC and MA-PPO (Lowe et al., 2017; Yu et al., 2022), face severe scalability challenges as the number of agents grows, primarily due to the well-known *curse of dimensionality*. A promising remedy is offered by mean-field theory, which approximates large-scale agent—environment interactions in the infinite-population limit (Huang et al., 2006). Two major areas of mean-field research are mean-field games (MFGs) (Huang et al., 2006; Lasry & Lions, 2007; Sen & Caines, 2019; Laurière et al., 2022), which focus on non-cooperative agents, and mean-field control (MFC) problems (Bensoussan et al., 2013; Gu et al., 2021), which study fully cooperative scenarios. In contrast, mixed collaborative—competitive scenarios that arise in many real-world domains, such as team sports (Gaviria Alzate et al., 2025) and social dilemmas (Leibo et al., 2017), remain relatively underexplored. To address this gap, we propose Mean-Field Multi-Agent Proximal Policy Optimization (MF-MAPPO), the first PPO-based *learning* algorithm tailored for mixed cooperative—competitive mean-field settings. Guided by existing theoretical results (Guan et al., 2024a), MF-MAPPO scales to hundreds or thousands of agents while preserving convergence guarantees and remaining agnostic to individual identities or private observations.

Mean-Field Teams. The single-team problem was explored in Arabneydi & Mahajan (2014), where agents share a common team reward (MFC). In contrast, Mahajan & Nayyar (2015) established optimality for finite-population games *only* in the LQG setting, while Sanjari et al. (2022) analyzed a two-team setting with continuous states and actions—unlike our finite state-space formulation that directly admits the familiar MDP-type structure. Similarly, *multi-population* MFGs (MP-MFGs) have been studied in the past but often restrict agent dynamics and policies to be independent of both other agents and other population distributions, see Perolat et al. (2021) and references therein. More

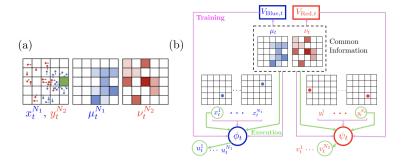


Figure 1: (a) Battlefield as a ZS-MFTG (b) Overview of the architecture of MF-MAPPO.

recently, Guan et al. (2024a) introduced *zero-sum mean-field team games* (ZS-MFTGs), modeling large-population teams that compete while cooperating internally. A common-information decomposition (Nayyar et al., 2013) enables team-size-independent learning with identical optimal team policies using MF feedback, unlike aforementioned MP-MFGs (open-loop MF policies). However, computing such policies numerically, especially in large state-action spaces using dynamic programming, is costly. By contrast, MF-MAPPO leverages shared actor–critic networks per team and uses only commonly accessible information (exact/estimated), ensuring tractability and scalability.

Mean-Field Theory and Learning. Recent advances in mean-field learning span Q-function—based methods, such as MF-Q and MF-AC (Yang et al., 2020) and DDPG-MFTG (Shao et al., 2024b), to value-function-based methods like Dec-POMFPPO (Cui et al., 2024) (MFC only). While DDPG-MFTG incorporates team games, it is restricted to simple grid worlds, unlike our focus on tightly coupled collaborative-competitive domains (ZS-MFTGs). We adopt it as a baseline and show that MF-MAPPO consistently outperforms it in stability and performance. Other related works include PMD-TD for MFGs (Yardim & He, 2024) and GAN-based ECA-Net for continuous-space attack-defense games (Wang et al., 2022), both differing in scope and structure. Moreover, Yang et al. (2020) define mean fields over neighboring actions rather than the full state space. MF-MAPPO extends PPO (Schulman et al., 2017) to competitive MFTGs, using team distributions as critic inputs for scalability, a shared actor-critic per team with a single buffer for efficiency, and simultaneous team training to avoid the inefficiencies of iterative best-response methods (Lanctot et al., 2017; Smith et al., 2021). Unlike prior MF methods that rely on infinite-population oracles (Shao et al., 2024a; Perolat et al., 2021; Carmona et al., 2021), MF-MAPPO is trained directly in finite-population simulators, making it suitable for realistic deployment. Finally, to standardize evaluation in large-scale MFTGs, we create novel MFTG benchmark environments (Constrained Rock-Paper-Scissors, Battlefield) going beyond existing ones that either focus only on MFGs (Guo et al., 2023) or omit MF coupling altogether (Terry et al., 2021; Zheng et al., 2018).

Mean-Field Estimation and Opponent Modeling. To enable reactive behavior to opponents' actions, most existing MF approaches assume centralized or exact knowledge of the opponent's MF, which is rarely practical. While Cui et al. (2024) considers partial observability, it is limited to MFC. Estimation methods such as kernel density estimation (Inoue et al., 2021) or normalizing flows (Perrin et al., 2021) struggle with discrete spaces and limited agent-visibility. Communication-based methods (Benjamin & Abate, 2025b;a) address these constraints but assume uniform estimates for unobserved states and rely on perfect multi-round communication while being restricted to fully cooperative or competitive regimes; noisy variants may even produce invalid distributions. We instead propose Dynamic-Projected Consensus (D-PC), a constrained consensus algorithm that ensures validity, exponential convergence, and bounded deviations when paired with a *gradient-regularized* MF-MAPPO policy. Gradient regularization naturally stabilizes MF-MAPPO in partially observable MFTGs. Experiments show that D-PC matches baseline performance and even outperforms them under limited communication, especially critical in adversarial settings requiring rapid adaptation (Richards et al., 2012), enhancing robustness and fault tolerance. To our knowledge, this is the first use of MF estimation for opponent modeling in competitive team settings.

Our contributions. The main contributions of our work can be summarized as: 1) MF-MAPPO, a scalable shared-actor-critic algorithm for large-scale MFTGs; 2) novel MFTG benchmarking environments (MFEnv) for validating MARL scalability; 3) A gradient-regularized extension of MF-MAPPO coupled with a decentralized mean-field estimation framework D-PC, with theoretical

performance guarantees in partially observable MFTGs; 4) comprehensive numerical experiments demonstrating MF-MAPPO and D-PC's superior performance and efficiency over existing baselines.

2 PROBLEM FORMULATION

2.1 ZERO-SUM MEAN-FIELD TEAM GAME

The zero-sum mean-field team game models a discrete-time stochastic game between two large teams of agents (Guan et al., 2024a). The Blue and Red teams consist of N_1 and N_2 identical agents for each team, with the total number of agents being $N=N_1+N_2$. Let $X_{i,t}^{N_1}\in\mathcal{X}$ and $U_{i,t}^{N_1}\in\mathcal{U}$ represent the state and action of Blue agent $i\in[N_1]$ at time t. Here, \mathcal{X} and \mathcal{U} are the finite state and action spaces of the Blue team. Similarly, $Y_{j,t}^{N_2}\in\mathcal{Y}$ and $V_{j,t}^{N_2}\in\mathcal{V}$ denote the state and action of Red agent $j\in[N_2]$ at time t. The joint state-action variables for the Blue and Red teams are denoted as $(\mathbf{X}_t^{N_1},\mathbf{U}_t^{N_1})$ and $(\mathbf{Y}_t^{N_2},\mathbf{V}_t^{N_2})$, respectively. We denote the space of probability measures over a set E as $\mathcal{P}(E)$. Below, $d_{\mathrm{TV}}(\mu,\mu')$ represents the total variation between $\mu,\mu'\in\mathcal{P}(E)$.

Definition 1. The empirical distributions (ED) for the Blue and Red teams are defined as

$$\mathcal{M}_{t}^{N_{1}}(x) = \frac{1}{N_{1}} \sum_{i=1}^{N_{1}} \mathbb{1}_{x}(X_{i,t}^{N_{1}}), \ x \in \mathcal{X}, \ \text{ and } \ \mathcal{N}_{t}^{N_{2}}(y) = \frac{1}{N_{2}} \sum_{i=1}^{N_{2}} \mathbb{1}_{y}(Y_{j,t}^{N_{2}}), \ y \in \mathcal{Y},$$
 (1)

where $\mathbbm{1}_a(b)=1$ if a=b and 0 otherwise. Specifically, $\mathcal{M}_t^{N_1}(x)$ gives the fraction of Blue agents at state x and, similarly, for $\mathcal{N}_t^{N_2}(y)$. We use $\mathcal{M}_t^{N_1}=\mathrm{Emp}_{\mu}(\mathbf{X}_t^{N_1})$ and $\mathcal{N}_t^{N_2}=\mathrm{Emp}_{\nu}(\mathbf{Y}_t^{N_2})$ to denote the EDs computed from the given joint states. Note that the Emp operators remove agent index information, so one *cannot* determine the state of a specific Blue agent i from $\mathcal{M}_t^{N_1}$.

We consider weakly-coupled dynamics where the dynamics of each individual agent is coupled with other agents through the EDs (Huang et al., 2006; Sanjari et al., 2022). For Blue agent i, its stochastic transition is governed by the transition kernel $f_t : \mathcal{X} \times \mathcal{U} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \to \mathcal{P}(\mathcal{X})$:

$$\mathbb{P}(X_{i,t+1}^{N_1} = x_{i,t+1}^{N_1} | U_{i,t}^{N_1} = u_{i,t}^{N_1}, \mathbf{X}_t^{N_1} = \mathbf{x}_t^{N_1}, \mathbf{Y}_t^{N_2} = \mathbf{y}_t^{N_2}) = f_t(x_{i,t+1}^{N_1} | x_{i,t}^{N_1}, u_{i,t}^{N_1}, \mu_t^{N_1}, \nu_t^{N_2}), \quad (2)$$

where $\mu_t^{N_1} = \operatorname{Emp}_{\mu}(\mathbf{x}_t^{N_1})$ and $\nu_t^{N_2} = \operatorname{Emp}_{\nu}(\mathbf{y}_t^{N_2})$. Similarly, the dynamics of Red agent j is governed by the transition kernel $g_t: \mathcal{Y} \times \mathcal{V} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \to \mathcal{P}(\mathcal{Y})$. All agents in the Blue team receive an identical weakly-coupled team reward, i.e., $r_t \triangleq r_t(\mu_t, \nu_t) : \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \to \mathbb{R}$. The Red agents receive $-r_t(\mu_t, \nu_t)$ as their rewards (zero-sum). We assume that the Blue team is maximizing while the Red team is minimizing and $r_t \in [-R_{\max}, R_{\max}]$ for all t.

Assumption 1 (Lipschitz Model). For all $x \in \mathcal{X}, u \in \mathcal{U}, \mu, \mu' \in \mathcal{P}(\mathcal{X}), \nu, \nu' \in \mathcal{P}(\mathcal{Y})$ and all t, there exist constants $L_f, L_r > 0$ such that $\sum_{x' \in \mathcal{X}} |f_t(x'|x, u, \mu, \nu) - f_t(x'|x, u, \mu', \nu')| \leq L_f (\mathrm{d_{TV}}(\mu, \mu') + \mathrm{d_{TV}}(\nu, \nu'))$ and $|r_t(\mu, \nu) - r_t(\mu', \nu')| \leq L_r (\mathrm{d_{TV}}(\mu, \mu') + \mathrm{d_{TV}}(\nu, \nu'))$. A similar assumption also holds for g_t .

Lipschitz continuity is commonly assumed (Huang et al., 2006; Gu et al., 2021), and at minimum uniform continuity is required; see Cui et al. (2024) for counterexamples.

The first grid in Figure 1(a) depicts the individual agents' local positions, with the target marked by the green cell. The subsequent grids illustrate the state distributions $\mu_t^{N_1}$ and $\nu_t^{N_2}$ of both teams. The agent interactions depend only on $\mu_t^{N_1}$ and $\nu_t^{N_2}$ (weakly-coupled) as described in (2).

We consider a mean-field sharing information structure (Arabneydi & Mahajan, 2015), where each agent's decision depends on its own state and the two team EDs. We start with assuming full observation of mean-fields and later relax this assumption. Specifically, the Blue and Red agents seek to construct mixed Markov policies $\phi_{i,t}: \mathcal{U} \times \mathcal{X} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \to [0,1], \text{ and } \psi_{j,t}: \mathcal{V} \times \mathcal{Y} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \to [0,1], \text{ where the Blue policy } \phi_{i,t}(u|x_{i,t}^{N_1}, \mu_t^{N_1}, \nu_t^{N_2}) \text{ dictates the probability that Blue agent } i \text{ selects action } u \text{ given its state } x_{i,t}^{N_1} \text{ and the observed/estimated team EDs } \mu_t^{N_1} \text{ and } \nu_t^{N_2}. \text{ Note that each agent's individual state is its private information.}$

Let $\Phi_t \left(\Psi_t \right)$ denote the set of individual Blue (Red) policies at time t. We define the Blue team policy $\phi_t^{N_1} = \left\{ \phi_{i,t} \right\}_{i=1}^{N_1}$ as the collection of the N_1 Blue agent individual policies, and denote the set of Blue team policies as $\Phi_t^{N_1} = \times_{N_1} \Phi_t$. Similarly, the Red team policy is denoted as $\psi_t^{N_2} \in \Psi_t^{N_2} = \times_{N_2} \Psi_t$. **Definition 2** (Identical team policy). The Blue team policy $\phi_t^{N_1} = \left(\phi_{1,t}^{N_1}, \ldots, \phi_{N_1,t}^{N_1}\right)$ is identical, if $\phi_{i_1,t} = \phi_{i_2,t}$ for all times t and all $i_1, i_2 \in [N_1]$. Φ represents the set of identical Blue team policies.

The definition extends naturally to the Red team, and Ψ denotes the set of identical Red team policies. The expected cumulative reward defines the performance of the team policy pair (ϕ^{N_1}, ψ^{N_2}) :

$$J^{N,\phi^{N_1},\psi^{N_2}}\left(\mathbf{x}_0^{N_1},\mathbf{y}_0^{N_2}\right) = \mathbb{E}_{\phi^{N_1},\psi^{N_2}}\left[\sum_{t=0}^{T} r_t(\mathcal{M}_t^{N_1},\mathcal{N}_t^{N_2}) \middle| \mathbf{x}_0^{N_1},\mathbf{y}_0^{N_2}\right]. \tag{3}$$

When the Blue team considers its worst-case performance, we have the following max-min optimization problem:

$$\underline{J}^{N*}(\mathbf{x}_0^{N_1}, \mathbf{y}_0^{N_2}) = \max_{\phi^{N_1} \in \Phi^{N_1}} \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^{N_1}, \psi^{N_2}}(\mathbf{x}_0^{N_1}, \mathbf{y}_0^{N_2}), \tag{4}$$

where \underline{J}^{N*} is the lower game value for the *finite-population* game. Similarly, the minimizing Red team considers a min-max optimization problem, which leads to the upper game value. Note that we allow both teams to follow *non-identical* team policies in (4).

2.2 Infinite-Population Solution

To reduce the complexity of team policy optimization domains in (4), the authors of Guan et al. (2024a) examined team behaviors under *identical team policies* at the *infinite-population* limit. It was shown that the team joint states can be represented using the team population distribution, which coincides with the state distribution of a *typical agent* referred to as the mean-fields (μ_t and ν_t for the Blue and Red teams, respectively). They also proved that MFs induced by identical team policies in an infinite-population game closely approximate the EDs induced by *non-identical* team policies in the corresponding finite-population game, which justifies the simplification of the optimization domain in (4) to identical team policies (also see Theorem 1). Furthermore, there is a one-to-one correspondence between infinite-population coordination policies (α , β) and local *identical* team policies (α , β) and local *identical* team policies (α , β) in the equivalent zero-sum coordinator game is measured by

$$J_{\infty}^{\alpha,\beta}(\mu_0,\nu_0) \equiv J_{\infty}^{\phi,\psi}(\mu_0,\nu_0) = \sum_{t=0}^{T} r_t(\mu_t,\nu_t), \tag{5}$$

where μ_t and ν_t follow a *deterministic* dynamics (Guan et al., 2024a) similar to the state distribution propagation of a controlled Markov chain. The worst-case performance of the Blue team in this infinite-population game is then given by the lower game value $\underline{J}_{\infty}^*(\mu_0,\nu_0) = \max_{\phi \in \Phi} \min_{\psi \in \Psi} J_{\infty}^{\phi,\psi}(\mu_0,\nu_0)$, where the optimization domain is restricted to *identical team policies*. Guan et al. (2024a) establishes guarantees that identical team policies resulting from the solution of this equivalent zero-sum *coordinator* game are still ϵ -optimal for the original max-min optimization problem in (4) where $\epsilon = \mathcal{O}(1/\sqrt{N})$ and $N = \min\{N_1, N_2\}$.

The infinite-population limit of large-population games offers several theoretical advantages, such as representing the population by a typical agent and deterministic dynamics that reduce (3) to the non-stochastic optimization of (5). Previous works (Shao et al., 2024a; Perolat et al., 2021; Carmona et al., 2021) depend on infinite-population oracles to obtain mean-field trajectories (μ_t, ν_t) in order to compute (5). This is rather unrealistic, since in practice, only finite-population simulations and local states ($\mathbf{x}_t^{N_1}, \mathbf{y}_t^{N_2}$) with actions ($\mathbf{u}_t^{N_1}, \mathbf{v}_t^{N_2}$) are available/observable. Moreover, a single coordinator policy $\alpha(\beta)$ defines a distribution over actions for each state conditioned on the mean-field, causing its dimensionality to scale with the joint state–action space (e.g., DDPG-MFTG), leading to high computational cost and degraded empirical performance (see Section 5). In summary, the infinite-population model is both impractical (due to oracle dependency) and computationally intractable (due to policy size). Thus, we turn to finite-population simulators and derive guarantees of optimality, scalability, and convergence of the policy gradient to the infinite-population ZS-MFTG.

The next result quantifies the level of suboptimality for the Blue team when it deploys the optimal identical policy learned directly from the solution of *finite*-population ZS-MFTG.

Theorem 1. The value of the optimal identical Blue team policy ϕ^* obtained from the finite population game is within ϵ of the finite-population lower game value defined in (4). Formally, for all joint states \mathbf{x}^{N_1} and \mathbf{y}^{N_2} ,

$$\min_{\psi^{N_2}} J^{N,\phi^*,\psi^{N_2}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) \ge \underline{J}^{N*}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{N}}\right), \text{ where } \underline{N} = \min\{N_1,N_2\}.$$
 (6)

Theorem 1 provides a principled justification for learning identical finite-population team policies in competitive–collaborative team games even when being exploited by non-identical opponent team strategies. Its motivation and proof build on the performance guarantees of the ZS-MFTG in the infinite-population limit, i.e., the coordinator game. Moreover, the error vanishes as $N_1, N_2 \to \infty$, thereby recovering the well-studied infinite-population MF formulation (Huang et al., 2006). We detail this finite-population training paradigm in the next section.

3 Mean-Field Multi-Agent Proximal Policy Optimization

Motivated by Theorem 1, we present an algorithm to learn the *finite-population* optimal identical team policy. We build our algorithm based on the proximal policy optimization (PPO) framework due to its simplicity and effectiveness. While PPO has shown promising performance in cooperative tasks including MFC problems (Yu et al., 2022; Cui et al., 2024), its application in mixed competitive-collaborative scenarios is less studied. In the sequel, we introduce our key contribution: MF-MAPPO. We initialize two pairs of actor-critic networks, one for each team, deployed to learn the identical policy used by each team, see Figure 1(b). Specifically, we introduce a *minimally-informed critic* network by exploiting the MF information structure. The key point here is that we only require commonly accessible information for the critic network in order to learn the value function (Proposition 1). Further, the private information available to each agent only *individually* enters the actor during training. This results in neural networks that scale well with the number of agents. We present the team actor-critic networks from the Blue team's perspective, and due to symmetry results extend naturally to the Red team.

Minimally-Informed Critic. The MF-MAPPO critic network of the Blue team evaluates the value function $V_{\text{Blue}}(\mu,\nu)$, which depends only on the common information (MFs)—assumed to be available at the time of training—and is *independent* of the joint agent states and actions. We use the parameter vector ζ_{Blue} to parameterize the critic network while minimizing the MSE loss

$$L_{\text{critic}}(\zeta_{\text{Blue}}) = \frac{1}{|B|} \sum_{k=1}^{|B|} \left(V_{\text{Blue}}(\mu_k, \nu_k | \zeta_{\text{Blue}}) - \hat{R}_{\text{Blue},k} \right)^2, \tag{7}$$

where B refers to the mini-batch size and $R_{{\rm Blue},k}$ is the discounted reward-to-go for sample k. The following proposition results from weakly-coupled team rewards and the use of identical team policies-justifying the deployment of a minimally-informed critic network with only MF inputs.

Proposition 1. Let $\mu_t^{N_1}$, and $\nu_t^{N_2}$ denote the EDs of a finite-population game obtained from identical Blue and Red team policies $\phi_t \in \Phi_t$ and $\psi_t \in \Psi_t$, respectively. The team reward structure admits a critic that depends only on $\mu_t^{N_1}$ and $\nu_t^{N_2}$. Specifically, for each Blue team agent $i \in [N_1]$, the individual critic value function $V_{i,t}^{N_1,\phi_t}(x_{i,t},\mu_t^{N_1},\nu_t^{N_2})$ satisfies $V_{i,t}^{N_1,\phi_t}(x_{i,t},\mu_t^{N_1},\nu_t^{N_2}) = V_{\mathrm{Blue},t}^{N_1,\phi_t}(\mu_t^{N_1},\nu_t^{N_2})$, where $V_{\mathrm{Blue},t}^{N_1,\phi_t}(\mu_t,\nu_t)$ is the team-level critic.

Importantly, it reduces the learning problem to one critic network per team. Specifically, the shared team reward structure along with the assumption of homogeneous agents in each team enables us to evaluate the performance of a team's agent using the minimally-informed critic—even if the individual agent has additional local observations such as their actions and private states.

Shared-Team Actor. As discussed in earlier sections, the coordinator game is a useful theoretical construct but has limited practical value for real-world deployment since it relies on an infinite-population oracle for training and produces policies whose size scales poorly with the state-action space. We therefore directly train finite-population identical local policies, which preserve the mean-field structure while reducing complexity and improving tractability, with guarantees in Theorem 1. Not only is the approach computationally tractable in terms of the size of the policy, but is also more realistic in terms of sampling training data. We use a single actor network per team to learn identical

team policies. The actor optimizes a PPO-based objective with a decaying entropy bonus (Schulman et al., 2017; Huang et al., 2022), which promotes exploration and stabilizes learning in mean-field settings (Cui & Koeppl, 2022; Guan et al., 2022). Permutation invariance and identical team policies further allow a single replay buffer per team, reducing memory costs and simplifying experience collection. The PPO-based objective function of the Blue actor is given by:

$$L(\theta_{\text{Blue}}) = \frac{1}{|B|} \sum_{k=1}^{|B|} \left[\min \left(g_k(\theta_{\text{Blue}}) A_k, \text{clip}_{[1-\epsilon, 1+\epsilon]}(g_k(\theta_{\text{Blue}})) A_k \right) + \omega S(\phi_{\theta_{\text{Blue}}}(x_k, \mu_k, \nu_k)) \right], \quad (8)$$

where, $g(\theta) = \frac{\phi_{\theta}(u|x,\mu,\nu)}{\phi_{\theta^{\text{old}}}(u|x,\mu,\nu)}$, A_k is the generalized advantage function estimate (Schulman et al., 2018) and the tunable parameter ω weighs the contribution of the entropy term $S(\phi_{\theta}(x,\mu,\nu))$ and decays as training progresses.

3.1 THEORETICAL GUARANTEES

As described in Section 2, the theoretical benefits of MFTGs at the infinite-population limit remain of significant interest. Indeed, the following theorem shows that policy gradients obtained through finite-population training (using a finite-population simulator) converge to their infinite-population counterparts as the population size grows.

Theorem 2. The approximate policy gradient of the infinite-population Blue (Red) team coordinator policy α (β) computed using local policies from the finite-population ZS-MFTG via MF-MAPPO ($\hat{J}^{N_1}(\alpha_{\theta})$) uniformly tends to the true policy gradient as the population size increases, i.e., $\|\nabla_{\theta}J_{\infty}(\alpha_{\theta}) - \nabla_{\theta}\hat{J}^{N_1}(\alpha_{\theta})\|_2 \to 0$ as $(N_1, N_2) \to \infty$, where $\|\cdot\|_2$ is the 2-norm.

The results extend to the Red team. We next demonstrate the scalability of MF-MAPPO as a direct consequence of Theorem 1 and the infinite-population coordinator game, by showing that, under certain conditions, the learned team policies generalize to varying population sizes (\bar{N}_1, \bar{N}_2) while maintaining performance guarantees. Theorem 3 allows MF-MAPPO to be trained on a smaller population and deployed to larger teams without additional tuning, significantly reducing computational costs while maintaining performance consistency and generalizability across population sizes.

