

Mean-Field Approximation of Cooperative Constrained Multi-Agent Reinforcement Learning (CMARL)

Washim Uddin Mondal

WMONDAL@PURDUE.EDU

*Lyles School of Civil Engineering,
School of Industrial Engineering,
Purdue University,
West Lafayette, IN, 47907, USA*

Vaneet Aggarwal

VANEET@PURDUE.EDU

*School of Industrial Engineering,
School of Electrical and Computer Engineering,
Purdue University,
West Lafayette, IN, 47907, USA*

Satish V. Ukkusuri

SUKKUSUR@PURDUE.EDU

*Lyles School of Civil Engineering,
Purdue University,
West Lafayette, IN, 47907, USA*

Editor: Tor Lattimore

Abstract

Mean-Field Control (MFC) has recently been proven to be a scalable tool to approximately solve large-scale multi-agent reinforcement learning (MARL) problems. However, these studies are typically limited to unconstrained cumulative reward maximization framework. In this paper, we show that one can use the MFC approach to approximate the MARL problem even in the presence of constraints. Specifically, we prove that, an N -agent constrained MARL problem, with state, and action spaces of each individual agents being of sizes $|\mathcal{X}|$, and $|\mathcal{U}|$ respectively, can be approximated by an associated constrained MFC problem with an error, $e \triangleq \mathcal{O}\left(\frac{[\sqrt{|\mathcal{X}}| + \sqrt{|\mathcal{U}}|]}{\sqrt{N}}\right)$. In a special case where the reward, cost, and state transition functions are independent of the action distribution of the population, we prove that the error can be improved to $e = \mathcal{O}(\sqrt{|\mathcal{X}|}/\sqrt{N})$. Also, we provide a Natural Policy Gradient based algorithm, and prove that it can solve the constrained MARL problem within an error of $\mathcal{O}(e)$ with a sample complexity of $\mathcal{O}(e^{-6})$.

1. Introduction

Consider the situation of a policy-maker that is tasked with allocating a certain amount of budget to improve the health of a number of infrastructures in a post-disaster scenario. If, due to the repair, the health of a certain infrastructure improves, it provides the perception of a high reward to the planner. However, such improvement comes at the “cost” of the amount of money that is allocated to it. How should the planner allocate the money over a long period of time? To answer such a question is the same as assigning a policy to each infrastructure that states how much money should that infrastructure draw from the budget

pool based on the global state of all infrastructures. The aim of designing such policies should be to maximize the expected cumulative reward while obeying the budget constraint.

Strategizing the decisions of a large pool of interacting agents under certain constraints is a frequently appearing problem in many branches of social science and engineering. For example, such a framework can be applied to power-constrained wireless sensor networks (Buratti et al., 2009), energy harvesting communication networks (Wang et al., 2015), queuing systems with stability constraints (Xiang et al., 2014), etc. A common way to deal with such problems is to use the framework of constrained multi-agent reinforcement learning (CMARL). In cooperative CMARL, the goal is to maximize the aggregate reward of the whole population while obeying some specified constraints. Unfortunately, as the number of agents increases, the size of the joint state space increases exponentially, rendering the task of maximizing reward incredibly hard.

The phenomenon of state-space explosion is not unique to CMARL, unconstrained MARL problems also have to deal with the same issue. There have been several attempts to alleviate this obstacle. For example, one idea is to confine oneself to local policies. This approach has yielded a number of useful algorithms e.g., Independent Q-Learning (IQL) (Tan, 1993), centralized training and decentralized execution (CTDE) (Son et al., 2019; Sunehag et al., 2018; Rashid et al., 2020, 2018; Zhou et al., 2022) based algorithms, etc. Despite their empirical success, none of the above algorithms can provide theoretical guarantees. Moreover, it is difficult to train policies using these algorithms when the number of agents is very high. Another alternative is the framework of mean-field control (MFC) which builds upon the idea that in an infinite population of identical agents, studying one representative agent is sufficient to infer the statistics of their collective behavior (Angiuli et al., 2022). MFC has recently been applied to approximate N -agent unconstrained MARL problems with theoretical guarantees (Gu et al., 2021). In contrast to the existing approaches, the benefit of the MFC is that its optimality gap decreases with an increase in the number of agents. It is, however, still unknown whether similar approximation results can be established for CMARL problems. In this paper, we provide an affirmative answer to this question.