Theorem 3. Let \mathcal{G}_1 denote the finite-population game where the agents utilize the identical team policies ϕ_t^* and ψ_t^* derived from MF-MAPPO trained on \mathcal{G}_1 and let the finite-population game \mathcal{G}_2 with the same state-action space, dynamics, and rewards, but with population sizes \bar{N}_1 and \bar{N}_2 such that $\bar{N}_1/\bar{N}_2 = N_1/N_2$ and $\min(\bar{N}_1,\bar{N}_2) \geq \min(N_1,N_2)$. Then, (ϕ_t^*,ψ_t^*) remain ϵ -optimal for \mathcal{G}_2 .

4 MF-MAPPO FOR PARTIALLY OBSERVABLE MFTGS

To have strategies reactive to opponent's unexpected behaviors, one needs feedback on opponent's MF-distribution, which in practice, is often unavailable through direct means. We consider a partially observable ZS-MFTG, relevant to domains like competitive sports/battlefield, where decentralized decision-making relies on estimating the opponent's state distribution. The two main challenges are: 1) the sensitivity of MF policies (ϕ, ψ) to the MF (μ_t, ν_t) feedback and 2) constructing valid, performance-preserving MF estimates that can serve as inputs to the trained MF policies.

To address the first challenge, we introduce a gradient penalty to the MF-MAPPO objective (8), enforcing Lipschitz continuity in the mean-field and ensuring robustness: small estimation errors induce only minor changes in actions distributions. The following proposition formalizes this idea.

Proposition 2. If the log-probability of the Blue team policy is bounded such that for all $x \in \mathcal{X}, u \in \mathcal{U}, \mu \in \mathcal{P}(\mathcal{X}), \nu \in \mathcal{P}(\mathcal{Y}), \|\nabla_{\eta} \log \phi(u \mid x, \mu, \nu)\|_{2} \leq L_{\phi}/2|\mathcal{U}|$, where the gradient is taken with respect to $\eta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \triangleq [\mu^{\mathsf{T}}, \nu^{\mathsf{T}}]^{\mathsf{T}}$ and $L_{\phi} > 0$, then $\phi(u \mid x, \mu, \nu)$ is Lipschitz continuous with Lipschitz constant L_{ϕ} , i.e.,

$$\sum_{u} |\phi_{t}(u|x,\hat{\mu},\hat{\nu}) - \phi_{t}(u|x,\hat{\mu}',\hat{\nu}')| \leq L_{\phi} \left(d_{\text{TV}}(\hat{\mu},\hat{\mu}') + d_{\text{TV}}(\hat{\nu},\hat{\nu}') \right) \quad \forall \quad x \in \mathcal{X}.$$
 (9)

Similarly, we can define Lipschitz continuous policies for the Red team with constant L_{ψ} . This idea of penalizing the gradient of the policies was introduced in robotics to promote smooth and stable policies in order to aid sim-to-real transfer (Chen et al., 2024; Shin et al., 2025).

To address the second challenge, we require a filter that can estimate the opponent distribution at every time-step for each agent (e.g., $i \in [N_1]$ obtains an estimate of the Red team distribution at time t given by $\hat{\nu}_{i,t}^{N_2}$) with accuracy guaranteed within a bounded tolerance ensuring agent actions and overall performance (3) remain within acceptable limits. Note that we formulate the problem from the Blue team's perspective. The results extend naturally to the Red team's perspective. Let the full-information and estimated MF trajectories be $\{\mathcal{M}^{N_1}, \mathcal{N}^{N_2}\}$ and $\{\hat{\mathcal{M}}^{N_1}, \hat{\mathcal{N}}^{N_2}\}$, respectively. We measure estimator performance for gradient-regularized MF-MAPPO (GR-MF-MAPPO) team policies (ϕ_t^*, ψ_t^*) via the cumulative regret between fully and partially observable MF rewards as:

$$\Delta J(\phi_t^*, \psi_t^*) = \mathbb{E}_{\phi^*, \psi^*} \left[\left| \sum_{t=0}^T r_t(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}) - \sum_{t=0}^T r_t(\hat{\mathcal{M}}_t^{N_1}, \hat{\mathcal{N}}_t^{N_2}) \right| \right]. \tag{10}$$

In fact, any ϵ -accurate estimator can be utilized during the deployment of GR-MF-MAPPO.

Proposition 3. Consider a given ϵ -accurate estimator, i.e., $d_{TV}(\hat{\nu}_{i,t}^{N_2}, \hat{\nu}_{t}^{N_2}) < \epsilon$, for all i, t, where $\hat{\nu}_{t}^{N_2}$ is the true opponent MF at time t and $\epsilon > 0$. For the identical team-policy pair (ϕ_t^*, ψ_t^*) obtained via gradient-regularized MF-MAPPO and deployed using this estimator, the cumulative regret satisfies $\Delta J(\phi_t^*, \psi_t^*) \leq K\epsilon + \mathcal{O}(1/\sqrt{N})$ for some constant K > 0.

We emphasize that MF-MAPPO is modular and policy inputs can be swapped with different estimates (using estimation/prediction algorithms) and still have good performance. Gradient regularization is key to ensure that minor errors in estimation do not result in extreme changes in MF trajectories and performance.

In lieu of Proposition 3, we propose a communication network-based decentralized estimation filter, namely, Dynamic-Projected Consensus (D-PC). It is an extension of the control-theoretic constrained consensus problem (Nedić & Liu, 2016) and addresses the shortcomings of the estimation algorithm proposed in Benjamin & Abate (2025b), namely, estimation under limited communication rounds, and ensuring valid estimates in the presence of errors. Following the connected graph topology used in MFGs (Benjamin & Abate, 2025b), we define a state-based visibility graph $\mathcal{G}_t^{\text{viz}}$ and team-based communication graphs $\mathcal{G}^{\text{com}}_{\text{Blue},t}$ and $\mathcal{G}^{\text{com}}_{\text{Red},t}$. We also define the projection operator $\Omega_{\mathcal{R}(x)}[\eta] \triangleq \arg\min_{\omega \in \mathcal{R}(x)} \|\eta - \omega\|_2 \text{ for } \eta, \omega \in \mathbb{R}^{|\mathcal{Y}|} \text{ where } \mathcal{R}(x) \text{ is a closed and convex set. We}$ assume that all agents in the same state receive the same information, so naturally they have the same estimate, i.e., all Blue agents at state $x \in \mathcal{X}$ at time t have the same estimate $\hat{\nu}_{x,t}^{N_2}$ of the Red team MF. We consider two time scales: t for system dynamics (2) and τ for communication rounds. At time t, and $\tau=0$, each Blue agent at state $x\in\mathcal{X}$ holds a belief $\hat{\nu}_{x,t}^{\tau=0}$ consistent with $\mathcal{G}_t^{\mathrm{viz}}$. At communication round au, agents share estimates $\hat{\nu}_{x,t}^{ au-1}$ with neighbors defined by $\mathcal{G}_{\mathrm{Blue},t}^{\mathrm{com}}$ and perform for R_{com} communication rounds: a) weighted-average consensus; and b) a projection onto a closed and convex constraint set $\mathcal{R}(x)$. The set $\mathcal{R}(x)$ combines the Red team's MF components known with certainty by the Blue agents at state x—i.e., the observable states given by $\mathcal{G}_t^{\text{viz}}$ —with those that must be estimated. $\mathcal{R}(x)$ guarantees that operations such as information aggregation, averaging, or distributed communication do not alter the parts of the distribution that are known with certainty. See Appendix A for detailed definitions.

Theorem 4. *D-PC satisfies Proposition 3 with* $\epsilon = \mathcal{O}(e^{-cR_{\text{com}}})$ *with* c > 0.

5 NUMERICAL EXPERIMENTS

In this section, we evaluate MF-MAPPO across large-population scenarios using our custom-made benchmark simulation platform, MFEnv. Built as an extension of Gymnasium (Towers et al., 2024), MFEnv is developed specifically to facilitate research in MFTGs, supporting both finite-agent simulations and oracle-based infinite-population models. Unlike existing toolkits, MFEnv includes aggregate reward metrics, policy-versus-policy evaluation, and flexible APIs for custom *mean-field environments* that adhere to MF dynamics, rewards and information structures. We showcase MF-MAPPO's efficacy on two representative environments (1) a constrained-action variant of the classical rock-paper-scissors game (Raghavan, 1994), enabling validation against analytically computed equilibria and (2) a complex battlefield setting where Blue and Red teams engage in attack-defense tasks with higher-dimensional state and action spaces, requiring sophisticated team-level coordination. Additional results and environments (not limited to ZS-MFTGs) in Appendix E.

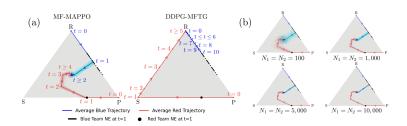


Figure 2: (a) 150 initializations of $\mu_{t=0} = [1,0,0]^T$ and $\nu_{t=0} = [0,1,0]^T$ for cRPS; $N_1 = N_2 = 1,000$ (b) Deploying MF-MAPPO trained on $N_1 = N_2 = 1,000$ to varying team sizes.

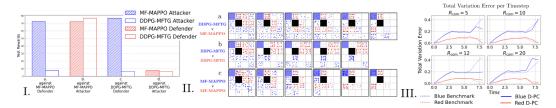


Figure 3: I. Average test rewards for Battlefield on a 4x4 grid (100 initializations); II. Example configuration; III. Comparing $d_{TV}(\cdot)$ for D-PC and Benchmark estimator for different R_{com} .

Constrained Rock-Paper-Scissors (cRPS). The state space of each individual agent is $\mathcal{S} = \{\mathbb{R}, \mathbb{P}, \mathbb{S}\}$, representing rock, paper, and scissors, respectively. We consider a non-trivial restriction of the action space to $\mathcal{A} = \{\mathbb{CW}, \mathbb{Stay}\}$ allowing agents to either move clockwise ($\mathbb{R} \to \mathbb{P}, \mathbb{P} \to \mathbb{S}, \mathbb{S} \to \mathbb{R}$) or remain idle, respectively. We assume deterministic transitions, where each action leads to a unique next state *deterministically*. At each time step t, the Blue (Red) team receives a team reward $r_t(\mu_t, \nu_t) = \mu_t^T A \nu_t (-\mu_t^T A \nu_t)$ where A is the standard RPS payoff matrix.

Figure 2(a) compares trajectories from MF-MAPPO and DDPG-MFTG (Shao et al., 2024a). Existing algorithms such as MADDPG and MAAC (Lowe et al., 2017) are unsuitable as baselines because their centralized action-value functions require the actions of all agents and the full state information, leading to poor scalability in the large-population settings considered in this work. We see that MF-MAPPO successfully reaches the equilibrium distribution. The transient time is attributed to finite-population approximation and entropy regularization. By contrast, DDPG-MFTG diverges, relying on a mean-field oracle, which is valid only in the infinite-population limit, and "central players" that map mean-field distributions to deterministic policies without clipping or regularization, making it unstable. Unlike multi-agent DDPG extensions (e.g., MADDPG), which consider other teams' local policies, DDPG-MFTG conditions only on its own, limiting inter-team awareness. Figure 2(b) shows MF-MAPPO's scalability, where larger populations reduce noise and variance, aligning with Theorem 3. MF-MAPPO updates every $T_{\rm rollout}$ steps, unlike DDPG-MFTG which updates per time step, making MF-MAPPO significantly faster and efficient (2h vs. 60h).

Battlefield Game. To fully test the capability of MF-MAPPO on a more complex scenario, we propose a grid-based battlefield game where an individual agent's dynamics is highly coupled with both teams' distributions. The Blue agents aim is to reach their targets without being deactivated, while the Red agents learn to guard them. Deactivation occurs when one of the opponents holds a numerical advantage within a cell, incentivizing both teams to aggregate to reduce the risk of being deactivated by a numerically superior opponent. The Blue team's reward depends on the fraction of agents active at the target, the Red team's reward follows from the zero-sum structure, and (to avoid degeneracy) the Red agents are restricted from entering the target. All experiments use $N_1 = N_2 = 100$ agents on grids with varied targets (lilac) and obstacles (black).

We compare MF-MAPPO and DDPG-MFTG by pitting them against each other on a 4×4 grid with full MF information. As shown in Figure 3I, MF-MAPPO consistently outperforms DDPG-MFTG across various initializations, achieving up to $10 \times$ higher rewards when attacking and lower/comparable rewards while defending. Figure 3II shows MF-MAPPO Red agents successfully cover both corridors and deactivate several Blue attackers II(a). Panels II(a) and II(b) highlight that DDPG-MFTG Blue agents do not aggressively pursue the target, illustrating their tendency

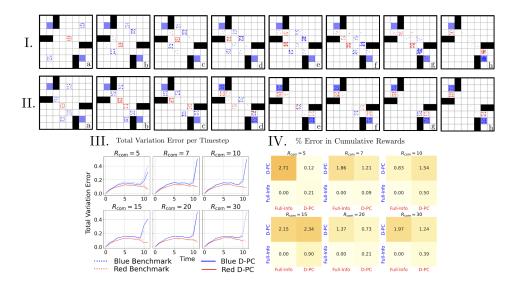


Figure 4: I. Red is concentrated; 30% Blue are at the bottom, rest are at the top II. Blue is evenly split, Red is concentrated III. Comparing $d_{TV}(\,\cdot\,)$ for D-PC and the Benchmark estimator for different communication rounds IV. % error in cumulative rewards: D-PC vs. fully observable ZS-MFTG.

to passively seek zero-reward outcomes rather than take goal-directed actions, unlike II(c) where MF-MAPPO agents exhibit coordinated maneuvering, forming coalitions to reach the target. While cases I and II utilize complete observability, case III evaluates the proposed D-PC estimator against the estimator in Benjamin & Abate (2025a) (Benchmark) under a gradient-penalized MF-MAPPO policy when Red has full information and Blue estimates Red's distribution. Both estimators yield comparable total variation errors relative to the full-information case, with D-PC showing advantages under low communication budgets ($R_{\rm com} < 20$). One can trivially show that the Benchmark satisfies Proposition 3 with $\epsilon = 1 - \mathcal{O}(R_{\rm com}/|\mathcal{X}|)$ and assumes uniform estimates for unobserved states, which degrades estimation accuracy under limited communication. It also relies on accurate information from neighbors to ensure validity of estimates. In contrast, D-PC exchanges inexact information but applies state-dependent corrections (projection), preserving privacy and robustness.

In Figure 4, the Red team faces a dilemma in deciding which target to defend, while the Blue team exploits this ambiguity. Due to DDPG-MFTG's high computational cost and frequent network updates, it is excluded from our analysis. With no other baselines for such large-scale complex MFTGs, we qualitatively assess MF-MAPPO's performance. Figures 4I–II illustrate how identical policies can generate heterogeneous team behaviors, with Blue adapting target selection and Red reallocating defenses, highlighting the flexibility of the mean-field approximation. Furthermore, D-PC again outperforms the Benchmark under limited communication (III) ($R_{\rm com} < 10$) and performs competitively otherwise. Cumulative rewards (IV) show only small deviations, consistent with Proposition 3 and Theorem 4, confirming that agents can rely on local observations with minimal communication.

6 Conclusion

We introduced MF-MAPPO, a novel MARL algorithm for large-population competitive team games that leverages finite mean-field approximation. With a minimally informed critic and shared team actor, MF-MAPPO scales efficiently while retaining performance, as shown against baselines such as DDPG-MFTG on cRPS and battlefield scenarios using the developed platform MFEnv. Despite shared policies, heterogeneous sub-population behaviors emerge, confirming that mean-field approximations do not hinder performance. We showed that MF-MAPPO naturally extends to partial observability via a simple gradient-regularized training scheme, and proposed D-PC, a decentralized mean-field estimator that ensures accuracy and strong performance when integrated with it. Empirically, D-PC outperforms baselines under limited communication. Limitations include MF-MAPPO's scaling with state dimensionality, which motivates future work on dimensionality reduction (e.g., kernel embeddings). Additional directions include extending D-PC to noisy settings and to more general, time-varying network topologies.

REFERENCES

- Jalal Arabneydi and Aditya Mahajan. Team optimal control of coupled subsystems with mean-field sharing. In *53rd IEEE Conference on Decision and Control*, pp. 1669–1674. IEEE, 2014.
- Jalal Arabneydi and Aditya Mahajan. Team-optimal solution of finite number of mean-field coupled LQG subsystems. In *54th IEEE Conference on Decision and Control*, pp. 5308–5313, Osaka, Japan, Dec. 15–18, 2015, 2015. IEEE.
- Patrick Benjamin and Alessandro Abate. Networked communication for decentralised cooperative agents in mean-field control. *arXiv preprint arXiv:2503.09400*, 2025a.
- Patrick Benjamin and Alessandro Abate. Networked communication for mean-field games with function approximation and empirical mean-field estimation, 2025b. URL https://arxiv.org/abs/2408.11607.
- Alain Bensoussan, Jens Frehse, Phillip Yam, et al. *Mean field games and mean field type control theory*, volume 101. Springer, 2013.
- Daniel S Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of markov decision processes. *Mathematics of operations research*, 27(4): 819–840, 2002.
- René Carmona, Mathieu Laurière, and Zongjun Tan. Model-free mean-field reinforcement learning: Mean-field mdp and mean-field q-learning, 2021. URL https://arxiv.org/abs/1910.12802.
- Zixuan Chen, Xialin He, Yen-Jen Wang, Qiayuan Liao, Yanjie Ze, Zhongyu Li, S Shankar Sastry, Jiajun Wu, Koushil Sreenath, Saurabh Gupta, et al. Learning smooth humanoid locomotion through lipschitz-constrained policies. *arXiv preprint arXiv:2410.11825*, 2024.
- Kai Cui and Heinz Koeppl. Approximately solving mean field games via entropy-regularized deep reinforcement learning, 2022. URL https://arxiv.org/abs/2102.01585.
- Kai Cui, Christian Fabian, Anam Tahir, and Heinz Koeppl. Major-minor mean field multi-agent reinforcement learning. *arXiv preprint arXiv:2303.10665*, 2023.
- Kai Cui, Sascha Hauck, Christian Fabian, and Heinz Koeppl. Learning decentralized partially observable mean field control for artificial collective behavior, 2024. URL https://arxiv.org/abs/2307.06175.
- Samuel Jose Octavio Gaviria Alzate, Wilder Valencia Sánchez, and Elkin Arias Arias. A socio-constructivist framework for tactical development in team sports: fostering critical thinking through collaborative learning. *Frontiers in Psychology*, 16:1610750, 2025.
- Haotian Gu, Xin Guo, Xiaoli Wei, and Renyuan Xu. Mean-field controls with q-learning for cooperative marl: Convergence and complexity analysis, 2021. URL https://arxiv.org/abs/2002.04131.
- Yue Guan, Mi Zhou, Ali Pakniyat, and Panagiotis Tsiotras. Shaping large population agent behaviors through entropy-regularized mean-field games. In *2022 American Control Conference (ACC)*, pp. 4429–4435, Atlanta, USA, June 08–10, 2022, 2022. IEEE, IEEE.
- Yue Guan, Mohammad Afshari, and Panagiotis Tsiotras. Zero-sum games between mean-field teams: Reachability-based analysis under mean-field sharing. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(9):9731–9739, Mar. 2024a. doi: 10.1609/aaai.v38i9.28831. URL https://ojs.aaai.org/index.php/AAAI/article/view/28831.
- Yue Guan, Mohammad Afshari, and Panagiotis Tsiotras. Zero-sum games between large-population teams: Reachability-based analysis under mean-field sharing, 2024b. URL https://arxiv.org/abs/2303.12243.
 - Xin Guo, Anran Hu, Matteo Santamaria, Mahan Tajrobehkar, and Junzi Zhang. Mfglib: A library for mean-field games. *arXiv preprint arXiv:2304.08630*, 2023.

- Minyi Huang, Roland P. Malhamé, and Peter E. Caines. Large population stochastic dynamic games:
 closed-loop McKean-Vlasov systems and the Nash certainty equivalence principle. *Communications in Information & Systems*, 6(3):221 252, 2006.
 - Shengyi Huang, Rousslan Fernand Julien Dossa, Antonin Raffin, Anssi Kanervisto, and Weixun Wang. The 37 implementation details of proximal policy optimization, 2022. URL https://iclr-blog-track.github.io/2022/03/25/ppo-implementation-details/. https://iclr-blog-track.github.io/2022/03/25/ppo-implementation-details/.
 - Daisuke Inoue, Yuji Ito, Takahito Kashiwabara, Norikazu Saito, and Hiroaki Yoshida. Model predictive mean field games for controlling multi-agent systems. In 2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 982–987. IEEE, 2021.
 - Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL https://arxiv.org/abs/1412.6980.
 - Marc Lanctot, Vinicius Zambaldi, Audrūnas Gruslys, Angeliki Lazaridou, Karl Tuyls, Julien Pérolat, David Silver, and Thore Graepel. A unified game-theoretic approach to multiagent reinforcement learning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 4193–4206, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
 - Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2 (1):229–260, 2007.
 - Mathieu Laurière, Sarah Perrin, Sertan Girgin, Paul Muller, Ayush Jain, Theophile Cabannes, Georgios Piliouras, Julien Pérolat, Romuald Élie, Olivier Pietquin, and Matthieu Geist. Scalable deep reinforcement learning algorithms for mean field games, 2022. URL https://arxiv.org/abs/2203.11973.
 - Joel Z Leibo, Vinicius Zambaldi, Marc Lanctot, Janusz Marecki, and Thore Graepel. Multi-agent reinforcement learning in sequential social dilemmas. *arXiv preprint arXiv:1702.03037*, 2017.
 - Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning, 2019. URL https://arxiv.org/abs/1509.02971.
 - Ryan Lowe, Yi Wu, Aviv Tamar, Jean Harb, Pieter Abbeel, and Igor Mordatch. Multi-agent actor-critic for mixed cooperative-competitive environments. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17, pp. 6382–6393, Red Hook, NY, USA, 2017. Curran Associates Inc. ISBN 9781510860964.
 - Aditya Mahajan and Ashutosh Nayyar. Sufficient statistics for linear control strategies in decentralized systems with partial history sharing. *IEEE Transactions on Automatic Control*, 60(8): 2046–2056, 2015. doi: 10.1109/TAC.2015.2398884.
 - Ashutosh Nayyar, Aditya Mahajan, and Demosthenis Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Transactions on Automatic Control*, 58(7):1644–1658, 2013.
 - Angelia Nedić and Ji Liu. On convergence rate of weighted-averaging dynamics for consensus problems. *IEEE Transactions on Automatic Control*, 62(2):766–781, 2016.
 - Julien Perolat, Sarah Perrin, Romuald Elie, Mathieu Laurière, Georgios Piliouras, Matthieu Geist, Karl Tuyls, and Olivier Pietquin. Scaling up mean field games with online mirror descent, 2021. URL https://arxiv.org/abs/2103.00623.
 - Sarah Perrin, Mathieu Laurière, Julien Pérolat, Matthieu Geist, Romuald Élie, and Olivier Pietquin. Mean field games flock! the reinforcement learning way, 2021. URL https://arxiv.org/abs/2105.07933.