Establishing MFC-based approximation bounds for CMARL is particularly challenging for the following reason. Unlike in unconstrained MARL, the set of feasible policies of CMARL may not be identical to the sets of feasible policies of its associated constrained MFC (CMFC) problem. In general, a feasible policy for CMARL may overshoot the constraints of CMFC and vice versa. Clearly, the sets of feasible policies of CMARL and CMFC may partially overlap but none is a subset of another. The lack of structural hierarchy between the feasible policy sets makes it difficult to compare the optimal values of CMARL and CMFC. The trick is not to deal with the original CMARL, and CMFC problems but to consider variants of those with smaller feasible sets. We show that these variants can be judiciously chosen to enforce hierarchy, and thereby create a way to compare values.

1.1 Our Contribution

We consider an N -agent CMARL problem where at each instant the agents receive rewards as well as incur costs depending on their joint states, and actions. The goal is to maximize the time-discounted sum of rewards (called reward value or simple value) while ensuring that the discounted cumulative cost lies below a certain threshold. We show that the stated N -agent

CMARL problem can be well approximated by a CMFC problem with an appropriately adjusted constraint. In particular, our result (Lemma 1) states that the optimal value of the stated CMFC is at most at the distance of $\mathcal{O}(e)$ from the optimal CMARL value where $e = \frac{1}{\sqrt{N}}[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|}]$. The terms $|\mathcal{X}|, |\mathcal{U}|$ denote the sizes of state, and action spaces of individual agents respectively. We also show that if the optimal policy obtained by solving the CMFC is adopted into the N -agent system, then it does not violate the constraint of CMARL, and yields an N -agent cumulative reward that is $\mathcal{O}(e)$ error away from the optimal N -agent value (Theorem 2). In a special case where the reward, cost, and state transition functions are independent of the action distribution of the population, we prove that the error improves to, $e = \sqrt{|\mathcal{X}|}/\sqrt{N}$ (Theorem 4).

The key idea behind Lemma 1 is a novel sandwiching technique. Specifically, we consider three distinct problems— a CMARL with zero constraint bound (target problem), a CMFC with slightly more restrictive constraint, and a CMARL with even more restrictive constraint. We prove that the optimal value for the CMFC problem must be trapped between the optimal values of the CMARL problems with a small margin of error. Next, we establish that the CMARL values themselves must be close for large N . This forces the CMFC value to lie in the vicinity of our target value.

Finally, using the global convergence result of (Ding et al., 2020), we devise a natural policy gradient-based primal-dual (NPG-PD) algorithm and show that the policy obtained from the stated algorithm satisfies the constraint of CMARL and approximates the optimal N -agent value within an error of $\mathcal{O}(e)$. The sample complexity of the algorithm is shown to be $\mathcal{O}(e^{-6})$ (Theorem 5).

1.2 Related Works

Unconstrained Single Agent RL: One of the early attempts to address the problem of single agent RL was made by the tabular Q-learning (Watkins and Dayan, 1992), and subsequently by the SARSA algorithm (Rummery and Niranjan, 1994). However, such approaches work only when the state space is small. Neural Network-based Deep Q-learning (Mnih et al., 2015), and deep policy gradient algorithms (Li et al., 2019) have been introduced relatively recently to tackle large state space. One cannot, however, scale these approaches to a large number of agents due to the exponential blow-up of their joint state space.

Constrained Single Agent RL: Decision making under constraint is typically modeled via Constrained Markov Decision Problems (CMDPs) (Chow et al., 2017). Several policy gradient (PG) or direct policy search methods have been proposed to solve CMDPs (Achiam et al., 2017; Bhatnagar and Lakshmanan, 2012). However, the convergence guarantees of these algorithms are local. Recently, a series of PG-type algorithms have been proposed that ensure global convergence (Paternain et al., 2022; Ding et al., 2020; Mondal and Aggarwal, 2024a). The advantage of the mean-field approach is that it effectively reduces a multi-agent problem to a single-agent problem, thereby allowing us to utilize these existing guarantees of CMDPs.

MFC-based Approximation of MARL: There have been a number of recent advances that show that MFC well-approximates MARL problems in the regime of a large population. (Gu et al., 2021) showed that an N -agent homogeneous MARL problem can be approximated by MFC within $\mathcal{O}(\frac{1}{\sqrt{N}})$ error. (Mondal et al., 2022a) later extended this idea to heterogeneous

MARL. Their approaches, however, cannot be directly adopted for CMARL. MFC has also been used to design local policies (Mondal et al., 2022b) and approximate cooperative MARL with non-uniform interactions (Mondal et al., 2022c).

Algorithms for MFC: In the model-free set up, both Q -learning based (Angiuli et al., 2022; Gu et al., 2021), and PG-based algorithms (Carmona et al., 2019) have been proposed to solve MFC problems. Model-based algorithms are also available in the literature (Pásztor et al., 2023). However, none of these algorithms apply to CMFC problems.

Application of MFC: MFC has found its application in a multitude of social and engineering applications ranging from epidemic control (Watkins et al., 2016; Lee et al., 2021), ride-hailing (Al-Abbasi et al., 2019), network traffic engineering (Geng et al., 2020), congestion management (Wang et al., 2020) etc.

2. Model for Cooperative CMARL

We consider a collection of N agents interacting with each other in discrete time $t \in \{0, 1, \dots\}$. The state of i -th agent at time t is denoted as x_t^i which can take values from the finite state space, \mathcal{X} . The joint state of N -agents at time t is symbolised as $\mathbf{x}_t^N \triangleq \{x_t^i\}_{i \in \{1, \dots, N\}}$. Upon observing \mathbf{x}_t^N , the i -th agent chooses an action u_t^i from the finite action set, \mathcal{U} according to the following probability law: $u_t^i \sim \pi_i(\mathbf{x}_t^N)$. The symbol, $\pi_i(\cdot)$, is defined to be the policy of the i -th agent which is assumed to be a function of the form $\pi_i : \mathcal{X}^N \rightarrow \mathcal{P}(\mathcal{U})$ where $\mathcal{P}(\cdot)$ defines a probability simplex over its argument. The joint action of N -agents at time t is indicated as $\mathbf{u}_t^N \triangleq \{u_t^i\}_{i \in \{1, \dots, N\}}$. As a consequence of these actions, the i -th agent receives a reward $r_i(\mathbf{x}_t^N, \mathbf{u}_t^N)$, a constraint cost $c_i(\mathbf{x}_t^N, \mathbf{u}_t^N)$ at time t , and its state changes in the next time step according to the following transition rule: $x_{t+1}^i \sim P_i(\mathbf{x}_t^N, \mathbf{u}_t^N)$. Here we implicitly assume that the elements of \mathbf{u}_t^N are conditionally independent given \mathbf{x}_t^N and the elements of \mathbf{x}_{t+1}^N are conditionally independently given $(\mathbf{x}_t^N, \mathbf{u}_t^N)$.

Note that, the reward, cost, and state transition law of an agent are not only affected by its own state, and action but also by states, and actions of other agents. Such interlacing is one of the main obstacles in solving large multi-agent problems. To ease the analysis, we shall assume a special structure of the reward, cost, and transition functions that are routinely adopted in the mean-field literature (Gu et al., 2021; Angiuli et al., 2022). Specifically, we assume that for some $r, c : \mathcal{X} \times \mathcal{U} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{U}) \rightarrow \mathbb{R}$, and $P : \mathcal{X} \times \mathcal{U} \times \mathcal{P}(\mathcal{X}) \times \mathcal{P}(\mathcal{U}) \rightarrow \mathcal{P}(\mathcal{X})$, the following holds $\forall i \in \{1, \dots, N\}$,

$$r_i(\mathbf{x}_t^N, \mathbf{u}_t^N) = r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) \tag{1}$$

$$P_i(\mathbf{x}_t^N, \mathbf{u}_t^N) = P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) \tag{2}$$

$$c_i(\mathbf{x}_t^N, \mathbf{u}_t^N) = c(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) \tag{3}$$

where $\boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N$ denote the empirical N -agent state, and action distributions respectively. Formally,