- Eduardo Pignatelli, Johan Ferret, Matthieu Geist, Thomas Mesnard, Hado van Hasselt, Olivier Pietquin, and Laura Toni. A survey of temporal credit assignment in deep reinforcement learning, 2024. URL https://arxiv.org/abs/2312.01072.
 - TES Raghavan. Zero-sum two-person games. *Handbook of game theory with economic applications*, 2:735–768, 1994.
 - Pam Richards, Dave Collins, and Duncan RD Mascarenhas. Developing rapid high-pressure team decision-making skills. the integration of slow deliberate reflective learning within the competitive performance environment: A case study of elite netball. *Reflective Practice*, 13(3):407–424, 2012.
 - Sina Sanjari, Naci Saldi, and Serdar Yüksel. Nash equilibria for exchangeable team against team games, their mean field limit, and role of common randomness. *arXiv preprint arXiv:2210.07339*, 2022.
 - John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms, 2017. URL https://arxiv.org/abs/1707.06347.
 - John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation, 2018. URL https://arxiv.org/abs/1506.02438.
 - Nevroz Sen and Peter E Caines. Mean field games with partial observation. *SIAM Journal on Control and Optimization*, 57(3):2064–2091, 2019.
 - Kai Shao, Jiacheng Shen, Chijie An, and Mathieu Laurière. Reinforcement learning for finite space mean-field type games. *arXiv preprint arXiv:2409.18152*, 2024a.
 - Kai Shao, Jiacheng Shen, Chijie An, and Mathieu Laurière. Reinforcement learning for finite space mean-field type games, 2024b. URL https://arxiv.org/abs/2409.18152.
 - Jaeyong Shin, Woohyun Cha, Donghyeon Kim, Junhyeok Cha, and Jaeheung Park. Spectral normalization for lipschitz-constrained policies on learning humanoid locomotion. *arXiv* preprint *arXiv*:2504.08246, 2025.
 - Max Olan Smith, Thomas Anthony, and Michael P. Wellman. Iterative empirical game solving via single policy best response, 2021. URL https://arxiv.org/abs/2106.01901.
 - Sylvain Sorin. A First Course on Zero-Sum Repeated Games, volume 37. Springer Science & Business Media, Berlin, 2002.
 - J Terry, Benjamin Black, Nathaniel Grammel, Mario Jayakumar, Ananth Hari, Ryan Sullivan, Luis S Santos, Clemens Dieffendahl, Caroline Horsch, Rodrigo Perez-Vicente, et al. Pettingzoo: Gym for multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 34: 15032–15043, 2021.
 - Mark Towers, Ariel Kwiatkowski, Jordan Terry, John U Balis, Gianluca De Cola, Tristan Deleu, Manuel Goulão, Andreas Kallinteris, Markus Krimmel, Arjun KG, et al. Gymnasium: A standard interface for reinforcement learning environments. *arXiv preprint arXiv:2407.17032*, 2024.
 - Guofang Wang, Ziming Li, Wang Yao, and Sikai Xia. A multi-population mean-field game approach for large-scale agents cooperative attack-defense evolution in high-dimensional environments. *Mathematics*, 10(21):4075, 2022.
- Lin Xiao, Stephen Boyd, and Sanjay Lall. Distributed average consensus with time-varying metropolis weights. *Automatica*, 1:1–4, 2006.
- Yaodong Yang, Rui Luo, Minne Li, Ming Zhou, Weinan Zhang, and Jun Wang. Mean field multiagent reinforcement learning, 2020. URL https://arxiv.org/abs/1802.05438.
 - Batuhan Yardim and Niao He. Exploiting approximate symmetry for efficient multi-agent reinforcement learning, 2024. URL https://arxiv.org/abs/2408.15173.

Chao Yu, Akash Velu, Eugene Vinitsky, Jiaxuan Gao, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of ppo in cooperative multi-agent games. *Advances in Neural Information Processing Systems*, 35:24611–24624, 2022.

Lianmin Zheng, Jiacheng Yang, Han Cai, Ming Zhou, Weinan Zhang, Jun Wang, and Yong Yu. Magent: A many-agent reinforcement learning platform for artificial collective intelligence. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018.

Algorithm 1 Mean-Field Multi-Agent Proximal Policy Optimization (MF-MAPPO)

```
 \begin{array}{l} \textbf{Initialize: NN parameters } \{\theta_{\text{Blue}}, \zeta_{\text{Blue}}\} \text{ and } \{\theta_{\text{Red}}, \zeta_{\text{Red}}\}; \text{ step size sequences } \{\alpha_m\} \text{ and } \{\beta_m\}; \\ \textbf{for } i=1,2,\ldots \textbf{do} \\ & (\phi_{\theta_{\text{Blue}}}^{\text{old}}, \psi_{\theta_{\text{Red}}}^{\text{old}}) \leftarrow (\phi_{\theta_{\text{Blue}}}, \psi_{\theta_{\text{Red}}}) \\ \textbf{for } t=0,1,\ldots,T_{\text{rollout}} \textbf{do} \\ & \text{Sample joint actions} \\ & u_{i,t} \sim \phi_{\theta_{\text{Blue}}^{\text{old}}}(x_{i,t}, \mu_t^{N_1}, \nu_t^{N_2}), v_{j,t} \sim \psi_{\theta_{\text{Red}}^{\text{old}}}(y_{j,t}, \mu_t^{N_1}, \nu_t^{N_2}) \\ & \text{Step environment according to kernels } (f_t, g_t) \\ & \text{Collect samples } (\mathbf{x}_{t+1}^{N_1}, \mathbf{y}_{t+1}^{N_2}, \mu_{t+1}^{N_1}, \nu_{t+1}^{N_2}, \mathbf{u}_t^{N_1}, \mathbf{v}_t^{N_2}, r_t) \\ & \textbf{end for} \\ & \textbf{for } K \text{ epochs } \textbf{do} \\ & \text{Update } \{\theta_{\text{Blue}}, \zeta_{\text{Blue}}\} \text{ and } \{\theta_{\text{Red}}, \zeta_{\text{Red}}\} \text{ using } (7\text{-8}) \\ & \textbf{end for} \\ & \textbf{Return: } (\phi_{\theta_{\text{Blue}}}, \psi_{\theta_{\text{Red}}}) \\ \end{array}
```

A DETAILS OF D-PC

Definition 3 (Visibility Graph). The bipartite state-based visibility graph at time t, denoted by $\mathcal{G}_t^{\mathrm{viz}} = (\mathcal{X}, \mathcal{Y}, \mathcal{E}_t^{\mathrm{viz}})$, is an undirected graph where (i) The vertex sets consist of the states $x \in \mathcal{X}$ and $y \in \mathcal{Y}$; (ii) An undirected edge $(x, y) \in \mathcal{E}_t^{\mathrm{viz}}$ exists if and only if all Blue agents at state x are mutually visible with all Red agents at state y at time t; that is, every Blue agent at state x is visible to every Red agent at state y at time t.

Before describing the communication architecture, we define $\mathcal{X}_t^o \subseteq \mathcal{X}$ such that $\hat{\mu}_t(x) > 0$ if and only if $x \in \mathcal{X}_t^o$.

Definition 4 (Communication Graph). The state-based communication graph at time t, denoted by $\mathcal{G}^{\text{com}}_{\text{Blue},t} = (\mathcal{X}^o_t, \mathcal{E}^{\text{com}}_t)$, is an undirected graph where (i) The vertex set \mathcal{X}^o_t consists of the states $x \in \mathcal{X}^o_t$; (ii) An undirected edge $(x_p, x_q) \in \mathcal{E}^{\text{com}}_t$ exists if and only if all agents at state x_p can communicate with all agents at state x_q at time t.

We denote the set of neighbors of vertex v of $\mathcal{G}_t^{\mathrm{viz}}$ and $\mathcal{G}_{\mathrm{Blue},t}^{\mathrm{com}}$ respectively, as $\mathcal{N}_t^{\mathrm{viz}}(v)$ and $\mathcal{N}_{\mathrm{Blue},t}^{\mathrm{com}}(v)$. $\mathcal{Y}_t^o \subseteq \mathcal{Y}$, $\mathcal{G}_{\mathrm{Red},t}^{\mathrm{com}}$ and $\mathcal{N}_{\mathrm{Red},t}^{\mathrm{com}}(v)$ are defined similarly. The following assumption on the topology of the communication graph will simplify the mean-field estimator design problem.

Assumption 2. We assume that $\mathcal{G}^{com}_{\mathrm{Blue},t}$ and $\mathcal{G}^{com}_{\mathrm{Red},t}$ are connected for all t.

At time t, we assume that each agent $i \in [N]$ observes the population distribution of its opponent's visible states $y \in \mathcal{N}_t^{\mathrm{viz}}(x_{i,t})$. Thus, from the perspective of agent i, the ED $\hat{\nu}_{i,t}^{N_2}$ is defined on the constrained simplex $\hat{\nu}_{i,t}^{N_2} \in \mathcal{R}(x_{i,t}) \triangleq \{ \eta \in \mathcal{P}(\mathcal{Y}) \mid \eta(z) = \hat{\nu}_t^{N_2}(z), \ \forall \ z \in \mathcal{N}_t^{\mathrm{viz}}(x_{i,t}) \}$. given by $\mathcal{R}(x_{i,t})$ combines the Red team's mean-field components known with certainty by Blue agent i—i.e., the observable states $\mathcal{N}_t^{\mathrm{viz}}(x_{i,t})$ —with those that must be estimated.

Proposition 4. For all $x \in \mathcal{X}$, $\mathcal{R}(x)$ is a closed and convex set.

It can be easily seen that the true mean-field distribution is an element of this constrained simplex for each agent, i.e., $\hat{\nu}_t^{N_2} \in \mathcal{R}(x_{i,t})$ for all $i \in [N_1]$ and time t. The following assumption imposes uniformity of the estimate of the ED across agents in a given state.

Assumption 3. We assume that $\hat{\nu}_{i,t}^{N_2} = \hat{\nu}_{j,t}^{N_2} \triangleq \hat{\nu}_{x,t}^{N_2}$ for all agents satisfying $x_{i,t}^{N_1} = x_{j,t}^{N_1} = x$.

Definition 5. Let $x \in \mathcal{X}$. The projection operator $\Omega_{\mathcal{R}(x)} : \mathbb{R}^{|\mathcal{Y}|} \to \mathcal{R}(x)$ where $\mathcal{R}(x)$ is a closed and convex set is defined as $\Omega_{\mathcal{R}(x)}[\eta] = \omega^* \triangleq \arg\min_{\omega \in \mathcal{R}(x)} \|\eta - \omega\|_2$, $\eta \in \mathbb{R}^{|\mathcal{Y}|}$. where $\|\cdot\|_2$ is the standard Euclidean norm.

At each communication round, D-PC performs the following two steps:

1. Weighted average consensus:

758 759 760

761

762 763

764 765 766

767

768

769 770

771

772

773 774

775

776 777

778

779

781 782

783 784

785

786

787

788

789

795

796

797

798

799

800

801

802

803

804

805 806

807 808

809

$$\xi_{x,t}^{\tau} = \sum_{z \in \mathcal{N}_{\mathrm{Blue},t}^{\mathrm{com}}(x) \cup x} w_{(x,z)}^{\mathrm{Blue}} \hat{\nu}_{z,t}^{\tau-1}, \quad \sum_{z} w_{(x,z)}^{\mathrm{Blue}} = 1, \quad x \in \mathcal{X}^{o}. \tag{A.1}$$

2. Projection onto the constraint set:

$$\hat{\nu}_{x,t}^{\tau} = \Omega_{\mathcal{R}(x)} \left[\xi_{x,t}^{\tau} \right], \ x \in \mathcal{X}^{o}. \tag{A.2}$$

 $\hat{\nu}_{x,t}^{\tau} = \Omega_{\mathcal{R}(x)} \left[\xi_{x,t}^{\tau} \right], \ x \in \mathcal{X}^{o}. \tag{A.2}$ **Assumption 4.** The matrix $W^{\text{Blue}} \in \mathbb{R}^{|\mathcal{X}^{o}| \times |\mathcal{X}^{o}|}$ formed by the non-negative weights is symmetric and doubly-stochastic. Furthermore, it respects the sparsity structure of the communication graph \mathcal{G}^{com} , that is, $w_{(x,z)}^{\text{Blue}} > 0$ if and only if $z \in \mathcal{N}^{\text{com}}(x) \cup x$. W^{Red} is constructed similarly.

To satisfy Assumption 4, one may use the well-known Metropolis Matrix (Xiao et al., 2006). We define the smallest non-zero entry of α as

$$\theta \triangleq \min_{(x,y)} \{ w_{(x,y)} > 0 \}. \tag{A.3}$$

For clarity of the D-PC pseudocode present in Algorithm A, the estimation is presented from the Blue team's perspective which models the distribution of its opponent Red team. One can simultaneously run this estimation algorithm from the Red team's perspective.

Algorithm 2 Dynamic Projected Consensus for Mean-Field Estimation (**D-PC**)

```
Initialize: \hat{\mu}_{t=0}^{N_1} = \mu_{t=0}^{N_1} and \hat{\nu}_{t=0}^{N_2} = \nu_{t=0}^{N_2}, identical MF-MAPPO trained Lipschitz policies (\phi^*, \psi^*) and graphs \mathcal{G}_{t=0}^{\mathrm{viz}} and (\mathcal{G}_{\mathrm{Blue},t}^{\mathrm{com}}, \mathcal{G}_{\mathrm{Red},t}^{\mathrm{com}})
for t = 0, \dots, T do
       Update the graphs \mathcal{G}_t^{\mathrm{viz}} and (\mathcal{G}_{\mathrm{Blue},t}^{\mathrm{com}},\mathcal{G}_{\mathrm{Red},t}^{\mathrm{com}}).
        Receive reward r_t(\hat{\mu}_t^{N_1}, \hat{\nu}_t^{N_2})
       Set \tau = 0
Initialize \hat{\nu}_{x,t}^{N_2,\tau=0} (\hat{\mu}_{y,t}^{N_1,\tau=0}) for all x \in \mathcal{X}_t^o (y \in \mathcal{Y}_t^o) using \mathcal{G}_{t=0}^{\mathrm{viz}}
        Compute matrix W^{\text{Blue}}(W^{\text{Red}})
        for \tau < R_{\rm com} do
                Communicate mean-field estimate to neighbors based on \mathcal{G}_{\text{Blue }t}^{\text{com}}(\mathcal{G}_{\text{Red }t}^{\text{com}})
                \tau = \tau + 1
                for all x \in \mathcal{X}_t^o (y \in \mathcal{Y}_t^o) do
                        Compute weighted average consensus (A.1)
                        Project result onto constraint set (A.2)
                end for
        end for
        Return: \hat{\nu}_{x,t}^{N_2,R_{\mathrm{com}}}(\hat{\mu}_{y,t}^{N_1,R_{\mathrm{com}}})
```

Remark 1. One could alternatively cast this problem as a Partially Observable Markov Decision Process (PO-MDP) (Bernstein et al., 2002), where the environment provides arbitrary observations. However, our focus is on policies that explicitly depend on the mean-field distributions of the two teams, for two key reasons (1) Distinct opponent strategies induce different mean-field trajectories, and closed-loop feedback on these distributions enables an appropriate response to the strategies deployed by the opponent (Guan et al., 2024a) and (2) Access to team-level distributions facilitates credit assignment (Pignatelli et al., 2024) in MARL by allowing agents to reason about which collective distributions are optimal. Thus, having access to MF information—particularly that of the opponent—is desirable. Yet, it is unrealistic to assume that the environment directly provides these distributions, which motivates the design of estimation algorithms. Finally, unlike general PO-MDP formulations that often require maintaining long observation histories, our setting involves population sizes for which such history-based tracking is computationally infeasible in terms of memory.

В PROOF OF THEORETICAL RESULTS

Theorem 5 (Guan et al. (2024a)). The optimal identical Blue team policy $\phi_{\infty}^* \in \Phi$ obtained from the equivalent zero-sum coordinator game is ϵ -optimal Blue team policy. Formally, for all joint states \mathbf{x}^{N_1} and \mathbf{y}^{N_2} ,

$$\min_{\psi^{N_2} \in \Psi^{N_2}} J^{N,\phi_{\infty}^*,\psi^{N_2}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) \ge \underline{J}^{N*}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{\underline{N}}}\right), \text{ where } \underline{N} = \min\{N_1,N_2\}.$$
(B.4)

Theorem 1. The value of the optimal identical Blue team policy ϕ^* obtained from the finite population game is within ϵ of the finite-population lower game value defined in (4). Formally, for all joint states \mathbf{x}^{N_1} and \mathbf{y}^{N_2} ,

$$\min_{\psi^{N_2}} J^{N,\phi^*,\psi^{N_2}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) \ge \underline{J}^{N*}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{N}}\right), \text{ where } \underline{N} = \min\{N_1,N_2\}.$$
 (6)

Sketch of Proof. The proof follows by restricting the optimization domain in 4 to that of identical team policies and then using the definition of the \max operator in tandem with Theorem 5.

Proof. We have the following definition of the lower game value for the finite-population ZS-MFTG:

$$\underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) = \max_{\phi^{N_1} \in \Phi^{N_1}} \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^{N_1}, \psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}).$$
(B.5)

Note that the maximization for the Blue team is being performed over the set of all team policies Φ^{N_1} , including identical as well non-identical team policies. If we restrict ourselves to the set of identical team policies $\Phi \subseteq \Phi^{N_1}$ it follows immediately that

$$\underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \ge \max_{\phi^{N_1} \in \Phi} \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^{N_1}, \psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}).$$
(B.6)

Suppose that ϕ^* is the optimal identical Blue team policy obtained from the finite population game. It follows from (B.5) that

$$\underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \ge \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^*, \psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}). \tag{B.7}$$

Furthermore, let $\phi_{\infty}^* \in \Phi$ be the optimal identical local Blue team policy obtained from the equivalent zero-sum infinite-population coordinator game (recall, one-to-one correspondence between coordinator policy and local identical team policy). By the definition of the optimality of ϕ^* in the space of identical team policies,

$$\min_{\psi^{N_2} \in \Psi^{N_2}} J^{N,\phi^*,\psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \ge \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N,\phi_{\infty}^*,\psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}).$$
(B.8)

Using Theorem 5, and using (B.8), yields the following sequence of inequalities,

$$\begin{split} \underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) &\geq \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^*, \psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \\ &\geq \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N, \phi^*_{\infty}, \psi^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \\ &\geq \underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{N}}\Big), \end{split}$$

where $\underline{N} = \min\{N_1, N_2\}$, thereby completing the proof.

Theorem 2. The approximate policy gradient of the infinite-population Blue (Red) team coordinator policy α (β) computed using local policies from the finite-population ZS-MFTG via MF-MAPPO ($\hat{J}^{N_1}(\alpha_{\theta})$) uniformly tends to the true policy gradient as the population size increases, i.e., $\|\nabla_{\theta}J_{\infty}(\alpha_{\theta})-\nabla_{\theta}\hat{J}^{N_1}(\alpha_{\theta})\|_2 \to 0$ as $(N_1,N_2)\to\infty$, where $\|\cdot\|_2$ is the 2-norm.

Sketch of Proof. In MF-MAPPO we directly train the finite-population local policies (ϕ_t, ψ_t) . As a result, the gradient of the coordinator policy $\nabla_{\theta}J(\alpha_{\theta})$ only has action samples $u_{i,t}^{N_1}, i \in [N_1]$ from the finite-population local policies $\phi_t(\cdot|x_{i,t}^{N_1}, \mu_t^{N_1}, \nu_t^{N_2})$ and not $\pi_t^{N_1}(\cdot|x_{i,t}^{N_1}) \sim \alpha(\mu_t^{N_1}, \nu_t^{N_2})$ itself.

Thus, our first step is to construct an approximation of $\pi_t^{N_1}$ using the obtained finite-population action samples $u_{i,t}^{N_1}, i \in [N_1]$, resulting in an approximate finite-population policy gradient expression given by $\nabla_{\theta} \hat{J}^{N_1}(\alpha_{\theta})$. Coupled with existing mean-field approximation results (Cui et al., 2023; Guan et al., 2024a; Shao et al., 2024a) for the infinite-population limit $N_1, N_2 \to \infty$ (LLN, etc.) we show convergence.

Proof. From Guan et al. (2024a) it follows that there exists a one-to-one correspondence between the *deterministic* infinite-population coordinator policies and the local identical team policies followed by the finite-population Blue and Red agents. We extend this formulation to potentially *stochastic* infinite-population coordinator policies and the local finite-population policies as follows: a stochastic Blue coordination policy $\alpha \in \mathcal{A}$ induces an identical team policy $\phi \in \Phi$ according to the rule

$$\phi_t(u_t|x_t, \mu_t^{\rho}, \nu_t^{\rho}) = \int_{\pi \in \Pi} \pi_t(u_t|x_t) \alpha_t(\pi_t|\mu_t^{\rho}, \nu_t^{\rho}) d\pi \quad \forall \mu_t \in \mathcal{P}(\mathcal{X}), \ \nu_t \in \mathcal{P}(\mathcal{Y}), \ x_t \in \mathcal{X} \text{ and } u_t \in \mathcal{U},$$
(B.9)

where $\pi_t(u_t|x_t)$ corresponds to the identical local policies for all states $x \in \mathcal{X}$ prescribed by the coordinator, i.e., $\pi_t \sim \alpha(\mu_t^{\rho}, \nu_t^{\rho})$.

It is important to note that in MF-MAPPO we directly work with the finite-population local policies (ϕ,ψ) , unlike prior works that train on the infinite-population coordinator policy, see Section 1. Consequently, the gradient of the coordinator policy $\nabla_{\theta}J(\alpha_{\theta})$ computed using this finite-population approximation (denoted $\nabla_{\theta}\hat{J}^{N_1}(\alpha_{\theta})$) only has action samples $u_{i,t}^{N_1}$, $i\in[N_1]$ from the finite-population local policies $\phi_t(\cdot|x_{i,t}^{N_1},\mu_t^{N_1},\nu_t^{N_2})$ and not direct access to $\pi_t^{N_1}\sim\alpha(\mu_t^{N_1},\nu_t^{N_2})$. Thus, our first step is to construct an approximation of the policy based on the obtained action samples. We define this approximate policy from the obtained samples as follows,

$$\hat{\pi}_{t}^{N_{1}}(u|x) = \begin{cases} \frac{\sum_{i=1}^{N_{1}} \mathbb{1}_{x} \left(X_{i,t}^{N_{1}}\right) \mathbb{1}_{u} \left(U_{i,t}^{N_{1}}\right)}{N_{1} \mathcal{M}_{t}^{N_{1}}(x)} & \text{if } \mathcal{M}_{t}^{N_{1}}(x) > 0, \\ 1/|\mathcal{U}| & \text{if } \mathcal{M}_{t}^{N_{1}}(x) = 0. \end{cases}$$
(B.10)

We can similarly define $\hat{\sigma}_t^{N_2}$ for the Red team. Using this constructed empirical policy $\hat{\pi}_t^{N_1}$ from the sampled actions $\mathbf{u}_t^{N_1}$, the policy gradient $\nabla_{\theta}\hat{J}^{N_1}(\alpha_{\theta})$ for the finite-population ZS-MFTG is,

$$\nabla_{\theta} \hat{J}^{N_1}(\alpha_{\theta}) = \sum_{t=0}^{\infty} \gamma^t \mathbb{E}_{\mathbf{u}^{N_1} \sim \phi, \mathbf{v}^{N_2} \sim \psi} \left[Q^{N, \alpha, \beta}(\mu_t^{N_1}, \nu_t^{N_2}, \hat{\pi}_t^{N_1}, \hat{\sigma}_t^{N_2}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_t^{N_1} | \mu_t^{N_1}, \nu_t^{N_2}) \right],$$
(B.11)

where,

$$\begin{split} Q^{N,\alpha,\beta}(\mu_t^{N_1},\nu_t^{N_2},\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2}) \\ &= \mathbb{E}_{\mathbf{u}^{N_1}\sim\phi,\mathbf{v}^{N_2}\sim\psi} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(\mu_{\tau}^{N_1},\nu_{\tau}^{N_2}) \big| \mu_{\tau=0}^{N_1} = \mu_t^{N_1}, \nu_{\tau=0}^{N_2} = \nu_t^{N_2}, \pi_{\tau=0}^{N_1} = \hat{\pi}_t^{N_1}, \sigma_{\tau=0}^{N_2} = \hat{\sigma}_t^{N_2} \right]. \end{split} \tag{B.12}$$

Recall that \mathbf{u}^{N_1} and \mathbf{v}^{N_2} enter (B.11) and (B.12) through the construction of $\hat{\pi}_t^{N_1}$ (B.10) and $\hat{\sigma}_t^{N_2}$ respectively. Furthermore, we have the following expression for the gradient of the infinite-population coordinator game:

$$\nabla_{\theta} J(\alpha_{\theta}) = \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\pi \sim \alpha, \sigma \sim \beta} \left[Q^{\alpha, \beta}(\mu_{t}, \nu_{t}, \pi_{t}, \sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t} | \mu_{t}, \nu_{t}) \right], \tag{B.13}$$

where

$$Q^{\alpha,\beta}(\mu,\nu,\pi,\sigma) = \mathbb{E}_{\pi\sim\alpha,\sigma\sim\beta} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(\mu_{\tau},\nu_{\tau}) \middle| \mu_{\tau=0} = \mu, \nu_{\tau=0} = \nu, \pi_{\tau=0} = \pi, \sigma_{\tau=0} = \sigma \right].$$
(B.14)

While (B.12) and (B.14) utilize the joint coordinator policies (α, β) , the policy gradients (B.11) and (B.13) are taken for each team *individually*. Consequently the policy gradient analysis is done

on a per-team basis. We focus on the policy gradient with respect to the Blue team but the analysis of the Red team's policy gradient is symmetric and can be proved by an identical approach.

The following lemma relates the expectation in (B.11) to (α, β) using (B.9).