$$\boldsymbol{\mu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta(x_t^i = x), \quad \forall x \in \mathcal{X}, \quad \text{and} \quad \boldsymbol{\nu}_t^N = \frac{1}{N} \sum_{i=1}^N \delta(u_t^i = u), \quad \forall u \in \mathcal{U} \tag{4}$$

where $\delta(\cdot)$ is the indicator function. The equations (1) – (3) are applicable to a system where agents are identical and exchangeable (Mondal et al., 2022a). These relations suggests

that the effect of the population on any agent is mediated only through the mean-field distributions. We would also like to point out that the functions, r , P , and c are identical for all agents. Therefore, the system of agents described above is essentially homogeneous. As a consequence, we can rewrite the policy functions, π_i as follows for some $\pi : \mathcal{X} \times \mathcal{P}(\mathcal{X}) \rightarrow \mathcal{P}(\mathcal{U})$, and $\forall i \in \{1, \dots, N\}$,

$$\pi_i(\mathbf{x}_t^N) = \pi(x_t^i, \boldsymbol{\mu}_t^N) \quad (5)$$

In simple words, due to homogeneity, we can describe the policy of an agent as a (stochastic) process of choosing its action upon observing its own state, and the state distribution of the whole population. The policy of an agent at time t is denoted as π_t where its sequence is denoted as $\boldsymbol{\pi} \triangleq \{\pi_t\}_{t \in \{0, 1, \dots\}}$. For a given initial joint state, \mathbf{x}_0^N , the N -agent reward, and cost-value (indicated respectively as V_N^R , V_N^C) of a policy-sequence $\boldsymbol{\pi}$ are defined as follows.

$$V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}) \triangleq \sum_{t=0}^{\infty} \gamma^t \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) \right] \quad (6)$$

$$V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}) \triangleq \sum_{t=0}^{\infty} \gamma^t \mathbb{E} \left[\frac{1}{N} \sum_{i=1}^N c(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) \right] \quad (7)$$

where the expectation is computed over all joint state-action trajectories induced by $\boldsymbol{\pi}$, and $\gamma \in [0, 1)$ is a discount factor. Let the set of all admissible policies be Π , and $\Pi_\infty \triangleq \Pi \times \Pi \times \dots$ be the collection of all admissible policy sequences. The goal of N -agent CMARL is to solve the following optimization problem for a given initial joint state \mathbf{x}_0^N .

$$\sup_{\boldsymbol{\pi} \in \Pi_\infty} V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}) \quad \text{subject to : } V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}) \leq 0 \quad (\text{CMARL})$$

We shall denote the solution to (CMARL) as $\boldsymbol{\pi}_N^*$. The existence of an optimal solution is guaranteed via Slater's condition described later. Obtaining a solution of (CMARL), is, in general, difficult. In the next section, we shall demonstrate how $\boldsymbol{\pi}_N^*$ can be approximated via the constrained mean field control (CMFC) approach.

3. Model for Constrained Mean-Field Control (CMFC)

A mean-field system comprises an infinite collection of agents. Due to homogeneity, it is sufficient to track only a representative agent to describe the behavior of the population. Let the state, and action of this representative at time t be denoted as $x_t \in \mathcal{X}$, $u_t \in \mathcal{U}$ respectively while the distributions of states, and actions of the whole population be denoted as $\boldsymbol{\mu}_t^\infty \in \mathcal{P}(\mathcal{X})$, and $\boldsymbol{\nu}_t^\infty \in \mathcal{P}(\mathcal{U})$. Observe that, for a given policy sequence, $\boldsymbol{\pi} \triangleq \{\pi_t\}_{t \in \{0, 1, \dots\}}$, the action distribution, $\boldsymbol{\nu}_t^\infty$ can be derived from the state distribution, $\boldsymbol{\mu}_t^\infty$ as follows.

$$\boldsymbol{\nu}_t^\infty = \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) \triangleq \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^\infty) \boldsymbol{\mu}_t^\infty(x) \quad (8)$$

Similarly, the state distribution at time $t+1$ can be obtained from $\boldsymbol{\mu}_t^\infty$ using the following update equation.

$$\boldsymbol{\mu}_{t+1}^\infty = P^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) \triangleq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} P(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \quad (9)$$

Finally, the average reward, and constraint cost at time t can be computed as

$$r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) = \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} r(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \quad (10)$$

$$c^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) = \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} c(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \quad (11)$$

For an initial state distribution, $\boldsymbol{\mu}_0^\infty = \boldsymbol{\mu}_0$, and a policy-sequence $\boldsymbol{\pi}$, we define the value functions related to reward, and constraint cost as follows.

$$V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}) = \sum_{t=0}^{\infty} \gamma^t r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) \quad (12)$$

$$V_\infty^C(\boldsymbol{\mu}_0, \boldsymbol{\pi}) = \sum_{t=0}^{\infty} \gamma^t c^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) \quad (13)$$

The goal of CMFC is to solve the following optimization problem.

$$\sup_{\boldsymbol{\pi} \in \Pi_\infty} V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}) \quad \text{subject to : } V_\infty^C(\boldsymbol{\mu}_0, \boldsymbol{\pi}) \leq 0 \quad (\text{CMFC})$$

where Π_∞ , as stated in section 2, is the set of admissible policy sequences. In the forthcoming section, we shall establish that the solution of a variant of (CMFC) closely approximates the solution of (CMARL) for large N .

4. Main Result: Approximation of CMARL via CMFC

Before stating the approximation result, we would like to enlist the set of assumptions needed to prove it. Our first assumption is on the reward, cost, and state transition function.

Assumption 1 *There exists constants $M_R, M_C, L_R, L_C, L_P > 0$ such that the following relations hold $\forall x \in \mathcal{X}, \forall u \in \mathcal{U}, \forall \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \mathcal{P}(\mathcal{X}),$ and $\forall \boldsymbol{\nu}_1, \boldsymbol{\nu}_2 \in \mathcal{P}(\mathcal{U})$.*

- (a) $|r(x, u, \boldsymbol{\mu}_1, \boldsymbol{\nu}_1)| \leq M_R,$
- (b) $|c(x, u, \boldsymbol{\mu}_1, \boldsymbol{\nu}_1)| \leq M_C,$
- (c) $|r(x, u, \boldsymbol{\mu}_1, \boldsymbol{\nu}_1) - r(x, u, \boldsymbol{\mu}_2, \boldsymbol{\nu}_2)| \leq L_R [|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2|_1 + |\boldsymbol{\nu}_1 - \boldsymbol{\nu}_2|_1],$
- (d) $|c(x, u, \boldsymbol{\mu}_1, \boldsymbol{\nu}_1) - c(x, u, \boldsymbol{\mu}_2, \boldsymbol{\nu}_2)| \leq L_C [|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2|_1 + |\boldsymbol{\nu}_1 - \boldsymbol{\nu}_2|_1],$
- (e) $|P(x, u, \boldsymbol{\mu}_1, \boldsymbol{\nu}_1) - P(x, u, \boldsymbol{\mu}_2, \boldsymbol{\nu}_2)|_1 \leq L_P [|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2|_1 + |\boldsymbol{\nu}_1 - \boldsymbol{\nu}_2|_1]$

where $|\cdot|_1$ denotes L_1 -distance.

Assumption 1(a), and 1(b) states that the reward, r , and the cost function, c are bounded. The very definition of the transition function, P makes it bounded. Hence, it is not listed as an assumption. On the other hand, Assumption 1(c)-(e) dictate that the functions r , c , and P are Lipschitz continuous with respect to their state, and action distribution arguments. Assumption 1 is common in the literature (Angiuli et al., 2022; Carmona et al., 2018). Our next assumption is on the class of admissible policies, Π .