Lemma 1. Given a function $g: \mathcal{U} \to \mathbb{R}$ of the joint actions $\mathbf{u}_t^{N_1}$ such that $u_{i,t}^{N_1} \sim \phi(\cdot|\mathbf{x}_{i,t}^{N_1}, \mu_t^{N_1}, \nu_t^{N_2})$, or equivalently, $\mathbf{u}_t^{N_1} \sim \phi(\cdot|\mathbf{x}_t^{N_1}, \mu_t^{N_1}, \nu_t^{N_2})$,

$$\mathbb{E}_{\mathbf{u}_t^{N_1} \sim \phi} \left[g(\mathbf{u}_t^{N_1}) \right] = \mathbb{E}_{\pi \sim \alpha} \left[\mathbb{E}_{\mathbf{u}_t^{N_1} \sim \pi} \left[g(\mathbf{u}_t^{N_1}) \right] \right]$$

Proof. By the definition of expectation,

$$\mathbb{E}_{\mathbf{u}_{t}^{N_{1}} \sim \phi} \left[g(\mathbf{u}_{t}^{N_{1}}) \right] = \sum_{\mathbf{u}_{t}^{N_{1}}} g(\mathbf{u}_{t}^{N_{1}}) \phi(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}, \mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}),$$

where the sum is taken over all the possible joint actions that can be sampled from ϕ . By using the definition of ϕ from (B.9),

$$\begin{split} \sum_{\mathbf{u}_{t}^{N_{1}}} g(\mathbf{u}_{t}^{N_{1}}) \phi(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}, \boldsymbol{\mu}_{t}^{N_{1}}, \boldsymbol{\nu}_{t}^{N_{2}}) &= \sum_{\mathbf{u}_{t}^{N_{1}}} g(\mathbf{u}_{t}^{N_{1}}) \underbrace{\int_{\boldsymbol{\pi} \in \boldsymbol{\Pi}} \boldsymbol{\pi}_{t}(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}) \alpha_{t}(\boldsymbol{\mu}_{t}^{N_{1}}, \boldsymbol{\nu}_{t}^{N_{2}}) d\boldsymbol{\pi}}_{= \mathbb{E}_{\boldsymbol{\pi} \sim \boldsymbol{\alpha}} \left[\boldsymbol{\pi}_{t}(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}) \right]} \\ &= \sum_{\mathbf{u}_{t}^{N_{1}}} g(\mathbf{u}_{t}^{N_{1}}) \mathbb{E}_{\boldsymbol{\pi} \sim \boldsymbol{\alpha}} \left[\boldsymbol{\pi}_{t}(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}) \right] \\ &= \sum_{\mathbf{u}_{t}^{N_{1}}} \mathbb{E}_{\boldsymbol{\pi} \sim \boldsymbol{\alpha}} \left[g(\mathbf{u}_{t}^{N_{1}}) \boldsymbol{\pi}_{t}(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}) \right] \\ &= \mathbb{E}_{\boldsymbol{\pi} \sim \boldsymbol{\alpha}} \left[\sum_{\mathbf{u}_{t}^{N_{1}}} g(\mathbf{u}_{t}^{N_{1}}) \boldsymbol{\pi}_{t}(\mathbf{u}_{t}^{N_{1}} | \mathbf{x}_{t}^{N_{1}}) \right] , \end{split}$$

where the last two expressions follow from linearity (and finiteness) of expectation. It follows then,

$$\mathbb{E}_{\mathbf{u}_t^{N_1} \sim \phi} \left[g(\mathbf{u}_t^{N_1}) \right] = \mathbb{E}_{\pi \sim \alpha} \left[\mathbb{E}_{\mathbf{u}_t^{N_1} \sim \pi} \left[g(\mathbf{u}_t^{N_1}) \right] \right]$$

We also explicitly state results regarding the boundedness of terms appearing in (B.11)-(B.14).

Remark 2. As mentioned in the main text, $r_t \in [-R_{\max}, R_{\max}]$. Consequently, the Q-functions $Q^{\alpha,\beta}$ and $Q^{N,\alpha,\beta}$ from (B.14) and (B.12) for all $\gamma \in (0,1)$ are uniformly bounded with,

$$\|Q^{\alpha,\beta}\|_{\infty} \leq \frac{R_{\max}}{1-\gamma}, \ \ and \ \ \|Q^{N,\alpha,\beta}\|_{\infty} \leq \frac{R_{\max}}{1-\gamma}.$$

We further assume continuity with respect to the local policies, i.e.,

Assumption 5. The Q-functions $Q^{\alpha,\beta}$ and $Q^{N,\alpha,\beta}$ from (B.14) and (B.12) are continuous in inputs (π,σ) and $(\hat{\pi}^{N_1},\hat{\sigma}^{N_2})$.

It is pertinent to note that the action-space Π and Σ are continuous. Thus, although Assumption 5 seems restrictive, the use of function approximators to train $Q^{\alpha,\beta}$ and $Q^{N,\alpha,\beta}$ like neural networks (endowed with continuous activation functions) in such continuous action-space settings, ensures that the resultant functions are continuous in their inputs Lillicrap et al. (2019).

The following assumption, as done in many works (Cui et al., 2024; 2023), states that the gradient of coordinator policy is continuous and uniformly bounded, i.e.,

Assumption 6. The log-gradient $\nabla \log \alpha_{\theta}(\pi_t | \mu_t, \nu_t)$ is L_{∇} -Lipschitz continuous and uniformly bounded.

Furthermore, we have the following result from Guan et al. (2024a):

Corollary 1. Let $\mathbf{X}_t^{N_1}$, $\mathbf{Y}_t^{N_2}$, $\mathcal{M}_t^{N_1}$, and $\mathcal{N}_t^{N_2}$ be the joint states and the corresponding EDs of a finite-population game. Denote the next Blue ED induced by an identical Blue team policy $\phi_t \in \Phi_t$ as $\mathcal{M}_{t+1}^{N_1}$. Then, the following holds:

$$\mathbb{E}_{\phi_t} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \mu_{t+1} \right) \middle| \mathbf{X}_t^{N_1}, \mathbf{Y}_t^{N_2} \right] \le \frac{|\mathcal{X}|}{2} \sqrt{\frac{1}{N_1}},$$

where $\mu_{t+1} = \mathcal{M}_t^{N_1} F_t(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}, \phi_t)$.

We are interested in deriving an explicit result for $\|\nabla_{\theta}J(\alpha_{\theta}) - \nabla_{\theta}\hat{J}^{N_{1}}(\alpha_{\theta})\|_{2}$ as the population sizes $(N_{1}, N_{2}) \to \infty$. In particular, using shorthand notation $\mathbb{E}_{\alpha,\beta}\left[\cdot\right] \triangleq \mathbb{E}_{\pi \sim \alpha,\sigma \sim \beta}\left[\cdot\right]$ and $\mathbb{E}_{\phi,\psi}\left[\cdot\right] \triangleq \mathbb{E}_{\mathbf{u}^{N_{1}} \sim \phi, \mathbf{v}^{N_{2}} \sim \psi}\left[\cdot\right]$)

$$\begin{split} & \|\nabla_{\theta} J(\alpha_{\theta}) - \nabla_{\theta} \hat{J}^{N_{1}}(\alpha_{\theta})\|_{2} \\ & = \left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta} \left[Q^{\alpha,\beta}(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t}|\mu_{t},\nu_{t}) \right] \\ & - \gamma^{t} \mathbb{E}_{\phi,\psi} \left[Q^{N,\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}}|\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) \right] \right\|_{2}. \end{split}$$

For the rest of the proof, we implicitly assume $\|\cdot\| \equiv \|\cdot\|_2$ to denote the 2-norm. By Lemma 1, we can rewrite the above as

$$\begin{split} & \|\nabla_{\theta} J(\alpha_{\theta}) - \nabla_{\theta} \hat{J}^{N_{1}}(\alpha_{\theta})\| \\ & = \left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta} \left[Q^{\alpha,\beta}(\mu_{t}, \nu_{t}, \pi_{t}, \sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t} | \mu_{t}, \nu_{t}) \right. \\ & \left. - \mathbb{E}_{\mathbf{u} \sim \pi, \mathbf{v} \sim \sigma} \left[Q^{N,\alpha,\beta}(\mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}, \hat{\pi}_{t}^{N_{1}}, \hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}} | \mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}) \right] \right] \right\|, \end{split}$$

where $\hat{\pi}_t^{N_1}, \hat{\sigma}_t^{N_2}$ depend on the random variables $\mathbf{u}^{N_1} \sim \pi, \mathbf{v}^{N_2} \sim \sigma$ respectively. We add and subtract the term $\mathbb{E}_{\mathbf{u}^{N_1} \sim \pi, \mathbf{v}^{N_2} \sim \sigma} \left[Q^{\alpha,\beta}(\mu_t^{N_1}, \nu_t^{N_2}, \hat{\pi}_t^{N_1}, \hat{\sigma}_t^{N_2}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_t^{N_1} | \mu_t^{N_1}, \nu_t^{N_2}) \right]$ and split the gradient terms as follows.

$$\|\nabla_{\theta}J(\alpha_{\theta}) - \nabla_{\theta}\hat{J}^{N_{1}}(\alpha_{\theta})\|$$

$$\leq \left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta} \left[Q^{\alpha,\beta}(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t}|\mu_{t},\nu_{t}) \right] \right\|$$

$$- \mathbb{E}_{\mathbf{u}^{N_{1}} \sim \pi,\mathbf{v}^{N_{2}} \sim \sigma} \left[Q^{\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}}|\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) \right] \right\|$$

$$+ \left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta,\mathbf{u}^{N_{1}} \sim \pi,\mathbf{v}^{N_{2}} \sim \sigma} \left[\left(Q^{\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) - Q^{N,\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \right) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}}|\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) \right] \right\|.$$
(B.15)

We analyze the two terms individually. For the first term,

$$\left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta} \left[Q^{\alpha,\beta}(\mu_{t}, \nu_{t}, \pi_{t}, \sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t} | \mu_{t}, \nu_{t}) \right. \\ \left. - \mathbb{E}_{\mathbf{u}^{N_{1}} \sim \pi, \mathbf{v}^{N_{2}} \sim \sigma} \left[Q^{\alpha,\beta}(\mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}, \hat{\pi}_{t}^{N_{1}}, \hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}} | \mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}) \right] \right] \right\|, \tag{B.16}$$

Since the first term in the above expression is a constant with respect to the random variables $\mathbf{u}^{N_1}, \mathbf{v}^{N_2}$ and the expectation of a constant is the value itself, we can equivalently write (B.16)

$$\left\| \sum_{t=0}^{\infty} \gamma^{t} \mathbb{E}_{\alpha,\beta,\mathbf{u}^{N_{1}} \sim \pi,\mathbf{v}^{N_{2}} \sim \sigma} \left[Q^{\alpha,\beta}(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t}|\mu_{t},\nu_{t}) - Q^{\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}}|\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) \right] \right\|,$$
(B.17)

We can split the sum from t=0 to t=T and t=T+1 to $t=\infty$, By uniform bounds on $Q^{\alpha,\beta}$ from Remark 2 and on $\nabla_{\theta} \log \alpha_{\theta}(\;\cdot\;|\;\mu_t^{N_1},\nu_t^{N_2}), \nabla_{\theta} \log \alpha_{\theta}(\;\cdot\;|\;\mu_t,\nu_t)$ (Assumption 6), the sum from t=T+1 to $t=\infty$ goes to zero for T large enough (tail sequence). This enables us to focus on the summation from t=0 to t=T, namely,

$$\left\| \sum_{t=0}^{T} \gamma^{t} \mathbb{E}_{\alpha,\beta,\mathbf{u}^{N_{1}} \sim \pi,\mathbf{v}^{N_{2}} \sim \sigma} \left[Q^{\alpha,\beta}(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) \nabla_{\theta} \log \alpha_{\theta}(\pi_{t}|\mu_{t},\nu_{t}) - Q^{\alpha,\beta}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \nabla_{\theta} \log \alpha_{\theta}(\hat{\pi}_{t}^{N_{1}}|\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) \right] \right\|.$$
(B.18)

Define $h(\mu, \nu, \pi, \sigma) \triangleq Q^{\alpha,\beta}(\mu, \nu, \pi, \sigma) \nabla_{\theta} \log \alpha_{\theta}(\pi|\mu, \nu)$. Assumptions 5 and 6 ensure continuity of the function $h(\mu, \nu, \pi, \sigma)$ with respect to all its inputs. Using this definition of h, rewrite (B.18)

$$\left\| \sum_{t=0}^{T} \gamma^{t} \mathbb{E}_{\alpha,\beta,\mathbf{u} \sim \pi,\mathbf{v} \sim \sigma} \left[h(\mu_{t}, \nu_{t}, \pi_{t}, \sigma_{t}) - h(\mu_{t}^{N_{1}}, \nu_{t}^{N_{2}}, \hat{\pi}_{t}^{N_{1}}, \hat{\sigma}_{t}^{N_{2}}) \right] \right\|. \tag{B.19}$$

Using the definition of the expectation operators,

$$\begin{split} & \left\| \sum_{t=0}^{T} \gamma^{t} \mathbb{E}_{\alpha,\beta,\mathbf{u}^{N_{1}} \sim \pi,\mathbf{v}^{N_{2}} \sim \sigma} \left[h(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) - h(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \right] \right\| \\ & = \left\| \sum_{t=0}^{T} \gamma^{t} \int_{\pi} \int_{\sigma} h(\mu_{t},\nu_{t},\pi_{t},\sigma_{t}) \beta_{t}(\mu_{t},\nu_{t}) \alpha_{t}(\mu_{t},\nu_{t}) d\sigma d\pi \\ & - \int_{\pi} \int_{\sigma} \left[\sum_{\mathbf{u}^{N_{1}}} \sum_{\mathbf{v}^{N_{2}}} h(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}},\hat{\pi}_{t}^{N_{1}},\hat{\sigma}_{t}^{N_{2}}) \pi(\mathbf{u}^{N_{1}}|\mathbf{x}_{t}^{N_{1}}) \sigma(\mathbf{v}^{N_{2}}|\mathbf{y}_{t}^{N_{1}}) \right] \beta_{t}(\mu_{t},\nu_{t}) \alpha_{t}(\mu_{t},\nu_{t}) d\sigma d\pi \right\| \end{split}$$

We already have, by weak LLN that $d_{TV}\left(\mu_t^{N_1},\mu_t\right) \leq \mathcal{O}(1/\sqrt{N_1})$ and $d_{TV}\left(\nu_t^{N_2},\nu_t\right) \leq \mathcal{O}(1/\sqrt{N_2})$, with the error going to zero as the population size becomes large, i.e., $N_1,N_2\to\infty$. A similar argument using weak LLN can be made for $(\pi_t^{N_1},\sigma_t^{N_2})\to(\pi_t,\sigma_t)$ as $N_1,N_2\to\infty$. Furthermore, as the population size becomes large, the empirical policies constructed $(\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2})\to(\pi_t^{N_1},\sigma_t^{N_2})$ because the number of samples of $\mathbf{u}_t^{N_1}$ and $\mathbf{v}_t^{N_2}$ scale with (N_1,N_2) . Thus, it follows that $(\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2})\to(\pi_t,\sigma_t)$ as the team sizes become large and, coupled with the continuity of h with respect to $\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2}$, the following simplification holds:

$$\begin{split} \sum_{\mathbf{u}} \sum_{\mathbf{v}} h(\mu_t^{N_1}, \nu_t^{N_2}, \hat{\pi}_t^{N_1}, \hat{\sigma}_t^{N_2}) \pi(\mathbf{u} | \mathbf{x}_t^{N_1}) \sigma(\mathbf{v} | \mathbf{y}_t^{N_1}) &\rightarrow \sum_{\mathbf{u}} \sum_{\mathbf{v}} h(\mu_t, \nu_t, \pi_t, \sigma_t) \pi(\mathbf{u} | \mathbf{x}_t^{N_1}) \sigma(\mathbf{v} | \mathbf{y}_t^{N_1}) \\ &= h(\mu_t, \nu_t, \pi_t, \sigma_t) \sum_{\mathbf{u}} \pi(\mathbf{u}^{N_1} | \mathbf{x}_t^{N_1}) \sum_{\mathbf{v}} \sigma(\mathbf{v}^{N_2} | \mathbf{y}_t^{N_1}) \\ &= h(\mu_t, \nu_t, \pi_t, \sigma_t) \end{split}$$

For the second term in (B.15), we first focus on the term

$$\left(Q^{\alpha,\beta}(\mu_t^{N_1},\nu_t^{N_2},\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2}) - Q^{N,\alpha,\beta}(\mu_t^{N_1},\nu_t^{N_2},\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2})\right).$$

The expectation in (B.12) is with respect to $\mathbf{u}^{N_1} \sim \phi, \mathbf{v}^{N_2} \sim \psi$. It follows from Lemma 1, we can rewrite (B.12)as

$$\begin{split} &Q^{N,\alpha,\beta}(\mu_t^{N_1},\nu_t^{N_2},\hat{\pi}_t^{N_1},\hat{\sigma}_t^{N_2}) \\ &= \mathbb{E}_{\pi\sim\alpha,\sigma\sim\beta,\mathbf{u}^{N_1}\sim\pi,\mathbf{v}^{N_2}\sim\sigma} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(\mu_{\tau}^{N_1},\nu_{\tau}^{N_2}) \big| \mu_0^{N_1} = \mu_t^{N_1}, \nu_0^{N_2} = \nu_t^{N_2}, \pi_0^{N_1} = \hat{\pi}_t^{N_1}, \sigma_0^{N_2} = \hat{\sigma}_t^{N_2} \right] \end{split}$$

¹Note that as $N_1 \to \infty$, $\hat{\pi}_t^{N_1}(\cdot|x) \neq \pi_t^{N_1}(\cdot|x)$ for states x such that $\mu(x) = 0$ by the very definition of the approximation (B.10). However, it is easy to check that the mean-field trajectories remain the same and the unoccupied states have no role to play in the evolution.

On the other hand, the expectation in (B.14) is with respect to α, β . Using a similar argument as (B.17), we can write (B.14) as

$$Q^{\alpha,\beta}(\mu_t^{N_1}, \nu_t^{N_2}, \hat{\pi}_t^{N_1}, \hat{\sigma}_t^{N_2})$$

$$= \mathbb{E}_{\pi \sim \alpha, \sigma \sim \beta, \mathbf{u}^{N_1} \sim \pi, \mathbf{v}^{N_2} \sim \sigma} \left[\sum_{\tau=0}^{\infty} \gamma^{\tau} r(\mu_{\tau}, \nu_{\tau}) \middle| \mu_0 = \mu_t^{N_1}, \nu_0 = \nu_t^{N_2}, \pi_0 = \hat{\pi}_t^{N_1}, \sigma_0 = \hat{\sigma}_t^{N_2} \right].$$

We apply the following lemma along with Assumption 6 to establish convergence of the second term in (B.15).

Lemma 2. Under Assumption 1 and Lipschitz policies B.17, $Q^{N,\alpha,\beta} \to Q^{\alpha,\beta}$ as $N_1, N_2 \to \infty$.

Proof. Assumption 1 allows the $Q^{N,\alpha,\beta}, Q^{\alpha,\beta}$ to have compact support. Thus, it suffices to prove pointwise convergence (Cui et al., 2023) for each τ , i.e.,

$$\mathbb{E}\left[\gamma^{\tau}r(\mu_{\tau}^{N_{1}},\nu_{\tau}^{N_{2}})\big|\mu_{0}^{N_{1}}=\mu_{t}^{N_{1}},\nu_{0}^{N_{2}}=\nu_{t}^{N_{2}},\pi_{0}^{N_{1}}=\hat{\pi}_{t}^{N_{1}},\sigma_{0}^{N_{2}}=\hat{\sigma}_{t}^{N_{2}}\right]\rightarrow$$

$$\mathbb{E}\left[\gamma^{\tau}r(\mu_{\tau},\nu_{\tau})\big|\mu_{0}=\mu_{t}^{N_{1}},\nu_{0}=\nu_{t}^{N_{2}},\pi_{0}=\hat{\pi}_{t}^{N_{1}},\sigma_{0}=\hat{\sigma}_{t}^{N_{2}}\right]$$
(B.20)

as $N_1, N_2 \to \infty$, where the expectation is taken with respect to $\alpha, \beta, \mathbf{u} \sim \pi, \mathbf{v} \sim \sigma$ (dropped for readability). We follow an induction approach similar to the one presented in Cui et al. (2023). It trivially holds at $\tau = 0$. Let us assume that (B.20) holds at time step τ . For subsequent $\tau + 1$,

$$\begin{split} &\gamma^{\tau} \Big\| \mathbb{E} \left[r(\mu_{\tau+1}, \nu_{\tau+1}) \big| \mu_0 = \mu_t^{N_1}, \nu_0 = \nu_t^{N_2}, \pi_0 = \hat{\pi}_t^{N_1}, \sigma_0 = \hat{\sigma}_t^{N_2} \right] \\ &- \mathbb{E} \left[r(\mu_{\tau+1}^{N_1}, \nu_{\tau+1}^{N_2}) \big| \mu_0^{N_1} = \mu_t^{N_1}, \nu_0^{N_2} = \nu_t^{N_2}, \pi_0^{N_1} = \hat{\pi}_t^{N_1}, \sigma_0^{N_2} = \hat{\sigma}_t^{N_2} \right] \Big\| \\ &\leq \gamma^{\tau} \Big\| \mathbb{E} \left[r\Big(\mu_{\tau} F(\mu_{\tau}, \nu_{\tau}, \pi_{\tau}), \nu_{\tau} G(\mu_{\tau}, \nu_{\tau}, \sigma_{\tau}) \Big) \big| \mu_0 = \mu_t^{N_1}, \nu_0 = \nu_t^{N_2}, \pi_0 = \hat{\pi}_t^{N_1}, \sigma_0 = \hat{\sigma}_t^{N_2} \right] \\ &- \mathbb{E} \left[r\Big(\mu_\tau^{N_1} F(\mu_\tau^{N_1}, \nu_\tau^{N_2}, \phi_\tau), \nu_\tau^{N_2} G(\mu_\tau^{N_1}, \nu_\tau^{N_2}, \psi_\tau) \Big) \big| \mu_0^{N_1} = \mu_t^{N_1}, \nu_0^{N_2} = \nu_t^{N_2}, \pi_0^{N_1} = \hat{\pi}_t^{N_1}, \sigma_0^{N_2} = \hat{\sigma}_t^{N_2} \right] \Big\| \\ &+ \gamma^{\tau} \Big\| \mathbb{E} \left[r\Big(\mu_\tau^{N_1} F(\mu_\tau^{N_1}, \nu_\tau^{N_2}, \phi_\tau), \nu_\tau^{N_2} G(\mu_\tau^{N_1}, \nu_\tau^{N_2}, \psi_\tau) \Big) \big| \mu_0^{N_1} = \mu_t^{N_1}, \nu_0^{N_2} = \nu_t^{N_2}, \pi_0^{N_1} = \hat{\pi}_t^{N_1}, \sigma_0^{N_2} = \hat{\sigma}_t^{N_2} \right] \\ &- \mathbb{E} \left[r(\mu_{\tau+1}^{N_1}, \nu_{\tau+1}^{N_2}) \big| \mu_0^{N_1} = \mu_t^{N_1}, \nu_0^{N_2} = \nu_t^{N_2}, \pi_0^{N_1} = \hat{\pi}_t^{N_1}, \sigma_0^{N_2} = \hat{\sigma}_t^{N_2} \right] \Big\| \end{aligned}$$

The convergence of the first term follows from the weak LLN Cui et al. (2023) and a similar argument as used for (B.19). By using the Lipshcitzness of the reward function (Assumption 1) and Corollary 1, the second term in the expression is bounded by $\mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$, where $N = \min(N_1, N_2)$ and goes to zero as $(N_1, N_2) \to \infty$.

It follows from the convergence of the two individual terms in B.15 that
$$\|\nabla_{\theta}J(\alpha_{\theta}) - \nabla_{\theta}\hat{J}^{N_1}(\alpha_{\theta})\|$$
 as $(N_1, N_2) \to \infty$

Theorem 3. Let \mathcal{G}_1 denote the finite-population game where the agents utilize the identical team policies ϕ_t^* and ψ_t^* derived from MF-MAPPO trained on \mathcal{G}_1 and let the finite-population game \mathcal{G}_2 with the same state-action space, dynamics, and rewards, but with population sizes \bar{N}_1 and \bar{N}_2 such that $\bar{N}_1/\bar{N}_2 = N_1/N_2$ and $\min(\bar{N}_1,\bar{N}_2) \geq \min(N_1,N_2)$. Then, (ϕ_t^*,ψ_t^*) remain ϵ -optimal for \mathcal{G}_2 .

Sketch of Proof. The key idea of the proof is to exploit the fact that the equivalent infinite-population games for \mathcal{G}_1 and \mathcal{G}_2 are the same as $N_1/N_2 = \bar{N}_1/\bar{N}_2$ (Guan et al., 2024a). We first compute the performance of ϕ^* , the optimal identical team policy derived for the \mathcal{G}_1 using MF-MAPPO (finite-population training), when applied to the infinite-population coordinator setting. We then compare the performance of the aforementioned ϕ^* for \mathcal{G}_2 with the infinite-population coordinator. Then, equating the common equivalent coordinators for \mathcal{G}_1 and \mathcal{G}_2 , the result follows.

Proof. Consider game \mathcal{G}_1 . Let us restrict ourselves to the set of identical team policies $\Phi \subseteq \Phi^{N_1}$ and suppose that ϕ^* is the optimal identical Blue team policy obtained from the (N_1, N_2) finite population game via MF-MAPPO. From Theorem 1, we already know that

$$\min_{\boldsymbol{\psi}^{N_2} \in \boldsymbol{\Psi}^{N_2}} J^{N, \boldsymbol{\phi}^*, \boldsymbol{\psi}^{N_2}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) \geq \underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{N}}\Big).$$

Denote $J^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) \triangleq \min_{\psi^{N_2} \in \Psi^{N_2}} J^{N,\phi^*,\psi^{N_2}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2})$, where ψ^{N_2*} is the potentially non-identical policy that solves the optimization problem.

Now, for \mathcal{G}_2 we want to show that

$$\min_{\boldsymbol{\psi}^{\bar{N}_2} \in \boldsymbol{\Psi}^{\bar{N}_2}} J^{\bar{N}, \phi^*, \boldsymbol{\psi}^{\bar{N}_2}}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) \geq \underline{J}^{\bar{N}*}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{N}}\Big), \tag{B.21}$$

where ϕ^* corresponds to the identical policy derived from \mathcal{G}_1 .

We know that the $\frac{N_1}{N_2} = \frac{\bar{N}_1}{\bar{N}_2}$. Thus, the infinite-population coordinator games remain identical for both \mathcal{G}_1 and \mathcal{G}_2 . We prove the theorem in the following two steps.

Step 1. We first show that for ϕ^* from \mathcal{G}_1 we have the following inequality for all joint states $\mathbf{x}^{N_1} \in \mathcal{X}^{N_1}$ and $\mathbf{y}^{N_2} \in \mathcal{Y}^{N_2}$,

$$J_{\infty}^{\phi^*, \tilde{\psi}_{\infty}^*}(\mu^{N_1}, \nu^{N_2}) \ge J^{N, \phi^*, \psi^{N_2*}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{N}}\right), \tag{B.22}$$

where $\mu^{N_1} = \operatorname{Emp}_{\mu}(\mathbf{x}^{N_1})$ and $\nu^{N_2} = \operatorname{Emp}_{\nu}(\mathbf{y}^{N_2})$ and $\tilde{\psi}_{\infty}^* \in \Psi$ is given by,

$$J_{\infty}^{\phi^*, \tilde{\psi}_{\infty}^*}(\mu_0, \nu_0) = \min_{\psi \in \Psi} J_{\infty}^{\phi^*, \psi}(\mu_0, \nu_0).$$
 (B.23)

The above equation (B.23) determines the optimal Red team response in the *infinite-population* domain, when the Blue team plays ϕ^* derived from \mathcal{G}_1 . We prove this through an inductive argument.

Base case: At the terminal timestep T, the two value functions are the same. Thus, we have, for all joint states $\mathbf{x}_T^{N_1} \in \mathbf{X}^{N_1}$ and $\mathbf{y}_T^{N_2} \in \mathbf{Y}^{N_2}$, that

$$J_{T}^{\phi^*,\tilde{\psi}_{T}^*}(\mu_T^{N_1},\nu_T^{N_1}) = J_{T}^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}_T^{N_1},\mathbf{y}_T^{N_2}) = r_t(\mu_T^{N_1},\nu_T^{N_1}).$$

Inductive hypothesis: Assume that, at time step t+1, the following holds for all $\mathbf{x}_{t+1}^{N_1} \in \mathbf{X}^{N_1}$ and $\mathbf{y}_{t+1}^{N_2} \in \mathbf{Y}^{N_2}$,

$$J_{\infty,t+1}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_{t+1}^{N_1},\nu_{t+1}^{N_1}) \geq J_{t+1}^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}_{t+1}^{N_1},\mathbf{y}_{t+1}^{N_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{N}}\Big). \tag{B.24}$$

Induction: Consider arbitrary $\mathbf{x}_t^{N_1} \in \mathbf{X}^{N_1}$ and $\mathbf{y}_t^{N_2} \in \mathbf{Y}^{N_2}$.

For simplicity, we do not emphasize the correspondence between the joint states and the EDs for the rest of the proof, as it is clear from the context. For the identical team policies $(\phi_t^*, \psi_t) \in \Phi \times \Psi$, at time step t, denote

$$\begin{split} \mu_{t+1}^{\phi^*} &\triangleq \mu_t^{N_1} F(\mu_t^{N_1}, \nu_t^{N_2}, \phi_t^*) \\ \nu_{t+1}^{\psi} &\triangleq \nu_t^{N_2} G(\mu_t^{N_1}, \nu_t^{N_2}, \psi_t). \end{split} \tag{B.25}$$

 For notational simplicity, we drop the conditions $\mathbf{X}_t^{N_1} = \mathbf{x}_t^{N_1}$ and $\mathbf{Y}_t^{N_2} = \mathbf{y}_t^{N_2}$ in the following derivations. Then, we have

$$\begin{aligned} & J_{t}^{1190} & J_{t}^{N,\phi^{*},\psi^{N_{2}^{*}}}(\mathbf{x}_{t}^{N_{1}},\mathbf{y}_{t}^{N_{2}}) \\ & = \min_{\psi_{t}^{N_{2}} \in \Psi_{t}^{N_{2}}} r(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathbb{E}_{\phi_{t}^{*},\psi_{t}^{N_{2}}} \left[J_{t+1}^{N,\phi^{*},\psi^{N_{2}^{*}}}(\mathbf{X}_{t+1}^{N_{1}},\mathbf{Y}_{t+1}^{N_{2}}) \right] \\ & = \min_{\psi_{t}^{N_{2}} \in \Psi_{t}^{N_{2}}} r(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathbb{E}_{\phi_{t}^{*},\psi_{t}^{N_{2}}} \left[J_{\infty,t+1}^{\phi^{*},\psi^{N_{2}^{*}}}(\mathbf{X}_{t+1}^{N_{1}},\mathbf{Y}_{t+1}^{N_{2}}) \right] \\ & \stackrel{(ii)}{\leq} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \min_{\psi_{t} \in \Psi_{t}} \mathbb{E}_{\phi_{t}^{*},\psi_{t}} \left[J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mathcal{M}_{t+1}^{N_{1}},\mathcal{N}_{t+1}^{N_{2}}) \right] + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \\ & \stackrel{(iii)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) + \min_{\psi_{t} \in \Psi_{t}} \mathbb{E}_{\phi_{t}^{*},\psi_{t}} \left[J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mathcal{M}_{t+1}^{N_{1}},\mathcal{N}_{t+1}^{N_{2}}) - J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) \right] \\ & \stackrel{(iii)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) + \min_{\psi_{t} \in \Psi_{t}} \mathbb{E}_{\phi_{t}^{*},\psi_{t}} \left[J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mathcal{M}_{t+1}^{N_{1}},\mathcal{N}_{t+1}^{N_{2}}) - J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) \right] \\ & \stackrel{(iii)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) + \min_{\psi_{t} \in \Psi_{t}} \mathbb{E}_{\phi_{t}^{*},\psi_{t}} \left[J_{\infty,t+1}^{\phi^{*},\psi_{t}^{*}}(\mathcal{M}_{t+1}^{N_{1}},\mathcal{N}_{t+1}^{N_{2}}) - J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) \right] \\ & \stackrel{(iii)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) + \min_{\psi_{t} \in \Psi_{t}} \mathbb{E}_{\phi_{t}^{*},\psi_{t}} \left[J_{\infty,t+1}^{\phi^{*},\psi_{t}^{*}}(\mathcal{M}_{t+1}^{N_{1}},\mathcal{N}_{t+1}^{N_{2}}) + J_{\infty,t+1}^{\phi^{*},\psi_{t}^{*}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) \right] \\ & \stackrel{(iv)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) + \min_{\psi_{t} \in \Psi_{t}} J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \\ & \stackrel{(vii)}{=} r_{t}(\mu_{t}^{N_{1}},\nu_{t}^{N_{2}}) + \min_{\psi_{t} \in \Psi_{t}} J_{\infty,t+1}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t+1}^{\phi^{*}},\nu_{t+1}^{\psi_{t}}) + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \\ & \stackrel{(vii)}{=} J_{\infty,t}^{\phi^{*},\tilde{\psi}^{*}_{\infty}}(\mu_{t}^{N_{1$$

For inequality (i), we used the inductive hypothesis; for inequality (ii), we reduced the optimization domain of the Red team to the set of *identical* team policies; inequality (iv) is a result of the Lipschitz continuity of the value function (Guan et al., 2024a); for equality (v), we use the fact that the mean-fields $(\mu_{t+1}^{\phi^*}, \nu_{t+1}^{\psi_t})$ are induced deterministically from the distributions at time t; inequality (vi) holds as a consequence of Corollary 1; (vii) converts the optimization domain from the policy space $\psi_t \in \Psi_t$ to the corresponding reachable set $\mathbb{R}_{\nu,t}(\mu_t^{N_1}, \nu_t^{N_2}) \triangleq \{\nu_{t+1} | \exists \psi_t \in \Psi_t \text{ s.t. } \nu_{t+1} = \nu_t G_t(\mu_t, \nu_t, \psi_t)\}$ following (Guan et al., 2024a); (vii) follows from the definition of $J_{\infty,t}^{\phi^*,\tilde{\psi}^*_\infty}(\mu_t^{N_1}, \nu_t^{N_2})$. Thus, $J^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) = J_0^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2})$ and (B.22) follows.

Step 2. We show that for ϕ^* from \mathcal{G}_1 we have the following inequality for all joint states $\mathbf{x}^{\bar{N}_1} \in \mathcal{X}^{\bar{N}_1}$ and $\mathbf{y}^{\bar{N}_2} \in \mathcal{Y}^{\bar{N}_2}$ in \mathcal{G}_2 ,

$$\min_{\psi^{\bar{N}_2} \in \Psi^{\bar{N}_2}} J^{\bar{N},\phi^*,\psi^{\bar{N}_2}}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) \ge J_{\infty}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu^{\bar{N}_1}, \nu^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\bar{N}}}\Big), \tag{B.26}$$

where $\mu^{\bar{N}_1} = \operatorname{Emp}_{\mu}(\mathbf{x}^{\bar{N}_1})$, $\nu^{\bar{N}_2} = \operatorname{Emp}_{\nu}(\mathbf{y}^{\bar{N}_2})$ and $\underline{\bar{N}} = \min(\bar{N}_1, \bar{N}_2)$. The proof is constructed based on induction. Fix an arbitrary Red team policy $\psi^{\bar{N}_2} \in \Psi^{\bar{N}_2}$.

Base case: At the terminal timestep T, since there is no decision to be made, both value functions are equal to the terminal reward and are thus the same. Formally, for all $\mathbf{x}_T^{\bar{N}_1} \in \mathbf{X}^{\bar{N}_1}$ and $\mathbf{y}_T^{\bar{N}_2} \in \mathbf{Y}^{\bar{N}_2}$,

$$J_T^{\bar{N},\phi^*,\psi^{\bar{N}_2}}(\mathbf{x}_T^{\bar{N}_1},\mathbf{y}_T^{\bar{N}_2}) = J_{\infty,T}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_T^{N_1},\nu_T^{N_1}) = r_t(\mu_T^{\bar{N}_1},\nu_T^{\bar{N}_1}),$$

where $\mu_T^{\bar{N}_1} = \mathrm{Emp}_{\mu}(\mathbf{x}_T^{\bar{N}_1})$ and $\nu_T^{\bar{N}_2} = \mathrm{Emp}_{\nu}(\mathbf{y}_T^{\bar{N}_2})$. For simplicity, we do not emphasize the correspondence between the joint states and the EDs for the rest of the proof, as it is clear from the context.

Inductive hypothesis: Assume that at t+1, the following holds for all joint states $\mathbf{x}_{t+1}^{\bar{N}_1} \in \mathbf{X}^{\bar{N}_1}$ and $\mathbf{y}_{t+1}^{\bar{N}_2} \in \mathbf{Y}^{\bar{N}_2}$:

$$J_{t+1}^{\bar{N},\phi^*,\psi^{\bar{N}_2}}(\mathbf{x}_{t+1}^{\bar{N}_1},\mathbf{y}_{t+1}^{N_2}) \ge J_{\infty,t+1}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_{t+1}^{\bar{N}_1},\nu_{t+1}^{\bar{N}_1}) - \mathcal{O}\left(\frac{1}{\sqrt{\bar{N}}}\right). \tag{B.27}$$

Induction: At timestep t, consider an arbitrary pair of joint states $\mathbf{x}_t^{\bar{N}_1} \in \mathbf{X}^{\bar{N}_1}$ and $\mathbf{y}_t^{\bar{N}_2} \in \mathbf{Y}^{\bar{N}_2}$, and their corresponding EDs $\mu_t^{\bar{N}_1}$ and $\nu_t^{\bar{N}_2}$.

For the identical team policy $\phi^* \in \Phi$, denote $\mu_{t+1}^{\phi^*}$) from (B.25). Furthermore, from Theorem 1 in Guan et al. (2024a) there exists a $\nu_{\mathrm{apprx},t+1}^{\psi_t^{\bar{N}_2}}$ for the Red team policy $\psi_t^{N_2}$ such that

$$\mathbb{E}_{\psi_t^{\bar{N}_2}}\left[\mathrm{d}_{\mathrm{TV}}\left(\mathcal{N}_{t+1}^{\bar{N}_2}, \nu_{\mathrm{apprx}, t+1}^{\psi_t^{\bar{N}_2}}\right) \big| \mathbf{X}_t^{\bar{N}_1} = \mathbf{x}_t^{\bar{N}_1}, \mathbf{Y}_t^{\bar{N}_2} = \mathbf{y}_t^{\bar{N}_2}\right] \leq \frac{|\mathcal{Y}|}{2} \sqrt{\frac{1}{\bar{N}_2}}. \tag{B.28}$$

Then, for all joint states $\mathbf{x}_t^{ar{N}_1} \in \mathbf{X}^{ar{N}_1}$ and $\mathbf{y}_t^{ar{N}_2} \in \mathbf{Y}^{ar{N}_2}$, we have

$$\begin{split} J_{t}^{\bar{N},\phi^{*},\psi^{\bar{N}_{2}}}(\mathbf{x}_{t}^{\bar{N}_{1}},\mathbf{y}_{t}^{\bar{N}_{2}}) \\ &= r_{t}(\mu_{t}^{\bar{N}_{1}},\nu_{t}^{\bar{N}_{2}}) + \mathbb{E}_{\phi^{*},\psi^{\bar{N}_{2}}} \left[J_{t+1}^{\bar{N},\phi^{*},\psi^{\bar{N}_{2}}}(\mathbf{X}_{t+1}^{\bar{N}_{1}},\mathbf{Y}_{t+1}^{\bar{N}_{2}}) \right] \\ &\stackrel{(i)}{\geq} r_{t}(\mu_{t}^{\bar{N}_{1}},\nu_{t}^{\bar{N}_{2}}) + \mathbb{E}_{\phi^{*},\psi^{\bar{N}_{2}}} \left[J_{\infty,t+1}^{\phi^{*},\tilde{\psi}_{\infty}^{*}}(\mathcal{M}_{t+1}^{\bar{N}_{1}},\mathcal{N}_{t+1}^{\bar{N}_{2}}) \right] - \mathcal{O}\left(\frac{1}{\sqrt{\underline{N}}}\right) \\ &= r_{t}(\mu_{t}^{\bar{N}_{1}},\nu_{t}^{\bar{N}_{2}}) - \mathcal{O}\left(\frac{1}{\sqrt{\underline{N}}}\right) + \mathbb{E}_{\phi^{*},\psi^{\bar{N}_{2}}} \left[J_{\infty,t+1}^{\phi^{*},\tilde{\psi}_{\infty}^{*}}(\mathcal{M}_{t+1}^{\bar{N}_{1}},\mathcal{N}_{t+1}^{\bar{N}_{2}}) \right] \end{split}$$

$$-J_{\infty,t+1}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_{t+1}^{\phi^*},\nu_{t}^{\psi_{t}^{\bar{N}_{2}}}) + J_{\infty,t+1}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_{t+1}^{\phi^*},\nu_{t}^{\psi_{t}^{\bar{N}_{2}}}) \Big] \\ \stackrel{\text{(ii)}}{\geq} r_{t}(\mu_{t}^{\bar{N}_{1}},\nu_{t}^{\bar{N}_{2}}) + J_{\infty,t+1}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_{t+1}^{\phi^*},\nu_{apprx,t+1}^{\psi_{t}^{\bar{N}_{2}}}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) - L_{J,t+1}\underbrace{\mathbb{E}_{\phi^*}\Big[\mathrm{d}_{\mathrm{TV}}\big(\mathcal{M}_{t+1}^{\bar{N}_{1}},\mu_{t+1}^{\phi^*}\big)\Big]}_{=\mathcal{O}(\frac{1}{\sqrt{N_{1}}}) \text{ due to Corollary 1}}$$

$$-\underbrace{L_{J,t+1}\mathbb{E}_{\psi^{N_2}}\left[\mathrm{d}_{\mathrm{TV}}\big(\mathcal{N}_{t+1}^{N_2},\nu_{\mathrm{apprx},t+1}^{\psi_t^{N_2}}\big)\right]}_{=\mathcal{O}(\frac{1}{\sqrt{\tilde{N}_2}})\text{ due to(B.28)}}$$

$$\overset{\text{(iii)}}{\geq} r_{t}(\mu_{t}^{\bar{N}_{1}}, \nu_{t}^{\bar{N}_{2}}) + J_{\infty, t+1}^{\phi^{*}, \tilde{\psi}_{\infty}^{*}}(\mu_{t+1}^{\phi^{*}}, \nu_{\text{apprx}, t+1}^{\psi_{t}^{\bar{N}_{2}}}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big)$$

$$\overset{\text{(iv)}}{\geq} r_t(\mu_t^{\bar{N}_1}, \nu_t^{\bar{N}_2}) + \min_{\nu_{t+1} \in \mathbb{R}_{\nu,t}(\mu_t^{\bar{N}_1}, \nu_t^{\bar{N}_2})} J_{\infty, t+1}^{\phi^*, \tilde{\psi}_{\infty}^*}(\mu_{t+1}^{\phi^*}, \nu_{t+1}) - \mathcal{O}\Big(\frac{1}{\sqrt{\bar{N}}}\Big)$$

$$\stackrel{\text{(v)}}{=} J_{\infty,t}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu_t^{\bar{N}_1},\nu_t^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\bar{N}}}\Big).$$

For inequality (i), we used the inductive hypothesis; for inequality (ii), we utilized the Lipschitz continuity of the coordinator value function (Guan et al., 2024a) using mean-field distributions $(\mu_{t+1}^{\phi^*}, \nu_{\text{apprx},t+1}^{\psi_t^{N_2}})$ from (B.25) and the deterministic transitions from t to t+1; inequality (iv) follows from the definition of the min operator; and equality (v) follows from the definition of $J_{\infty,t}^{\phi^*,\tilde{\psi}^*_\infty}(\mu_t^{N_1}, \nu_t^{N_2})$ which completes the induction.

Since the Red team policy $\psi^{\bar{N}_2} \in \Psi^{\bar{N}_2}$ is arbitrary, we have that, for all joint states $\mathbf{x}^{\bar{N}_1} \in \mathbf{X}^{\bar{N}_1}$ and $\mathbf{v}^{\bar{N}_2} \in \mathbf{Y}^{\bar{N}_2}$.

$$\begin{split} \min_{\psi^{\bar{N}_2} \in \Psi^{\bar{N}_2}} J^{\bar{N}, \phi^*, \psi^{\bar{N}_2}}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) &= \min_{\psi^{\bar{N}_2} \in \Psi^{\bar{N}_2}} J_0^{\bar{N}, \phi^*, \psi^{\bar{N}_2}}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) \\ &\geq \underline{J}_{\infty}^{\phi^*, \tilde{\psi}_{\infty}^*}(\mu^{\bar{N}_1}, \nu^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\bar{N}}}\Big). \end{split}$$

Step 3. Combining (B.22) from Step 1 with Theorem 1, we have for all joint states $\mathbf{x}^{N_1} \in \mathcal{X}^{N_1}$ and $\mathbf{y}^{N_2} \in \mathcal{Y}^{N_2}$,

$$J_{\infty}^{\phi^*,\tilde{\psi}_{\infty}^*}(\mu^{N_1},\nu^{N_1}) \ge J^{N,\phi^*,\psi^{N_2*}}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$$
$$\ge \underline{J}^{N*}(\mathbf{x}^{N_1},\mathbf{y}^{N_2}) - \mathcal{O}\left(\frac{1}{\sqrt{N}}\right), \tag{B.29}$$

where $\mu^{N_1} = \operatorname{Emp}_{\mu}(\mathbf{x}^{N_1})$ and $\nu^{N_2} = \operatorname{Emp}_{\nu}(\mathbf{y}^{N_2})$.

Assuming that the initial distributions are the same for \mathcal{G}_1 and \mathcal{G}_2 , i.e., $\mu^{N_1} = \mu^{\bar{N}_1}$ and $\nu^{N_2} = \nu^{\bar{N}_2}$ we have the following sequence of inequalities

$$\begin{split} \min_{\boldsymbol{\psi}^{\bar{N}_2} \in \boldsymbol{\Psi}^{\bar{N}_2}} J^{\bar{N}, \phi^*, \boldsymbol{\psi}^{\bar{N}_2}}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) &\overset{\text{(i)}}{\geq} J^{\phi^*, \tilde{\psi}^*_{\infty}}_{\infty}(\boldsymbol{\mu}^{\bar{N}_1}, \boldsymbol{\nu}^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) \\ &\overset{\text{(ii)}}{\geq} \underline{J}^{N*}(\mathbf{x}^{N_1}, \mathbf{y}^{N_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) \\ &\overset{\text{(iii)}}{\geq} \underline{J}^{\phi^*_{\infty}, \psi^*_{\infty}}_{\infty}(\boldsymbol{\mu}^{N_1}, \boldsymbol{\nu}^{N_1}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big) \\ &\overset{\text{(iv)}}{\geq} \underline{J}^{\bar{N}*}(\mathbf{x}^{\bar{N}_1}, \mathbf{y}^{\bar{N}_2}) - \mathcal{O}\Big(\frac{1}{\sqrt{\underline{N}}}\Big), \end{split}$$

where $J_{\infty}^{\phi_{\infty}^*,\psi_{\infty}^*}(\mu_0,\nu_0) = \max_{\phi\in\Phi}\min_{\psi\in\Psi}J_{\infty}^{\phi,\psi}(\mu_0,\nu_0)$. Note that (i) follows from Step 2 (B.26); (ii) is a consequence of (B.29) and inequalities (iii) and (iv) follow from Lemmas 5 (applied on game \mathcal{G}_1) and 6 (applied on game \mathcal{G}_2) in Guan et al. (2024b) respectively, combined with the fact that $\min(\bar{N}_1,\bar{N}_2) > \min(N_1,N_2)$, completing the proof.

Proposition 2. If the log-probability of the Blue team policy is bounded such that for all $x \in \mathcal{X}, u \in \mathcal{U}, \mu \in \mathcal{P}(\mathcal{X}), \nu \in \mathcal{P}(\mathcal{Y}), \|\nabla_{\eta} \log \phi(u \mid x, \mu, \nu)\|_{2} \leq L_{\phi}/2|\mathcal{U}|$, where the gradient is taken with respect to $\eta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y}) \triangleq [\mu^{\mathsf{T}}, \nu^{\mathsf{T}}]^{\mathsf{T}}$ and $L_{\phi} > 0$, then $\phi(u \mid x, \mu, \nu)$ is Lipschitz continuous with Lipschitz constant L_{ϕ} , i.e.,

$$\sum_{u} |\phi_t(u|x,\hat{\mu},\hat{\nu}) - \phi_t(u|x,\hat{\mu}',\hat{\nu}')| \le L_\phi\left(\mathrm{d_{TV}}(\hat{\mu},\hat{\mu}') + \mathrm{d_{TV}}(\hat{\nu},\hat{\nu}')\right) \quad \forall \ x \in \mathcal{X}.$$
 (9)

Proof. We prove the proposition in two parts.