Notation	Interpretation	Used In
π_N^*	Optimal solution of (CMARL)	Theorem 2
π_0^*	Optimal solution of (CMARL)	Section 5
π_∞^*	Optimal solution of (G-CMFC) for $\zeta = -G_C$	Theorem 2, 4
π_1^*	Optimal solution of (G-CMFC) for $\zeta = -G_C$	Section 5
π_2^*	Optimal solution of (G-CMARL) for $\zeta = -2G_C$	Section 5
$V_N^R(\mathbf{x}_0^N, \pi)$, $V_N^C(\mathbf{x}_0^N, \pi)$	Reward, constraint values of N -agent system for initial states \mathbf{x}_0^N , and policy sequence π	Theorem 2, 4, 5, Lemma 3
$V_\infty^R(\boldsymbol{\mu}_0, \pi)$, $V_\infty^C(\boldsymbol{\mu}_0, \pi)$	Reward, constraint values of MFC system for initial distribution $\boldsymbol{\mu}_0$, and policy sequence π	Theorem 2, 4, 5, Lemma 3
$V_N^*(\mathbf{x}_0^N, \zeta)$, $V_\infty^*(\mathbf{x}_0^N, \zeta)$	Optimal objective values of (G-CMARL) and (G-CMFC) respectively with constraint ζ	Lemma 1, Theorem 4, 5 Section 5

Table 1: Notations. Note that $V_N^*(\mathbf{x}_0^N, -2G_C) = V_N^R(\mathbf{x}_0^N, \pi_2^*)$, $V_\infty^*(\boldsymbol{\mu}_0, -G_C) = V_\infty^R(\boldsymbol{\mu}_0, \pi_1^*) = V_\infty^R(\boldsymbol{\mu}_0, \pi_\infty^*)$, and $V_N^*(\mathbf{x}_0^N, 0) = V_N^R(\mathbf{x}_0^N, \pi_0^*) = V_N^R(\mathbf{x}_0^N, \pi_N^*)$.

Assumption 2 *There exists a constant $L_Q > 0$ such that $\forall \pi \in \Pi$ the following holds.*

$$|\pi(x, \boldsymbol{\mu}_1) - \pi(x, \boldsymbol{\mu}_2)| \leq L_Q |\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2|_1$$

$\forall x \in \mathcal{X}$, and $\forall \boldsymbol{\mu}_1, \boldsymbol{\mu}_2 \in \mathcal{P}(\mathcal{X})$.

Assumption 2 dictates that the admissible policies are Lipschitz continuous with respect to the state distributions. The assumption stated above is also common in the literature and holds for neural network-based policies with bounded weights (Gu et al., 2021; Pásztor et al., 2023). Before stating the main result, we need to introduce some notations. Consider the following optimization problem.

$$\sup_{\pi \in \Pi_\infty} V_N^R(\mathbf{x}_0^N, \pi) \quad \text{subject to : } V_N^C(\mathbf{x}_0^N, \pi) \leq \zeta \quad (\text{G-CMARL})$$

Problem (G-CMARL) is a generalization of the problem (CMARL). Specifically, it uses an arbitrary real, $\zeta \in \mathbb{R}$, as the constraint upper bound, instead of $\zeta = 0$ as considered in (CMARL). With slight abuse of notation, we define the optimal objective value of (G-CMARL) as $V_N^*(\mathbf{x}_0^N, \zeta)$. In a similar fashion, $V_\infty^*(\boldsymbol{\mu}_0, \zeta)$ denotes the optimal objective value of the following optimization problem for $\zeta \in \mathbb{R}$.

$$\sup_{\pi \in \Pi_\infty} V_\infty^R(\boldsymbol{\mu}_0, \pi) \quad \text{subject to : } V_\infty^C(\boldsymbol{\mu}_0, \pi) \leq \zeta \quad (\text{G-CMFC})$$

Clearly, (G-CMFC) generalizes the problem, (CMFC). Table 1 summarizes the notations used in the paper. The following Assumption is required (in addition to Assumption 1, 2) to establish the main result.

Assumption 3 *There exists $\zeta_0 > 0$ such that the set of feasible solutions of (G-CMARL) for $\zeta = -\zeta_0$ is non-empty. In other words, there exists a policy-sequence, $\pi \in \Pi_\infty$ such that $V_N^C(\mathbf{x}_0^N, \pi) \leq -\zeta_0$ for any given joint initial state, \mathbf{x}_0^N .*

Assumption 3 ensures that no pathological example with an empty feasible set arises in our analysis. This assumption is similar to Slater’s condition in the optimization literature. Below we state an intermediate result that will be useful to prove our main theorem.