Step 1. We first show that $\|\nabla_{\eta} \log \phi(u \mid x, \mu, \nu)\|_2 \leq \frac{L_{\phi}}{2|\mathcal{U}|}$ implies

$$\sum_{u} |\log \phi_{t}(u|x, \hat{\mu}, \hat{\nu}) - \log \phi_{t}(u|x, \hat{\mu}', \hat{\nu}')| \leq L_{\phi} \left(d_{\text{TV}}(\hat{\mu}, \hat{\mu}') + d_{\text{TV}}(\hat{\nu}, \hat{\nu}') \right) \quad \forall \quad x \in \mathcal{X}.$$
(B.30)

By the Definition of η , we can denote $\log \phi_t(u|x,\eta) \triangleq \log \phi_t(u|x,\hat{\mu},\hat{\nu})$. By the Mean-Value Theorem and the bounded gradient from Step 1 we have for all $x \in \mathcal{X}$, $u \in \mathcal{U}$, $\eta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})$,

$$\begin{aligned} |\log \phi_{t}\left(u|x,\eta\right) - \log \phi_{t}\left(u|x,\eta'\right)| &\leq \sup_{\eta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})} \left\|\nabla_{\eta} \log \phi(u \mid x,\eta)\right\|_{2} \left\|\eta - \eta'\right\|_{2} \\ &\leq \sup_{\eta \in \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{Y})} \left\|\nabla_{\eta} \log \phi(u \mid x,\eta)\right\|_{2} \left\|\eta - \eta'\right\|_{1} \\ &\leq \frac{L_{\phi}}{2|\mathcal{U}|} \cdot 2\left(\mathrm{d}_{\mathrm{TV}}\left(\hat{\mu},\hat{\mu}'\right) + \mathrm{d}_{\mathrm{TV}}\left(\hat{\nu},\hat{\nu}'\right)\right) \\ &= \frac{L_{\phi}}{|\mathcal{U}|} \left(\mathrm{d}_{\mathrm{TV}}\left(\hat{\mu},\hat{\mu}'\right) + \mathrm{d}_{\mathrm{TV}}\left(\hat{\nu},\hat{\nu}'\right)\right). \end{aligned}$$

Summing over all actions for any given $x \in \mathcal{X}$,

$$\sum_{u} |\log \phi_{t}(u|x,\eta) - \log \phi_{t}(u|x,\eta')| \leq \sum_{u} \frac{L_{\phi}}{|\mathcal{U}|} \left(d_{\text{TV}}(\hat{\mu}, \hat{\mu}') + d_{\text{TV}}(\hat{\nu}, \hat{\nu}') \right)
\leq L_{\phi} \left(d_{\text{TV}}(\hat{\mu}, \hat{\mu}') + d_{\text{TV}}(\hat{\nu}, \hat{\nu}') \right) \sum_{u} \frac{1}{|\mathcal{U}|}
\leq L_{\phi} \left(d_{\text{TV}}(\hat{\mu}, \hat{\mu}') + d_{\text{TV}}(\hat{\nu}, \hat{\nu}') \right),$$

completing the proof of Part 1.

 Step 2. Using Step 1, we show that (B.30) implies (9). Define $h(z) \triangleq e^z$ and values $a \triangleq \log \phi_t(u|x,\hat{\mu}',\hat{\nu}')$ and $b \triangleq \log \phi_t(u|x,\hat{\mu},\hat{\nu})$. It follows from the Mean-Value Theorem that for all $x \in \mathcal{X}$ and $u \in \mathcal{U}$,

$$|h(b) - h(a)| = |h'(c)||b - a|, \min(a, b) \le c \le \max(a, b)$$

$$|\phi_t(u|x, \hat{\mu}, \hat{\nu}) - \phi_t(u|x, \hat{\mu}', \hat{\nu}')| = e^c |\log \phi_t(u|x, \hat{\mu}, \hat{\nu}) - \log \phi_t(u|x, \hat{\mu}', \hat{\nu}')|$$

$$\le e^{\max(a, b)} |\log \phi_t(u|x, \hat{\mu}, \hat{\nu}) - \log \phi_t(u|x, \hat{\mu}', \hat{\nu}')|.$$

Note that $e^a = \phi_t(u|x, \hat{\mu}', \hat{\nu}') \le 1$ and $e^b = \phi_t(u|x, \hat{\mu}, \hat{\nu}) \le 1$. Consequently $e^{\max(a,b)} \le 1$ and we have for all $x \in \mathcal{X}$.

$$\begin{aligned} &|\phi_t\left(u|x,\hat{\mu},\hat{\nu}\right) - \phi_t\left(u|x,\hat{\mu}',\hat{\nu}'\right)| \leq |\log\phi_t\left(u|x,\hat{\mu},\hat{\nu}\right) - \log\phi_t\left(u|x,\hat{\mu}',\hat{\nu}'\right)| \\ &\sum_{u} |\phi_t\left(u|x,\hat{\mu},\hat{\nu}\right) - \phi_t\left(u|x,\hat{\mu}',\hat{\nu}'\right)| \leq \sum_{u} |\log\phi_t\left(u|x,\hat{\mu},\hat{\nu}\right) - \log\phi_t\left(u|x,\hat{\mu}',\hat{\nu}'\right)| \\ &\sum_{u} |\phi_t\left(u|x,\hat{\mu},\hat{\nu}\right) - \phi_t\left(u|x,\hat{\mu}',\hat{\nu}'\right)| \leq L_\phi\left(\mathrm{d}_{\mathrm{TV}}\big(\hat{\mu},\hat{\mu}'\big) + \mathrm{d}_{\mathrm{TV}}\big(\hat{\nu},\hat{\nu}'\big)\big), \end{aligned}$$

where the last inequality follows from (B.30)

Proposition 3. Consider a given ϵ -accurate estimator, i.e., $d_{TV}(\hat{\nu}_{i,t}^{N_2}, \hat{\nu}_t^{N_2}) < \epsilon$, for all i,t, where $\hat{\nu}_t^{N_2}$ is the true opponent MF at time t and $\epsilon > 0$. For the identical team-policy pair (ϕ_t^*, ψ_t^*) obtained via gradient-regularized MF-MAPPO and deployed using this estimator, the cumulative regret satisfies $\Delta J(\phi_t^*, \psi_t^*) \leq K\epsilon + \mathcal{O}(1/\sqrt{N})$ for some constant K > 0.

Sketch of Proof. MFs induced by identical team policies in an infinite-population game closely approximate the EDs induced in the corresponding finite-population game. Combining this consequence with Lipschitz continuity in the dynamics, rewards (Assumption 1) and policies (Proposition 2), we prove the statement.

Proof. For the performance guarantee, from (10), we have,

$$\Delta J(\phi_t^*, \psi_t^*) = \mathbb{E}_{\phi^*, \psi^*} \left[\left| \sum_{t=0}^T r_t(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}) - \sum_{t=0}^T r_t(\hat{\mathcal{M}}_t^{N_1}, \hat{\mathcal{N}}_t^{N_2}) \right| \left| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \right].$$

From Lipschitz rewards (Assumption 1), we have,

$$\Delta J(\phi_t^*, \psi_t^*; \Gamma)
= \mathbb{E}_{\phi^*, \psi^*} \left[\left| \sum_{t=0}^T r_t(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}) - \sum_{t=0}^T r_t(\hat{\mathcal{M}}_t^{N_1}, \hat{\mathcal{N}}_t^{N_2}) \right| \left| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \right] \right]
\leq \mathbb{E}_{\phi^*, \psi^*} \left[\sum_{t=0}^T L_r \left(d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right) \left| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \right] \right]
\leq L_r \sum_{t=0}^T \mathbb{E}_{\phi^*, \psi^*} \left[d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \right], \quad (B.31)$$

where the last inequality follows from the linearity of expectations.

Step 1. We first show the following for all time steps t.

$$\mathbb{E}_{\phi^*,\psi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right]$$

$$\leq \kappa_1 + \kappa_2 \left(d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right),$$

where
$$\kappa_1 \triangleq \frac{1}{2}(L_\phi + L_\psi)\epsilon + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$$
 and $\kappa_2 \triangleq (1 + L_f + \frac{1}{2}L_\phi + \frac{1}{2}L_\psi)$.

By the linearity property of expectations,

$$\mathbb{E}_{\phi^*,\psi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] + \mathbb{E}_{\psi^*} \left[d_{\text{TV}} \left(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right) \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \right] \\
= \mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+$$

We now bound the Blue team's mean-field, i.e., the first term on the RHS. The Red team can be bounded similarly (second term).

$$\mathbb{E}_{\phi^{*}} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_{1}}, \hat{\mathcal{M}}_{t+1}^{N_{1}} \right) \middle| \mathbf{x}_{t}^{N_{1}}, \hat{\mathbf{x}}_{t}^{N_{1}}, \mathbf{y}_{t}^{N_{2}}, \hat{\mathbf{y}}_{t}^{N_{2}} \right] \\
\leq \mathbb{E}_{\phi^{*}} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_{1}}, \mathcal{M}_{t+1} \right) \middle| \mathbf{x}_{t}^{N_{1}}, \hat{\mathbf{x}}_{t}^{N_{1}}, \mathbf{y}_{t}^{N_{2}}, \hat{\mathbf{y}}_{t}^{N_{2}} \right] \\
+ \mathbb{E}_{\phi^{*}} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}, \hat{\mathcal{M}}_{t+1} \right) \middle| \mathbf{x}_{t}^{N_{1}}, \hat{\mathbf{x}}_{t}^{N_{1}}, \mathbf{y}_{t}^{N_{2}}, \hat{\mathbf{y}}_{t}^{N_{2}} \right] \\
+ \mathbb{E}_{\phi^{*}} \left[d_{\text{TV}} \left(\hat{\mathcal{M}}_{t+1}^{N_{1}}, \hat{\mathcal{M}}_{t+1} \right) \middle| \mathbf{x}_{t}^{N_{1}}, \hat{\mathbf{x}}_{t}^{N_{1}}, \mathbf{y}_{t}^{N_{2}}, \hat{\mathbf{y}}_{t}^{N_{2}} \right], \tag{B.32}$$

where $\mathcal{M}_{t+1} = \mathcal{M}_t^{N_1} F(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}, \phi^*)$ and $\hat{\mathcal{M}}_{t+1} = \hat{\mathcal{M}}_t^{N_1} F(\hat{\mathcal{M}}_t^{N_1}, \hat{\mathcal{N}}_t^{N_2}, \phi^*, \Gamma_{\text{D-PC}})$ are short-hand notations for the next induced mean-field from the *infinite*-population *deterministic* dynamics (Guan et al., 2024a) with and without the estimator $\Gamma_{\text{D-PC}}$. From Corollary 1 in Guan et al. (2024a), we have

$$\mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \mathcal{M}_{t+1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \leq \mathcal{O} \left(\frac{1}{\sqrt{N}} \right),$$

$$\mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\hat{\mathcal{M}}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \leq \mathcal{O} \left(\frac{1}{\sqrt{N}} \right).$$

Thus, we are left to simplify

$$\mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}, \hat{\mathcal{M}}_{t+1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right]$$

$$= d_{\text{TV}} \left(\mathcal{M}_t^{N_1} F(\mathcal{M}_t^{N_1}, \mathcal{N}_t^{N_2}, \phi^*), \hat{\mathcal{M}}_t^{N_1} F(\hat{\mathcal{M}}_t^{N_1}, \hat{\mathcal{N}}_t^{N_2}, \phi^*, \Gamma_{\text{D-PC}}) \right),$$

where the expectation vanishes due to *deterministic* transitions under identical policies (ϕ^*, ψ^*) . Now,

$$2d_{\text{TV}}\left(\mathcal{M}_{t}^{N_{1}}F(\mathcal{M}_{t}^{N_{1}},\mathcal{N}_{t}^{N_{2}},\phi^{*}),\hat{\mathcal{M}}_{t}^{N_{1}}F(\hat{\mathcal{M}}_{t}^{N_{1}},\hat{\mathcal{N}}_{t}^{N_{2}},\phi^{*},\Gamma_{\text{D-PC}})\right)$$
(B.33)
$$=\sum_{x'\in\mathcal{X}}\left|\sum_{x\in\mathcal{X}}\sum_{u\in\mathcal{U}}f_{t}(x'|x,u,\mathcal{M}_{t}^{N_{1}},\mathcal{N}_{t}^{N_{2}})\phi_{t}^{*}(u|x,\mathcal{M}_{t}^{N_{1}},\mathcal{N}_{t}^{N_{2}})\mathcal{M}_{t}^{N_{1}}(x)\right.$$

$$-\sum_{x\in\mathcal{X}}\sum_{u\in\mathcal{U}}f_{t}(x'|x,u,\hat{\mathcal{M}}_{t}^{N_{1}},\hat{\mathcal{N}}_{t}^{N_{2}})\phi_{t}^{*}(u|x,\hat{\mathcal{M}}_{t}^{N_{1}},\hat{\mathcal{N}}_{x,t}^{N_{2}})\hat{\mathcal{M}}_{t}^{N_{1}}(x)\right|$$

$$\leq\sum_{x'\in\mathcal{X}}\sum_{x\in\mathcal{X}}\sum_{u\in\mathcal{U}}\left|f_{t}(x'|x,u,\mathcal{M}_{t}^{N_{1}},\mathcal{N}_{t}^{N_{2}})\phi_{t}^{*}(u|x,\mathcal{M}_{t}^{N_{1}},\mathcal{N}_{t}^{N_{2}})\mathcal{M}_{t}^{N_{1}}(x)\right.$$

$$-f_{t}(x'|x,u,\hat{\mathcal{M}}_{t}^{N_{1}},\hat{\mathcal{N}}_{t}^{N_{2}})\phi_{t}^{*}(u|x,\hat{\mathcal{M}}_{t}^{N_{1}},\hat{\mathcal{N}}_{x,t}^{N_{2}})\hat{\mathcal{M}}_{t}^{N_{1}}(x)\right|.$$

Firstly, we add and subtract $f_t(x'|x,u,\hat{\mathcal{M}}_t^{N_1},\hat{\mathcal{N}}_t^{N_2})\phi_t^*(u|x,\hat{\mathcal{M}}_t^{N_1},\hat{\mathcal{N}}_t^{N_2})\hat{\mathcal{M}}_t^{N_1}(x)$ to split absolute value as

$$2d_{\text{TV}}\big(\mathcal{M}_{t}^{N_{1}}F(\mathcal{M}_{t}^{N_{1}}, \mathcal{N}_{t}^{N_{2}}, \phi^{*}), \hat{\mathcal{M}}_{t}^{N_{1}}F(\hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}, \phi^{*}, \Gamma_{\text{D-PC}})\big) \leq \sum_{x' \in \mathcal{X}} \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} (I_{1} + I_{2}),$$

where,

1459
1460
$$I_{1} \triangleq \left| f_{t}(x'|x, u, \mathcal{M}_{t}^{N_{1}}, \mathcal{N}_{t}^{N_{2}}) \phi_{t}^{*}(u|x, \mathcal{M}_{t}^{N_{1}}, \mathcal{N}_{t}^{N_{2}}) \mathcal{M}_{t}^{N_{1}}(x) - f_{t}(x'|x, u, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}) \phi_{t}^{*}(u|x, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}) \hat{\mathcal{M}}_{t}^{N_{1}}(x) \right|$$
1461
1462
$$I_{2} \triangleq \left| f_{t}(x'|x, u, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}) \phi_{t}^{*}(u|x, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}) \hat{\mathcal{M}}_{t}^{N_{1}}(x) - f_{t}(x'|x, u, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{t}^{N_{2}}) \phi_{t}^{*}(u|x, \hat{\mathcal{M}}_{t}^{N_{1}}, \hat{\mathcal{N}}_{x,t}^{N_{2}}) \hat{\mathcal{M}}_{t}^{N_{1}}(x) \right|.$$

Similar to (9), we can simplify I_2 using Proposition 2 and obtain

$$\sum_{x' \in \mathcal{X}} \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} I_2 < L_{\phi} \epsilon.$$

We now individually bound the terms in I_1 using the definition of total variation distance, the identity

$$|abc - a'b'c'| \le |a - a'|bc + a'|b - b'|c + a'b'|c - c'|, \ a, b, c \ge 0,$$

Assumption 1 and Proposition 2 as follows:

$$\sum_{x' \in \mathcal{X}} \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} I_1 \leq 2 d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + \left(L_{\phi} + L_f \right) \left(d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right)$$

Thus we can substitute the expressions for I_1 and I_2 back in (B.33) and rewrite (B.32) to get

$$\mathbb{E}_{\phi^*} \left[d_{\text{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\
< d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) \\
+ \frac{1}{2} (L_{\phi} + L_f) \left(d_{\text{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + d_{\text{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right) + \frac{1}{2} L_{\phi} \epsilon + \mathcal{O} \left(\frac{1}{\sqrt{N}} \right) \tag{B.34}$$

By symmetry,

$$\mathbb{E}_{\psi^{*}}\left[d_{\text{TV}}\left(\mathcal{N}_{t+1}^{N_{2}}, \hat{\mathcal{N}}_{t+1}^{N_{2}}\right) \middle| \mathbf{x}_{t}^{N_{1}}, \hat{\mathbf{x}}_{t}^{N_{1}}, \mathbf{y}_{t}^{N_{2}}, \hat{\mathbf{y}}_{t}^{N_{2}}\right] \\
< d_{\text{TV}}\left(\mathcal{N}_{t}^{N_{2}}, \hat{\mathcal{N}}_{t}^{N_{2}}\right) \\
+ \frac{1}{2}(L_{\psi} + L_{f})\left(d_{\text{TV}}\left(\mathcal{M}_{t}^{N_{1}}, \hat{\mathcal{M}}_{t}^{N_{1}}\right) + d_{\text{TV}}\left(\mathcal{N}_{t}^{N_{2}}, \hat{\mathcal{N}}_{t}^{N_{2}}\right)\right) + \frac{1}{2}L_{\psi}\epsilon + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \tag{B.35}$$

Adding (B.34) and (B.35) we obtain,

$$\begin{split} &\mathbb{E}_{\phi^*,\psi^*} \left[\mathrm{d}_{\mathrm{TV}} \left(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \right) + \mathrm{d}_{\mathrm{TV}} \left(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \right) \middle| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \right] \\ &< (1 + L_f + \frac{1}{2} L_\phi + \frac{1}{2} L_\psi) \left(\mathrm{d}_{\mathrm{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + \mathrm{d}_{\mathrm{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right) + \frac{1}{2} (L_\phi + L_\psi) \epsilon + \mathcal{O} \left(\frac{1}{\sqrt{\underline{N}}} \right) \\ &= \kappa_1 + \kappa_2 \left(\mathrm{d}_{\mathrm{TV}} \left(\mathcal{M}_t^{N_1}(x), \hat{\mathcal{M}}_t^{N_1}(x) \right) + \mathrm{d}_{\mathrm{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \right). \end{split}$$

Step 2. Define

$$a_t \triangleq \mathbb{E}_{\phi^*, \psi^*} \left[\mathrm{d}_{\mathrm{TV}} \left(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \right) + \mathrm{d}_{\mathrm{TV}} \left(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \right) \middle| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \right]$$

At t=0, $\mathcal{M}_0^{N_1}=\hat{\mathcal{M}}_0^{N_1}$ and $\mathcal{N}_0^{N_2}=\hat{\mathcal{N}}_0^{N_2}$ as we begin the fully observable and estimated scenarios under the same initial joint states , i.e., $\mathbf{x}_0^{N_1}=\hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2}=\hat{\mathbf{y}}_0^{N_2}$. Thus, $a_0=0$. We proceed to show that $a_{t+1}\leq \kappa_1+\kappa_2 a_t$ where κ_1 and κ_2 are the same as Step 1. By the law of iterated expectations,

$$\begin{split} & \mathbb{E}_{\phi^*,\psi^*} \Big[\mathrm{d}_{\mathrm{TV}} \big(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \big) + \mathrm{d}_{\mathrm{TV}} \big(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \big) \Big| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \Big] \\ & = \mathbb{E}_{\phi^*,\psi^*} \left[\Big[\mathbb{E}_{\phi^*,\psi^*} \mathrm{d}_{\mathrm{TV}} \big(\mathcal{M}_{t+1}^{N_1}, \hat{\mathcal{M}}_{t+1}^{N_1} \big) + \mathrm{d}_{\mathrm{TV}} \big(\mathcal{N}_{t+1}^{N_2}, \hat{\mathcal{N}}_{t+1}^{N_2} \big) \Big| \mathbf{x}_t^{N_1}, \hat{\mathbf{x}}_t^{N_1}, \mathbf{y}_t^{N_2}, \hat{\mathbf{y}}_t^{N_2} \Big] \Big| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \Big] \\ & \leq \mathbb{E}_{\phi^*,\psi^*} \left[\kappa_1 + \kappa_2 \left(\mathrm{d}_{\mathrm{TV}} \big(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \big) + \mathrm{d}_{\mathrm{TV}} \big(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \big) \right) \Big| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \Big] \\ & \leq \mathbb{E}_{\phi^*,\psi^*} \left[\mathrm{d}_{\mathrm{TV}} \big(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \big) + \mathrm{d}_{\mathrm{TV}} \big(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \big) \Big| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \Big] \\ & = \kappa_1 + \kappa_2 \mathbb{E}_{\phi^*,\psi^*} \left[\mathrm{d}_{\mathrm{TV}} \big(\mathcal{M}_t^{N_1}, \hat{\mathcal{M}}_t^{N_1} \big) + \mathrm{d}_{\mathrm{TV}} \big(\mathcal{N}_t^{N_2}, \hat{\mathcal{N}}_t^{N_2} \big) \Big| \mathbf{x}_0^{N_1} = \hat{\mathbf{x}}_0^{N_1}, \mathbf{y}_0^{N_2} = \hat{\mathbf{y}}_0^{N_2} \Big] \\ & = \kappa_1 + \kappa_2 a_t, \end{aligned}$$

where we use the result obtained from Step 1 along with the linearity property of expectations. Using this new notation, (B.31) simplifies to

$$\Delta J(\phi_t^*, \psi_t^*) \le L_r \sum_{t=0}^T a_t$$

We have $a_0=0$, $a_1<\kappa_1$, $a_2<\kappa_1+\kappa_2a_1$ or $a_2<\kappa_1(\kappa_2+1)$, $a_3<\kappa_2a_2+\kappa_1$ or $a_3<\kappa_1(\kappa_2^2+\kappa_2+1)$ and so on. This can be written compactly for all t>0 as

$$a_t < \kappa_1 \sum_{\tau=0}^{t-1} \kappa_2^{\tau} = \kappa_1 \frac{\kappa_2^t - 1}{\kappa_2 - 1},$$

as $\kappa_2 = (1 + L_f + \frac{1}{2}L_\phi + \frac{1}{2}L_\psi) \neq 1$. This is a geometric sum and thus we have

$$\begin{split} \Delta J(\phi_t^*, \psi_t^*) &\leq L_r \sum_{t=0}^T a_t \\ &= L_r \sum_{t=1}^T a_t \\ &< L_r \sum_{t=1}^T \kappa_1 \frac{\kappa_2^t - 1}{\kappa_2 - 1} \\ &= \frac{L_r \kappa_1}{\kappa_2 - 1} \sum_{t=1}^T \kappa_2^t - 1 \\ &= \frac{L_r \kappa_1}{\kappa_2 - 1} \left(\kappa_2 \frac{\kappa_2^T - 1}{\kappa_2 - 1} - T \right) \\ &= \underbrace{\frac{1}{2} \frac{L_r (L_\phi + L_\psi)}{\kappa_2 - 1} \left(\kappa_2 \frac{\kappa_2^T - 1}{\kappa_2 - 1} - T \right)}_{K} \epsilon + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \frac{L_r}{\kappa_2 - 1} \left(\kappa_2 \frac{\kappa_2^T - 1}{\kappa_2 - 1} - T \right) \\ &= K \epsilon + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right). \end{split}$$

Theorem 4. *D-PC satisfies Proposition 3 with* $\epsilon = \mathcal{O}(e^{-cR_{\text{com}}})$ *with* c > 0.

Sketch of Proof. The exponential convergence of D-PC at each time-step t follows directly from Theorems 6-8 in Nedić & Liu (2016) as the number of communication rounds $R_{\rm com} \to \infty$. Therefore, for finite communication rounds, there exist constants $c_x, c_y > 0$ and $\rho_x, \rho_y \in (0,1)$ such that for any number of communication rounds $R_{\rm com}$ and any underlying distributions $(\hat{\mu}_t^{N_1}, \hat{\nu}_t^{N_2})$ D-PC satisfies

$$d_{\text{TV}}(\hat{\nu}_{x,t}^{N_2}, \hat{\nu}_t^{N_2}) < c_x \rho_x^{R_{\text{com}}} \triangleq \epsilon_x, \quad \text{for all } x \in \mathcal{X}^o \text{ and } t$$

$$d_{\text{TV}}(\hat{\mu}_{y,t}^{N_1}, \hat{\mu}_t^{N_1}) < c_y \rho_y^{R_{\text{com}}} \triangleq \epsilon_y, \quad \text{for all } y \in \mathcal{Y}^o \text{ and } t.$$
(B.36)

We define $\epsilon = \max(\epsilon_x, \epsilon_y)$.