Lemma 1 *Let \mathbf{x}_0^N denote the initial joint state in an N -agent system, and $\boldsymbol{\mu}_0$ be its empirical distribution. If assumptions 1 – 3 hold, then there exists a sufficiently large $N_0 > 0$ such that $\forall N \geq N_0$ the following inequality holds whenever $\gamma S_P < 1$.*

$$|V_N^*(\mathbf{x}_0^N, 0) - V_\infty^*(\boldsymbol{\mu}_0, -G_C)| \leq G_R + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1 - \gamma} \right) \right] \quad (14)$$

The terms G_R, G_C are defined as shown below.

$$G_J \triangleq \left(\frac{1}{1 - \gamma} \right) \left[\frac{M_J}{\sqrt{N}} + \frac{L_J \sqrt{|\mathcal{U}|}}{\sqrt{N}} \right] + \frac{[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|}]}{\sqrt{N}} \left(\frac{S_J C_P}{S_P - 1} \right) \left[\frac{1}{1 - \gamma S_P} - \frac{1}{1 - \gamma} \right] \quad (15)$$

where $C_P \triangleq 2 + L_P$, $S_J \triangleq (M_J + 2L_J) + L_Q(M_J + L_J)$, $S_P \triangleq (1 + 2L_J) + L_Q(1 + L_J)$, and $J \in \{R, C\}$. The terms $\{M_R, M_C, L_R, L_C, L_P, \zeta_0\}$ are defined in Assumptions 1, 2, and 3.

Lemma 1 provides a recipe for approximately solving the N -agent constrained MARL problem (CMARL). In particular, inequality (14) ensures that the optimal value function of N -agent problem (CMARL), and that of generalized constrained mean-field problem (G-CMFC) for $\zeta = -G_C$ are at most $\mathcal{O}(G_R + G_C)$ distance away from each other. Note that $G_R, G_C = \mathcal{O}\left(\frac{1}{\sqrt{N}}\right)$. Thus for sufficiently large N , obtaining a solution for (G-CMFC) for $\zeta = -G_C$ is approximately equivalent to solving the original N -agent problem (CMARL). We would like to mention that, $G_R, G_C = \mathcal{O}(\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|})$ where $|\mathcal{X}|, |\mathcal{U}|$ denote the sizes of state, and action spaces respectively. Therefore, as the state and action spaces of individual agents become larger, the approximation progressive becomes worse.

Although (14) establishes the proximity of the value-functions of N -agent, and infinite agent systems, it does not clarify whether the policy sequence obtained by solving (G-CMFC) for $\zeta = -G_C$ would obey the constraint of (CMARL) if the policy is adopted in an N -agent system. It also does not comment on how the said policy sequence would perform (in terms of cumulative reward) in an N -agent system in comparison to the optimal N -agent value. These important questions are answered in the following Theorem. The proof of Theorem 2 is relegated to Appendix A.

Theorem 2 *Let $\boldsymbol{\pi}_N^*, \boldsymbol{\pi}_\infty^*$ be solutions of (CMARL), and (G-CMFC) with $\zeta = -G_C$ respectively. If Assumptions 1 – 3 hold, then the following inequality holds for $\forall N \geq N_0$ whenever one has $\gamma S_P < 1$.*

$$V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_N^*) - V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \leq 2G_R + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1 - \gamma} \right) \right], \quad (16)$$

$$\text{and } V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \leq 0 \quad (17)$$

where $V_N^R(\cdot, \cdot)$, and $V_N^C(\cdot, \cdot)$ are defined in (6), and (7) respectively. The notations $\mathbf{x}_0^N, \boldsymbol{\mu}_0, G_R, G_C, S_P$, and N_0 carry the same meaning as stated in Lemma 1.

Inequalities (16), (17) dictate that if the policy-sequence, π_∞^* , obtained by solving the mean-field problem (G-CMFC) for $\zeta = -G_C$ is adopted in the N -agent system, then the value obtained for such a system is $\mathcal{O}(\frac{1}{\sqrt{N}})$ error away from the optimum. Also, π_∞^* does not violate (CMARL) constraint.

5. Proof of Lemma 1

The proof of Lemma 1 hinges on Lemma 3. The proof of Lemma 3 is given in Appendix B.