Proof. We begin by defining set regularity.

Definition 6 (Set Regularity (Nedić & Liu, 2016)). Let $Z \subseteq \mathbb{R}^n$ be a nonempty set. A (finite) collection of closed convex sets $\{Y_j\}_{j=1}^J \subseteq \mathbb{R}^n$ is regular (in Euclidian norm) with respect to the set Z, if there exists a constant $r \geq 1$ such that

$$\inf_{y \in Y} \|z - y\|_2 \le r \max_{i} \inf_{y \in Y_i} \|z - y\|_2, z \in Z,$$

where $Y \triangleq \bigcap_{j=1,...,J} Y_j$ is non-empty. We refer to the scalar r as the regularity constant. When the preceding relation holds with $Z = \mathbb{R}^n$, we say that the sets $\{Y_j\}_{j=1}^J$ are uniformly regular.

Proposition 5. The constraint sets $\mathcal{R}(x)$ for all $x \in \mathcal{X}^o$ are uniformly regular with regularity constant $\kappa \geq 1$, i.e.,

$$\inf_{\nu \in \mathcal{F}} \|\omega - \nu\|_2 \le \kappa \max_{x \in \mathcal{X}} \inf_{\nu_x \in \mathcal{R}(x)} \|\omega - \nu_x\|_2, \omega \in \mathbb{R}^{|\mathcal{X}|}.$$

Proof. Refer to Proposition 1 from Nedić & Liu (2016).

Following Theorem 8 from Nedić & Liu (2016), we have

$$\rho_x = \left(1 - \frac{\theta^2}{4|\mathcal{X}^o|\operatorname{diam}(\mathcal{G}_{\mathrm{Blue}}^{\mathrm{com}})(\kappa+1)^2}\right), \quad c_x = \frac{1}{2}\sqrt{|\mathcal{Y}|} \sum_{x \in \mathcal{X}^o} \|\hat{\nu}_x^{\tau=0} - \xi_x^{\tau=1}\|_2^2$$
$$\rho_y = \left(1 - \frac{\theta^2}{4|\mathcal{Y}^o|\operatorname{diam}(\mathcal{G}_{\mathrm{Red}}^{\mathrm{com}})(\kappa+1)^2}\right), \quad c_y = \frac{1}{2}\sqrt{|\mathcal{X}|} \sum_{x \in \mathcal{Y}^o} \|\hat{\mu}_y^{\tau=0} - \xi_y^{\tau=1}\|_2^2,$$

where κ is the regularity constant as in Proposition 5 and θ follows from (A.3).

C RPS AND CRPS SETUP

C.1 STATE SPACE

 We have three states in this representation of the game: rock, paper and scissors. We denote this state space as $\mathcal{S} = \{\mathbb{R}, \mathbb{P}, \mathbb{S}\}$. The empirical distribution of the Blue team is denoted by $\mu \in \mathcal{P}(\mathcal{S})$ and that of the Red team is denoted by $\nu \in \mathcal{P}(\mathcal{S})$. Since we have three states for each team, both EDs lie in a three-dimensional simplex denoted by $\mathcal{P}(\mathcal{S})$.

C.2 ACTION SPACE

C.2.1 RPS

At each state, we define three actions denoted by $A = \{CW, CCW, Stay\}$. These actions represent the ability of the agents to move from one state to the other in the following fashion:

- 1. CW denotes a clockwise cyclic action from one state to the other, i.e., from $R \to P$, $P \to S$, $S \to R$.
- 2. CCW denotes a counterclockwise cyclic movement, i.e., from $R \to S$, $S \to P$, $P \to R$.
- 3. Stay denotes the idle action (remain in the same state as before).

C.2.2 CRPS

At each state we have two actions denoted by $A = \{CW, Stay\}$. These actions represent the ability of the agents to move from one state to the other in the following fashion:

- 1. CW denotes a clockwise cyclic action from one state to the other, i.e., from $R \to P$, $P \to S$, $S \to R$.
- 2. Stay denotes the idle action (remain in the same state as before).

Thus, we cannot directly jump from R to S within a single step, but must go via P. Mathematically, S does not lie in the reachable set of R. The reachable set $\mathcal{R}(s)$ for each state s at a given time step under this modified action space is as follows

$$\mathcal{R}(R) = \{R, P\}, \quad \mathcal{R}(P) = \{P, S\}, \quad \mathcal{R}(S) = \{S, R\}.$$

The state-action space of both RPS and cRPS are presented in Figure C.1.

C.3 DYNAMICS AND TRANSITION PROBABILITIES

For both RPS and cRPS, we consider deterministic transitions $\mathcal{T}(s,a,s')$, which implies that given a state-action pair (s,a), the agent reaches a unique next state s' with certainty (no distribution over the reachable states). Thus, for state R and action Stay, the transition function $\mathcal{T}: \mathcal{S} \times \mathcal{A} \times \mathcal{S} \to \{0,1\}$ is given as:

$$\mathcal{T}(R, Stay, R) = 1$$
, $\mathcal{T}(R, Stay, P) = 0$, $\mathcal{T}(R, Stay, S) = 0$.

This implies that an agent in the state R upon taking the Stay action remains in state R. Similarly,

$$\mathcal{T}(R, CW, R) = 0, \quad \mathcal{T}(R, CW, P) = 1, \quad \mathcal{T}(R, CW, S) = 0.$$

represent the method to transition from R to P.

C.4 REWARD STRUCTURE

 In a two player RPS game, the reward matrix for Player 1 is defined as:

$$A = \left[\begin{array}{rrr} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{array} \right].$$

We extend this two-player framework to the multi-agent team game formulation. Define the pairwise reward for an agent at state $x \in \mathcal{S}$ within the Blue team and at state $y \in \mathcal{S}$ from the Red team as

$$r(x,y) \triangleq A_{xy}$$

where A_{xy} represents the element from the reward matrix A corresponding to the states x (row player) and y (column player). In lieu of the zero-sum structure, the reward for the agent at y with respect to x becomes -r(x,y). Thus, for each player $x_i \in \mathcal{S}$ and $i=1,2,\ldots,N_1$ in the Blue team and $y_j \in \mathcal{S}$ and $j=1,2,\ldots,N_2$ in the Red team, the reward for the Blue team can be defined as

$$R_{\mathrm{Blue}}(\mathbf{x}, \mathbf{y}) = \frac{1}{N_1} \sum_{i=1}^{N_1} \underbrace{\left[\frac{1}{N_2} \sum_{j=1}^{N_2} r(x_i, y_j) \right]}_{\text{reward for agent } i}.$$

Rewriting the term inside the square brackets as

$$\frac{1}{N_2} \sum_{j=1}^{N_2} r(x_i, y_j) = \frac{1}{N_2} \sum_{y \in \mathcal{S}} \sum_{j=1}^{N_2} r(x_i, y) \mathbf{1}_{y_j = y}$$

$$= A_{x_i s_0} \sum_{j=1}^{N_2} \frac{1}{N_2} \mathbf{1}_{y_j = s_0} + A_{x_i s_1} \sum_{j=1}^{N_2} \frac{1}{N_2} \mathbf{1}_{y_j = s_1}$$

$$+ A_{x_i s_2} \sum_{j=1}^{N_2} \frac{1}{N_2} \mathbf{1}_{y_j = s_2}$$

$$= A_{x_i s_0} \nu(s_0) + A_{x_i s_1} \nu(s_1) + A_{x_i s_2} \nu(s_2)$$

$$= A(x_i) \nu, \tag{C.37}$$

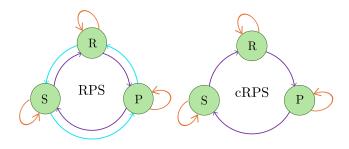


Figure C.1: State-action spaces for RPS and cRPS.

where $A(x_i)$ is the row of the reward matrix corresponding to state x_i . Using (C.37), the total Blue reward can be expressed as

$$R_{\text{Blue}}(\mathbf{x}, \mathbf{y}) = \frac{1}{N_1} \sum_{i=1}^{N_1} A(x_i) \nu$$

$$= \left(\frac{1}{N_1} \sum_{x \in \mathcal{S}} \sum_{i=1}^{N_1} A(x_i) \mathbf{1}_{x_i = x}\right) \nu$$

$$= \left(A(s_0) \mu(s_0) + A(s_1) \mu(s_1) + A(s_2) \mu(s_2)\right) \nu$$

$$= \mu^{\mathsf{T}} A \nu.$$

Proposition 6. With initial conditions $\mu_{t=0} = [1,0,0]^{\mathsf{T}}$ and $\nu_{t=0} = [0,1,0]^{\mathsf{T}}$, all mean-field optimal trajectories satisfy $\mu_t^* = \nu_t^* = [\frac{1}{3},\frac{1}{3},\frac{1}{3}]^{\mathsf{T}}$ for all $t \geq 2$, and $\mu_1^* = [0,1-\eta,\eta]^{\mathsf{T}}$ where $\eta \in [\frac{1}{3},\frac{2}{3}]$ and $\nu_1^* = [0,\frac{2}{3},\frac{1}{3}]^{\mathsf{T}}$. Furthermore, the unique game value is given by $-\frac{1}{3}$.

Proof. For the constrained RPS game under the stated initial condition, we cannot obtain the target distribution $\begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}^T$ after a single time step but this may be possible for $t \geq 2$. To this end, consider the following candidate trajectory respecting the transition dynamics:

$$\begin{split} & \mu_0^{\mathsf{T}} = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix}, \\ & \mu_1^{\mathsf{T}} = \begin{bmatrix} 1 - x_1 & x_1 & 0 \end{bmatrix}, \\ & \mu_2^{\mathsf{T}} = \begin{bmatrix} 1 - x_1 - x_2 & x_1 + x_2 - x_3 & x_3 \end{bmatrix}, \\ & \mu_t^{\mathsf{T}} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}, \quad \forall t > 2. \\ & \nu_0^{\mathsf{T}} = \begin{bmatrix} 0 & 1 & 0 \end{bmatrix}, \\ & \nu_1^{\mathsf{T}} = \begin{bmatrix} 0 & 1 - y_1 & y_1 \end{bmatrix}, \\ & \nu_2^{\mathsf{T}} = \begin{bmatrix} y_3 & 1 - y_1 - y_2 & y_1 + y_2 - y_3 \end{bmatrix}, \\ & \nu_t^{\mathsf{T}} = \begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}, \quad \forall t > 2. \end{split}$$

In order to respect the simplex structure for μ_t and ν_t , we have the following constraints at all times:

$$0 \le x_1, x_2, x_3 \le 1, \quad x_2 \le 1 - x_1, \quad x_3 \le x_1.$$

Similarly,

$$0 \le y_1, y_2, y_3 \le 1, \quad y_2 \le 1 - y_1, \quad y_3 \le y_1.$$

For the distribution at t=2 to be $\begin{bmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{bmatrix}^T$ for both teams, we get the additional constraints

$$x_3 = y_3 = \frac{1}{3},$$

$$x_1 + x_2 - x_3 = y_1 + y_2 - y_3 = \frac{1}{3},$$

$$\Rightarrow x_1 + x_2 = y_1 + y_2 = \frac{2}{3},$$

which implies that

$$x_1, y_1 \le \frac{2}{3},$$

since $x_2, y_2 \ge 0$. The constraints now take the form

$$\frac{1}{3} \le x_1 \le \frac{2}{3},\tag{C.38}$$

and similarly,

$$\frac{1}{3} \le y_1 \le \frac{2}{3}.\tag{C.39}$$

The objective function for cRPS is given by

$$J^{N,\phi,\psi}(\mu_0,\nu_0) = \mathbb{E}_{\phi,\psi}\Big[\sum_{t=1}^{T} \mu_t^{\mathsf{T}} A \nu_t \Big| \mu_0,\nu_0\Big], \tag{C.40}$$

which leads to the optimization problem

$$\max_{\phi^t} \min_{j,t} \quad J^{N_1,N_2,\phi^t,\psi^t}(\mu_0,\nu_0) = \mu_1^{\mathsf{T}} A \nu_1 \ . \tag{C.41}$$

Substituting $\mu_1^{\mathsf{T}} = [1-x_1, x_1, 0]$ and $\nu_1^{\mathsf{T}} = [0, 1-y_1, y_1]$ results in the following expression for the maximizing Blue team:

$$\max_{\phi^t} J^{N_1, N_2, \phi^t, \psi^t}(\mu_0, \nu_0) = x_1 + 2y_1 - 3x_1y_1 - 1$$
$$= x_1(1 - 3y_1) + (2y_1 - 1).$$

Since this equation is linear in x_1 , the solution to the maximization problem subject to the constraint (C.38) is

$$x_1 = \frac{1}{3}, \quad y_1 > \frac{1}{3},$$
 (C.42)

$$x_1 = \frac{2}{3}, \quad y_1 < \frac{1}{3},$$
 (C.43)

$$x_1 \in \left[\frac{1}{3}, \frac{2}{3}\right], \quad y_1 = \frac{1}{3}.$$
 (C.44)

Following the same approach for the minimizing Red team, we get the following objective,

$$\min_{\psi^t} J^{N_1, N_2, \phi^t, \psi^t}(\mu_0, \nu_0) = x_1 + 2y_1 - 3x_1y_1 - 1$$
$$= y_1(2 - 3x_1) + (x_1 - 1),$$

subject to the constraint (C.39), with the solution being:

$$y_1 = \frac{1}{3}, \quad x_1 < \frac{2}{3},$$
 (C.45)

$$y_1 = \frac{2}{3}, \quad x_1 > \frac{2}{3},$$
 (C.46)

$$y_1 \in \left[\frac{1}{3}, \frac{2}{3}\right], \quad x_1 = \frac{2}{3}.$$
 (C.47)

Constraint (C.38) ensures that (C.46) cannot hold, while constraint (C.39) similarly prevents (C.43) from holding.

Consider now the case when $y_1 > \frac{1}{3}$. From (C.42) it follows that $x_1 = \frac{1}{3}$. Conversely, if the Blue team commits to a distribution with $x_1 = \frac{1}{3}$, the Red team's best response given by (C.45) gives $y_1 = \frac{1}{3}$, resulting in an incentive for the Red team to deviate from $y_1 > \frac{1}{3}$. Thus, (C.42) does not constitute an optimal solution. Following a similar argument, it can be shown that (C.47) is not an optimal solution either, as illustrated below.

Assume that $x_1 = \frac{2}{3}$. From (C.47), $y_1 \in [\frac{1}{3}, \frac{2}{3}]$. Now, if the Red team announces that it will deploy the distribution $y_1 \in [\frac{1}{3}, \frac{2}{3}]$, the Blue team's response for x_1 follows from (C.42) and (C.44). We have already established that (C.42) is not an optimal solution. This implies that $x_1 \in [\frac{1}{3}, \frac{2}{3}]$ can be a possible response to the Red team. However it violates (C.47), where $x_1 = \frac{2}{3}$ follows from strict equality. Thus, (C.47) does not constitute an optimal solution as the Blue team has an incentive to deviate.

Now, suppose the Blue team announces a distribution where $x_1 \in [\frac{1}{3}, \frac{2}{3}]$. In this case, the Red team's optimal response, derived from (C.45) and (C.44), is $y_1 = \frac{1}{3}$. Conversely, if the Red team announces that its distribution will be $y_1 = \frac{1}{3}$, the Blue team will still follow $x_1 \in [\frac{1}{3}, \frac{2}{3}]$. Since neither team has

an incentive to deviate from these distributions, they form an optimal trajectory. Thus, the solution to the bilinear optimization problem for two-time step convergence takes the form:

$$\mu_1^* = \begin{bmatrix} 1 - x_1 \\ x_1 \\ 0 \end{bmatrix} \quad \text{and} \quad \nu_1^* = \begin{bmatrix} 0 \\ \frac{2}{3} \\ \frac{1}{3} \end{bmatrix},$$
 (C.48)

such that $x_1 \in \left[\frac{1}{3}, \frac{2}{3}\right]$, leading to a game value of $-\frac{1}{3}$. This establishes the distribution at t=1 and confirms the existence of a two-time step optimal trajectory, thereby proving the first part of the proposition.

Now note the following:

- 1. The original objective function (C.40) can be expressed in a bilinear form (similar to the expressions for μ_0, μ_1, μ_2 using x_1, x_2, x_3). This makes it concave in the first argument and convex in the second argument.
- 2. The mean-fields μ and ν lie on a simplex and are hence, compact and convex.

Thus, by the generalized version of von Neumann's minimax theorem Sorin (2002), we conclude that the game value is unique, proving the second part of the proposition 2 .

C.5 IMPLEMENTATION DETAILS AND HYPERPARAMETERS

The state distributions are represented as arrays that are concatenated together to form the global observation. This becomes the input to the critic network which consists of a single hidden layer of 64 neurons and two tanh activation functions. The output is a single value that is equal to the estimated value function. On the other hand, the actor-network consists of a single MLP layer of 64 neurons that is concatenated with the local agent observation. Additionally, the logits are converted to a probability distribution through a softmax layer. The dimension scales with $|\mathcal{A}|$. Both the actor and critic networks are initialized using orthogonal initialization (Huang et al., 2022).

The single-stage RPS game is trained for 5,000 time steps with the actor and critic learning rates set to 0.0005 and 0.001, respectively, which remain constant throughout training. The networks are updated using the ADAM optimizer Kingma & Ba (2017) every 50 time steps for 10 epochs and a PPO clip value of 0.1. The entropy is decayed from 0.01 to 0.001 geometrically. We use an episode length of 1 after which the rewards are bootstrapped.

Moreover, since we have a single "team" buffer and the input/output dimensions are small, we do not use a mini-batch based update. For cRPS we use an episode length of 10 after which the rewards are bootstrapped. cRPS is trained using 200,000 time steps (=20,000 episodes) and is updated every 100 time steps. The algorithm was trained on a single NVIDIA GeForce RTX 3070 GPU and the training times are given in Tables 1 and 2.

D BATTLEFIELD SETUP

D.1 STATE AND ACTION SPACE

We consider a large-scale two-team (Blue and Red) ZS-MFTG on an $n \times n$ grid world. The state of the i^{th} Blue agent is defined as the pair $x_i = (p_i^x, s_i^x)$ where $p_i^x \in \mathcal{S}_{\text{position}}$ denotes the position of the agent in the grid world and $s_i^x \in \mathcal{S}_{\text{status}} = \{0,1\}$ defines the status of the agent: 0 being inactive and 1 being active. Similarly, we define the state of the Red agent as $y_i = (p_i^y, s_i^y)$. The state spaces for the Blue and Red teams are denoted by $\mathcal{X} = \mathcal{Y} = \mathcal{S}_{\text{position}} \times \mathcal{S}_{\text{status}}$, respectively. The mean-fields of the Blue (μ) and Red (ν) teams are distributions over the joint position and status space, i.e., $\mu, \nu \in \mathcal{P}(\mathcal{S}_{\text{position}} \times \mathcal{S}_{\text{status}})$. The action spaces are given by $\mathcal{U} = \mathcal{V} = \{\text{Up}, \text{Down}, \text{Left}, \text{Right}, \text{Stay}\}$ for both teams, representing discrete movements in the grid world. The learned identical team

²Note: The optimal infinite horizon trajectory itself need not be unique (we have shown that x_1 can take a range of values).

policy assigns actions based on an agent's local position and status, as well as the observed mean-fields of both teams. In the following subsections, we elaborate on the weakly coupled transition dynamics and reward structure introduced in the game, followed by a detailed discussion of the training procedure and network architecture for MF-MAPPO in this example.

D.2 Interaction Between Agents

The transitions between states for agents belonging to both teams are characterized by their dynamics. These dynamics are probabilistic and depend on interactions among agents and are weakly coupled through their mean-field distributions. The weak coupling dynamics is keeping in line with the assumption in Guan et al. (2024a).

An agent at a given grid cell can be deactivated by the opponent team with a nonzero probability if the empirical mean-field of the opponent team at the grid cell supersedes that of the agent's own team. Similarly, a deactivated agent can be revived if the empirical mean-field of the agent's team is greater than the opponent's. This is referred to as numerical advantage. The total transition probability from state (p, s) to state (p', s') by taking an action a is given by

$$\mathbb{P}((p', s') | (p, s), a, \mu, \nu) = \mathbb{P}(p' | (p', s'), a) \, \mathbb{P}(s' | (p, s), \mu, \nu),$$

where the first term on the right-hand side corresponds to the deterministic position transition when the agent is active. The second term corresponding to the status transition is given by

$$\mathbb{P}(0 \mid (p, 1), \mu, \nu) = \text{clip}_{[0, 1]} (\alpha_x(\nu(p) - \mu(p))),$$

$$\mathbb{P}(1 \mid (p, 1), \mu, \nu) = 1 - \mathbb{P}(0 \mid (p, 1), \mu, \nu),$$

and

$$\mathbb{P}(1 | (p, 0), \mu, \nu) = \text{clip}_{[0,1]} (\beta_x(\mu(p) - \nu(p))),$$

$$\mathbb{P}(0 | (p, 0), \mu, \nu) = 1 - \mathbb{P}(1 | (p, 0), \mu, \nu),$$

where $\nu(p)-\mu(p)$ is the Red team's numerical advantage over the Blue team at p Similarly, the Blue team's numerical advantage over Red is given by $\mu(p)-\nu(p)$. α_x and β_x are tuning parameters to control the Red team's deactivation power and Blue team's reactivation power respectively. The Red team, being the defending team, is given a slight advantage in terms of higher deactivation power. This enables the possibility of capturing Blue team agents. However, to avoid degeneracy, the Red team agents are not allowed to enter the target. For our experiments, we assume $\alpha_x=15$, $\alpha_y=5$, and $\beta_x=\beta_y=0$.

D.3 REWARD STRUCTURE

The team rewards only depend on the mean-fields of the two teams. For the battlefield scenario, the Blue team agents receive a positive reward corresponding to the fraction of agents that reach the target alive. This is a one-time reward that depends on the change in the fraction of the population of the agents at the target, i.e., if $\mu_t|_{\text{target}} = \mu_{t+1}|_{\text{target}}$, then the team does not receive any positive reward. Each agent in the team receives an identical "team reward." The reward function is mathematically formulated as

$$R_{\text{Blue},t+1}(\mu,\nu) = \kappa \Delta \mu_{t+1}|_{\text{target}},$$

where,

$$\Delta \mu_{t+1}|_{\text{target}} = \mu_{t+1}(p^x = \text{Target}, s^x = 1) - \mu_t(p^x = \text{Target}, s^x = 1).$$

We have chosen $\kappa=100$ in our simulations (heavier emphasis on reaching the target). The Red team's reward is the negative of the Blue team since we have a zero-sum game. Each team aims to maximize its own expected reward.

D.4 IMPLEMENTATION AND HYPERPARAMETERS

The state distribution for a grid world of size $n \times n$ is represented as a three-dimensional array of size (2, n, n) for each team. The first layer depicts the mean-field of the agents over an $n \times n$ grid that are

alive and active, while the second layer gives information about the team's deactivated population. Each team's distribution is then concatenated together to form the global observation. This becomes the common information that is the input to the critic network which in our case is of size (4, n, n) as we have two teams. Both neural networks consist of two main parts: a convolutional block and a fully connected block.

For the critic, the first CNN layer is the input layer that takes the 4 channels and outputs 32 channels, with a kernel size of 3x3, stride of 1, and padding of 1. Followed by ReLU activation, we have a hidden layer that takes 32 channels and outputs 64 channels, with the same kernel size, stride, and padding. Lastly, after another ReLU activation, we have the output layer that takes 64 channels and outputs 64 channels, again with the same kernel size, stride, and padding. After another ReLU layer, the output of the CNN is passed through an MLP. Namely, a fully connected (dense) layer takes the flattened output of the convolutional block and reduces it to 128 units. Between the input and the output layers, we have a single tanh activation function.

On the other hand, the input to the actor-network is split into two CNN blocks: one to process the common information and one to process the local information. The local information channel, is an array of size (1, n, n) that locates the position of the agent with value +1 if it is active and -1 if it has been deactivated. This local information is passed through a single CNN layer that outputs 16 channels with a kernel size of 3x3, stride of 1, and padding of 1 while the common information is passed through two such layers with the output of 32 channels. Both outputs are then followed by a ReLU activation function and the latent representation of the common information combined with the local agent observation is then passed through an MLP architecture.

A fully connected (dense) layer takes the flattened output of the convolutional block and reduces it to 512 units. We have a single hidden layer that reduces the dimension further to 128 and then the output logits. The layers are separated by the tanh activation functions. Finally, the logits are converted to a probability distribution through a softmax layer. Both the actor and critic networks are initialized using orthogonal initialization (Huang et al., 2022). The architectures of the shared-team actor and minimally-informed critic networks for this example are shown in Figures D.2 and D.3 respectively.

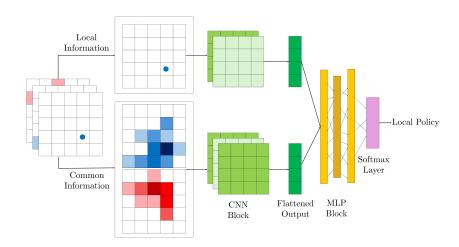


Figure D.2: MF-MAPPO: Shared-team actor for battlefield

All maps are trained using a single NVIDIA GeForce RTX 3070 GPU. The actor and critic learning rates are set to 0.0005 and 0.001 and both decay geometrically by a factor of 0.999. The networks are updated using the ADAM optimizer (Kingma & Ba, 2017) with two mini-batches for 10 epochs and a PPO clip value of 0.1. The entropy coefficient is initialized to 0.01 and decays with a factor of 0.995.

Maps 1 and 2 which are 4×4 grid worlds are trained for 5×10^6 and 4.5×10^6 time steps, respectively, and in both cases, the episode length is 20 time steps and the update frequency is every 500 time steps. The total training period is about one day. On the other hand, Map 3 being 8×8 in dimension,

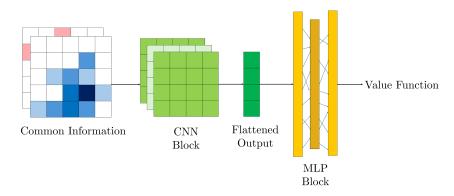


Figure D.3: MF-MAPPO: Minimally-informed critic for battlefield

has an episode length of 64, is trained for 9×10^6 time steps and its network is updated every 1,000 time steps. The total training period is approximately three days.

For evaluating D-PC, we employ MF-MAPPO to learn Lipschitz policies under a fully observable mean-field regime, and use these policies to evaluate the performance of our dynamic estimation algorithm, D-PC. We consider a time-invariant, edgeless visibility graph, wherein agents have access only to the mean-field distribution of their opponent corresponding to their current cell. For the communication graph, we assume a subgrid based topology that we detail in Section E.3. The communication graph is naturally time-varying since the number of states occupied by the agents need not be fixed. In accordance with Assumption 2 and the theoretical requirements of both D-PC and benchmark algorithms, we explicitly impose the constraint that the communication graph is connected at all time steps.

At each time step t, the initial Red team's estimate for the D-PC algorithm by the Blue team (i.e., at $\tau=0$) for every state x is taken as the projection of the estimate from the previous time step t-1 onto the constraint set $\mathcal{R}(x)$, that is,

$$\hat{\nu}_{x,t}^{N_2,\tau=0} = \Omega_{\mathcal{R}(x)} \left[\hat{\nu}_{x,t-1}^{N_2,R_{\text{com}}} \right], \quad x \in \mathcal{X}_t^o.$$

At t=0 and $\tau=0$, we assume $\hat{\nu}_{x,t=0}^{N_2,\tau=0}\sim \mathrm{unif}\left(\mathcal{R}(x)\right)$ for all $x\in\mathcal{X}_t^o$.

E ADDITIONAL RESULTS

In this section, we present additional simulation results from various environments present on MFEnv. In particular, we look at the standard version of Rock-Paper-Scissors and more experiments on the zero-sum battlefield game - both trained using MF-MAPPO. We also present a single team grid world navigation task wherein agents learn a Lipschitz constrained MF-MAPPO policy and use D-PC to estimate their empirical distributions. This grid world navigation task in particular highlights the applicability of the presented algorithms to general mean-field team settings, not restricted to competitive zero-sum team games.

E.1 ROCK-PAPER-SCISSORS (RPS)

We first extend the two-player Rock-Paper-Scissors (RPS) game to a game played between two populations as described in Appendix C. The Nash equilibrium for this population-based RPS game is the uniform population distribution [1/3, 1/3, 1/3] over the 3 states (Raghavan, 1994).

We compare MF-MAPPO with DDPG-MFTG (Shao et al., 2024b) based on the training time, average test rewards and attainment of the computed Nash distributions for $N_1=N_2=1,000$ agents. We exclude MADDPG (Lowe et al., 2017) from our comparison, as it scales poorly to hundreds or thousands of agents due to its reliance on all agents' local and global observations and actions as inputs to its critic networks. We include the training curve and performance for cRPS in Figure E.4 and Table 2 respectively, for reference.

From the learning curves in Figure E.4 one can see that the DDPG-MFTG algorithm failed to converge to the analytical game value of zero, while MF-MAPPO almost immediately attained the Nash game value. This corroborates with the results presented for cRPS in the main text. However, as shown in Table 1, MF-MAPPO does take slightly longer to train since, unlike DDPG-MFTG, since MF-MAPPO avoids mini-batch training, following Yu et al. (2022).

We tested the learned policy with a fixed initial distribution $\mu_{t=0} = [1,0,0]^T$ and $\nu_{t=0} = [0,1,0]^T$, and the resulting trajectories are visualized in Figure E.5. All simulations were run for 150 instances. The trajectories of the Blue and Red team ED are depicted in cyan and pink, respectively, alongside the mean trajectory. The randomness in these trajectories arises from the finite-population approximation under a stochastic optimal policy, resulting in stochastic EDs. As shown in Figures E.5, DDPG-MFTG diverges from the equilibrium whereas MF-MAPPO converges immediately.

E.2 BATTLEFIELD

E.2.1 VALIDATION CASES FOR MF-MAPPO

The following subsection qualitatively discusses the battlefield game for different map layouts. For these results, both teams are deploying policies trained using MF-MAPPO.

Map A. The first map is a simple 4×4 grid world with a single target that we use to validate our algorithm. The target is partially blocked by an obstacle, see Figure E.6. For the initial condition in Figure E.6, the Blue team is initially split into two equal groups. The Blue team decides to merge the two sub-groups of agents into a single group. With this formation, the Red team has zero numerical advantage over the Blue team when they encounter in (g), resulting in all Blue agents safely arriving at the target. In comparison, if the two Blue subgroups do not merge but move toward the target one at a time, it will lead to 50% of the Blue team population being deactivated (first subgroup), followed by the remaining 50% (second subgroup). This demonstrates how the observation of mean-field distributions guides rational decision-making.

Map B. This map is identical to the one presented in Section 5. In the first scenario (Figure E.7), 70% of the Blue agents start at cell [2, 2], and 30% at cell [1, 1], while the Red team is evenly split between cells [0, 1] and [3, 3]. Half of the Red team at [0, 1] successfully blocks the 30% Blue agent group from entering the left corridor due to its numerical advantage, which forces the Blue agents to opt for the right corridor. At the same time, the larger Blue group with 70% of the population utilized their numerical advantage over the half Red team at the top right and deactivated all the Red agents as shown in (c) and reached the target at time step (d). This allowed the smaller 30% group

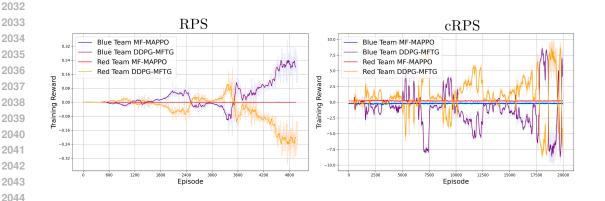


Figure E.4: Training curves for RPS and cRPS.

Table 1: Performance comparison for RPS

Approach	Training Time	Average Reward	NE Attained?
MF-MAPPO	5min 17s	0.0	✓
DDPG-MFTG	1min 34s	0.334	Х

2	0	5	2	
2	0	5	3	

Table 2: Performance comparison for cRPS

Approach	Training Time	Average Reward	NE Attained?
MF-MAPPO	2h 17min 15s	-0.331	✓
DDPG-MFTG	60h 49min 41s	3.774	Х

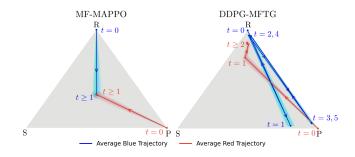


Figure E.5: ED trajectories induced by learned team policies on the state distribution simplex. Mean trajectories are averaged based on the 150 runs from fixed initialization $\mu_{t=0} = [1,0,0]^T$ and $\nu_{t=0} = [0,1,0]^T$; $N_1 = N_2 = 1,000$.

to follow through the same corridor without losing agents. In the second scenario (Figure E.8), with the Red team evenly distributed at the corners, 30% of the Blue agents start at cell [2, 2] and 70% at cell [3, 1]. The Red team's numerical advantage at [3, 3] forces the Blue agents to move around and regroup (Figures E.8(b)-E.8(f)). Once united, the Blue team's numerical advantage forces the Red subgroup at [0,1] to disperse to avoid deactivation, allowing Blue to reach the target.

Map C. In Figure E.9, we deploy the MF-MAPPO policies trained with $N_1=N_2=100$ (from the main text) to larger populations of $\bar{N}_1=\bar{N}_2=1500$ and $\bar{N}_1=\bar{N}_2=1000$ using the same initial condition as Figure 4I and II respectively. The teams continue to achieve their objectives as established in Theorem 3, exhibiting behaviors similar to $N_1=N_2=100$ setting. Section 5 presented scenarios featuring structured initial configurations for both teams over an 8×8 grid. It is important to emphasize, however, that the algorithm is trained on a diverse set of initial conditions for a given map, ranging from agents concentrated within a few selected cells to agents distributed randomly across the grid world. The following examples demonstrate that the teams are able to accomplish their objectives even in scenarios where agents are dispersed across the environment rather than clustered into just 1-2 subgroups using the same policy from Section 5. In Figure E.10, the Blue team is initialized randomly, and local subgroups of agents emerge and coordinate to reach the target. This behavior is particularly pronounced near the upper target, where a greater numerical advantage facilitates successful coalition formation (Figures E.10(c)–(e)).

Turning to the randomly distributed Red team agents in Figure E.11, it is observed that they concentrate near the two target entrances and successfully neutralize most Blue team subgroups.

E.2.2 Comparison of MF-MAPPO with Baseline

Initial Condition 1. Figure E.12I. compares the two algorithms against the baseline defending team. MF-MAPPO Blue agents exhibit coordinated maneuvering, forming coalitions to reach the target (a), whereas DDPG-MFTG Blue agents (b) show limited coordination, with only nearby agents reaching the target and distant agents failing to engage. In (c) we pit the Blue team against the MF-MAPPO defenders instead of the DDPG-MFTG defenders. The results align with those in Figure 3, where MF-MAPPO Blue agents effectively leverage their numerical advantage, enabling a larger number of agents to reach the target. (d) represents MF-MAPPO vs. MF-MAPPO to illustrate the goal strategies and expected behavior.

Initial Condition 2. We present a different initial condition for the 4×4 Battlefield game (Figure E.12II.), where, using similar arguments as in the previous case and the main text, we can establish the superiority of MF-MAPPO agents over the baseline DDPG-MFTG, whether MF-MAPPO serves as the attacker or the defender.

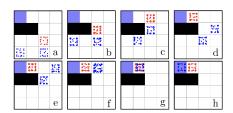


Figure E.6: Red is concentrated; Blue is evenly split.

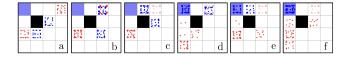


Figure E.7: Red is evenly split; 30% Blue are at [1, 1] and 70% are at [2, 2].

E.3 MEAN-FIELD ESTIMATION FOR GRID WORLD NAVIGATION USING D-PC

We consider a 9×9 grid with a target region at the center that is surrounded by penetrable obstacles (see Figure E.13). At a given cell, the population can penetrate an adjacent obstacle-state $x_{\rm obstacle}$ with a probability proportional to its distribution at that cell, according to

$$f(x_{\text{obstacle}}|x, u, \mu_t) \propto \exp^{b(\mu_t(x) - c)},$$
 (E.49)

where b>0 and $c\in(0,1)$. We run the algorithms for T=1,000 time steps for varying values of $R_{\rm com}$ under a subgrid-based communication topology. For each value of $R_{\rm com}$, we evaluate our algorithm as well as the benchmark over multiple seeds of different initial mean-field distributions $\mu_{t=0}^N=\hat{\mu}_{t=0}^N$ and report the average estimation error over these seeds.

We use an edgeless visibility graph and a subgrid-based communication graph. In the subgrid communication topology, an agent located at state x can exchange information with its neighboring states arranged in a 3×3 grid centered around x, as depicted in Figure E.13. Naturally, agents at the boundary have reduced connectivity: edge cells communicate within a 2×3 neighborhood, and corner cells within a 2×2 configuration, also shown in Figure E.13. We evaluate the estimation performance for communication rounds $R_{\rm com} \in \{1,2,3,4,5,6,7,8\}$ and consider population sizes of N=100 and N=1,000 (Recall, Theorem 3). From Figures E.14 and E.16 we see that up to approximately $R_{\rm com}=4$, the proposed algorithm consistently outperforms the benchmark, particularly in the early phases where communication is limited.

The benchmark algorithm (Benjamin & Abate, 2025b) assumes that each state directly acquires the exact distributions of its communication neighbors $\mathcal{N}^{\text{com}}(x)$. Consequently, it achieves better accuracy when $R_{\rm com}$ approaches diam $(\mathcal{G}_{\epsilon}^{\rm com})$ communication rounds. While our method converges slightly more gradually, it offers notable practical advantages: it avoids relying on direct access to neighbor distributions—thereby preserving privacy—and performs competitively even under tighter communication constraints. Furthermore, for $R_{\text{com}} < \text{diam}(\mathcal{G}_t^{\text{com}})$, the benchmark estimates unobserved states using a uniform distribution assumption—potentially leading to inaccurate representations of the actual mean-field. This inaccuracy is especially visible for small communication budgets ($R_{\rm com} = 1, 2, 3$). Although increasing rounds mitigates this error, in real-world scenarios where fast decision-making is critical such as in real-time opponent modeling, large-scale communication may be infeasible. In contrast, our algorithm remains tractable, communication-efficient, and robust across a range of practical deployment settings, offering a compelling trade-off between estimation quality and operational cost. The reward plots in Figures E.15 and E.17, corresponding to both population sizes, demonstrate that our method exhibits notably lower regret compared to the benchmark when evaluated against the fully observable policy. This improvement is particularly evident when $R_{\text{com}} < \frac{1}{2} \text{diam}(\mathcal{G}_t^{\text{com}})$. For higher number of communication rounds our performance remains competitive with the benchmark. The plots also show that the estimation error incurred reduces with an increase in population size, predominantly noticeable at the peak around t = 100. Furthermore, the increase in population size from N=100 to N=1,000 results in reduced noise in the total variation errors and smoother error plots, thereby resulting in smaller variance. Both

1.24		iii	Ei Ei a	DE		\$100 \$100 \$100 \$100 \$100 \$100 \$100 \$100	∴ !!! b	124 2-4	3-4	IVI ISI	·:.	19.1	107	課 六	d
	æ				10				И			ļ:	Į.		
鹼	6	124 2-2		ĮSĮ				134				28	1,1	••	
•			е				f	132	10		g				h

Figure E.8: Red is evenly split; 70% Blue are at [3, 1] and 30% are at [2, 2].

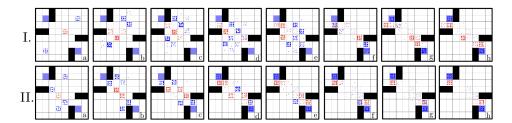


Figure E.9: I. 1500 agents in each team where Red is concentrated; 30% Blue are at the bottom, rest are at the top II. 1000 agents in each team with Blue evenly split; Red is concentrated.

these results corroborate with the finite population mean-field approximation guarantees presented from Theorem 5. Empirically, we assert that our algorithm is particularly well-suited for scenarios with limited communication budgets, especially when $R_{\rm com} < \frac{1}{2} {\rm diam}(\mathcal{G}_t^{\rm com})$, where it consistently delivers competitive accuracy with substantially reduced overhead.

A direct consequence of Theorem 4 is that increasing the number of communication rounds $R_{\rm com}$ enables the estimation error ϵ to be made arbitrarily small. In finite populations, however, stochasticity in the MF approximation causes the full-information and estimated trajectories, \mathcal{M}^{N_1} , \mathcal{N}^{N_2} and $\hat{\mathcal{M}}^{N_1}$, $\hat{\mathcal{N}}^{N_2}$, to vary across runs of GR-MF-MAPPO even under identical initial conditions. We confirm that D-PC achieves exponential convergence, which becomes most apparent in the infinite-population limit where state evolution is deterministic. Figure E.18 illustrates this behavior by plotting the estimation error at selected time steps against $R_{\rm com}$, consistent with Theorem 4.

E.4 Non-Lipschitz Policies and Estimation

While introducing partially observable ZS-MFTGs, we emphasized the importance of enforcing Lipschitz continuity of policies with respect to the mean-field distribution. In this section, we examine how varying the Lipschitz gradient penalty coefficient λ influences the accuracy of the mean-field estimation procedure. Specifically, we introduce the following term in the objective (8):

$$\lambda \mathbb{E}\left[\nabla_{\eta} \|\log \phi(u \mid x, \mu, \nu)\|^{2}\right], \tag{E.50}$$

where λ denotes the Lipschitz gradient penalty coefficient. Larger values of λ enforce stronger smoothness constraints on the policy, promoting more regular behavior across variations in the mean-field input. In contrast, setting $\lambda=0$ removes this constraint entirely. We conduct two sets of experiments in the finite-population grid world setting.

Battlefield. We utilize the same 4×4 battlefield introduced in Section 5 and consider the case wherein the Red team is estimating the Blue team's distribution and the Blue team has full information. We consider three distinct policies, each trained under a different value of λ , and allow the teams to estimate mean-field distributions that evolve accordingly. From Table 3 we observe that for most cases, $\Delta J(\phi_t^*, \psi_t^*)$ (represented as % for normalized values) is higher for the policy trained without a Lipschitz constraint and the error reduces as λ is increased, especially for larger values of R_{com} . This is because even small inaccuracies in the estimated distributions can lead to large discrepancies in the resulting mean-field, causing substantial deviations from the ideal fully observable behavior. Thus, enforcing Lipschitz continuity in policies offers tangible practical benefits by promoting smoother and more robust behaviors. Moreover, such policies tend to yield behaviors that

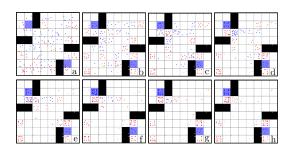


Figure E.10: Blue team is randomly spread around the map.

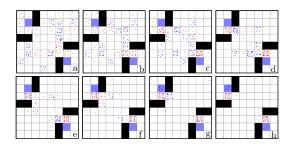


Figure E.11: Red team is randomly spread around the map.

are better aligned with the demands of real-world deployment, making them a compelling design choice when evaluating estimation performance.

Navigation. We consider the navigation problem defined above and now focus on the cumulative total variation error, i.e., $\sum_t d_{\rm TV}(\mu_t, \hat{\mu}_t)$. Figure E.19 show the performance of our proposed D-PC algorithm and the benchmark method from Benjamin & Abate (2025b). In both estimation settings, once again, gradient-regularization yields trajectories that closely resemble the fully-observable scenario (lower errors).

Table 3: Effect of gradient-regularization on $\Delta J(\phi_t^*, \psi_t^*)$

λ	$R_{\rm com} = 5$	$R_{\rm com} = 10$	$R_{\rm com} = 15$	$R_{\rm com} = 20$
0.0	0.39	0.20	1.15	2.10
0.005	0.09	0.65	0.47	0.27
0.01	0.49	0.26	0.38	0.03

F ROLE OF LLMS

The authors acknowledge the use of GPT-4 and GPT-5 for polishing the main text and for quick access to existing results in probability theory, in particular, for Theorem 2 and Proposition 3.

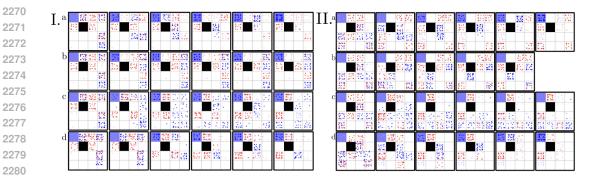


Figure E.12: a. MF-MAPPO Blue vs. DDPG-MFTG Red; b. DDPG-MFTG Blue vs. DDPG-MFTG Red; c. MF-MAPPO Red vs. DDPG-MFTG Blue; d. MF-MAPPO Red vs. MF-MAPPO Blue.

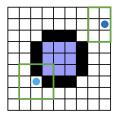


Figure E.13: Subgrid Communication Graph.

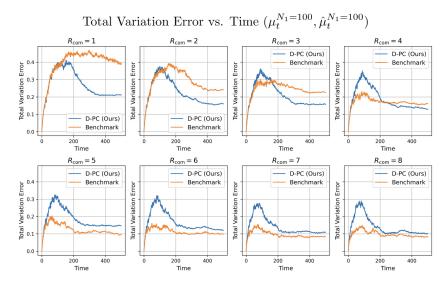


Figure E.14: Total variation error at each time step for different values of R_{com} with N = 100.

2327

2328

2329

2330

233123322333

23342335

2336

2337

2338233923402341

23422343

23542355

2356

2357

23582359

2360

2361

2362

2363

2364

2365

236623672368

236923702371

2372237323742375

Cumulative Regret upto Time t for N = 100 $R_{com} = 2$ $R_{com} = 1$ $R_{com} = 3$ $R_{com} = 4$ Error in Cumulative Rewards D-PC (Ours) Benchmark D-PC (Ours) Benchmark D-PC (Ours) Benchmark D-PC (Ours) Benchmark Error in Cumulative Error in Cumulative 200 Time 200 Time 200 Time 200 Time 400 400 400 400 $R_{com} = 5$ $R_{com} = 6$ $R_{com} = 7$ $R_{com} = 8$ in Cumulative Rewards D-PC (Ours) D-PC (Ours) D-PC (Ours) D-PC (Ours) Benchmark Benchmark Benchmark Benchmark Error in Cumulative 400 400

Figure E.15: Regret at each time step for different values of $R_{\rm com}$ with N=100.

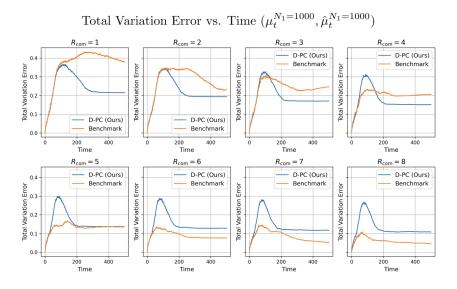


Figure E.16: Total variation error at each time step for different values of $R_{\rm com}$ with N=1000.

Cumulative Regret upto Time t for N = 1000 $R_{com} = 2$ $R_{com} = 3$ $R_{com} = 4$ in Cumulative Rewards D-PC (Ours) D-PC (Ours) D-PC (Ours) D-PC (Ours) Benchmark Error in Cumulative Error in Cumulative in Cun Time Time Time Time $R_{com} = 5$ $R_{com} = 6$ $R_{com} = 7$ $R_{com} = 8$ in Cumulative Rewards Rewards Rewards D-PC (Ours) D-PC (Ours) D-PC (Ours) Benchmark D-PC (Ours) Benchmark Benchmark Benchmark in Cumulative Error in Cumulative Time Time Time

Figure E.17: Regret at each time step for different values of R_{com} with N = 1000.

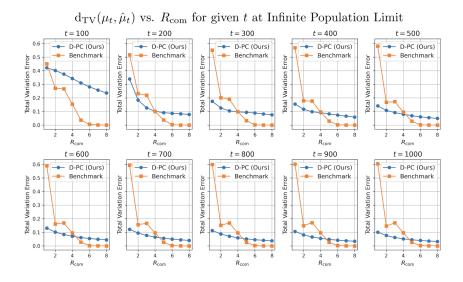


Figure E.18: Exponential convergence of D-PC at the infinite population limit, in the absence of stochasticity.

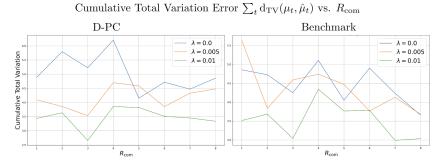


Figure E.19: Cumulative total variation error by various estimation algorithms in a 9×9 grid world for different values of λ under a subgrid communication graph with N = 1,000.