Lemma 3 *Let \mathbf{x}_0^N denote the initial joint state in an N -agent system, and $\boldsymbol{\mu}_0$ be its empirical distribution. If Assumption 1 – 3 hold, then there exists sufficiently large $N_0 > 0$ such that $\forall N \geq N_0$, and $\forall \boldsymbol{\pi} \in \Pi_\infty$, the following inequalities hold whenever $\gamma S_P < 1$.*

$$|V_N^J(\mathbf{x}_0^N, \boldsymbol{\pi}) - V_\infty^J(\boldsymbol{\mu}_0, \boldsymbol{\pi})| \leq G_J \quad (18)$$

where $J \in \{R, C\}$. The value functions $V_N^R, V_N^C, V_\infty^R, V_\infty^C$ are provided by (6), (7), (12), (13) respectively, and the terms G_R, G_C, S_P are defined in Lemma 1.

Intuitively, Lemma 3 dictates that the value function pairs $(V_N^R, V_\infty^R), (V_N^C, V_\infty^C)$ are close in the sense that for every policy-sequence $\boldsymbol{\pi}$, the differences $|V_N^J(\mathbf{x}_0^N, \boldsymbol{\pi}) - V_\infty^J(\boldsymbol{\mu}_0, \boldsymbol{\pi})|$, $J \in \{R, C\}$ are of the form $\mathcal{O}(\frac{1}{\sqrt{N}})$. Below we explain how such a result helps us to establish Lemma 1. Consider the three distinct constrained maximizations stated below.

- Problem 0: (G-CMARL) with $\zeta = 0$. It is the same as the original N -agent problem (CMARL).
- Problem 1: (G-CMFC) with $\zeta = -G_C$.
- Problem 2: (G-CMARL) with $\zeta = -2G_C$.

Let, $\Pi_\infty^k \subset \Pi_\infty$ be the set of admissible policy sequences that satisfy the constraint of k -th optimization problem stated above, $k \in \{0, 1, 2\}$. Note that, if $\boldsymbol{\pi} \in \Pi_\infty^2$, then

$$V_\infty^C(\boldsymbol{\mu}_0, \boldsymbol{\pi}) \stackrel{(a)}{\leq} V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}) + G_C \stackrel{(b)}{\leq} -G_C \quad (19)$$

Inequality (a) follows from (18) and (b) is a consequence of the fact that $\boldsymbol{\pi} \in \Pi_\infty^2$. Thus, $\Pi_\infty^2 \subset \Pi_\infty^1$. Similarly, if $\boldsymbol{\pi} \in \Pi_\infty^1$, then

$$V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}) \stackrel{(a)}{\leq} V_\infty^C(\boldsymbol{\mu}_0, \boldsymbol{\pi}) + G_C \stackrel{(b)}{\leq} 0 \quad (20)$$

which shows $\Pi_\infty^1 \subset \Pi_\infty^0$. Combining, we get $\Pi_\infty^2 \subset \Pi_\infty^1 \subset \Pi_\infty^0$. Finally, assume that $\boldsymbol{\pi}_k^* \in \Pi_\infty^k$ indicate the solution of k -th maximization problem described above, $k \in \{0, 1, 2\}$. Then,

$$V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_1^*) \stackrel{(a)}{\leq} V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_1^*) + G_R \stackrel{(b)}{\leq} V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_0^*) + G_R = V_N^*(\mathbf{x}_0^N, 0) + G_R, \quad (21)$$

$$V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_1^*) \stackrel{(c)}{\geq} V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_2^*) \stackrel{(d)}{\geq} V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_2^*) - G_R = V_N^*(\mathbf{x}_0^N, -2G_C) - G_R \quad (22)$$

Inequalities (a), and (d) follow from (18). Relation (b) is a consequence of the fact that π_0^* is a maximizer of $V_N^R(\mathbf{x}_0^N, \cdot)$ over Π_∞^0 , and $\pi_1^* \in \Pi_\infty^1 \subset \Pi_\infty^0$. Similarly, π_1^* is a maximizer of $V_\infty^R(\boldsymbol{\mu}_0, \cdot)$ over Π_∞^1 , and $\pi_2^* \in \Pi_\infty^2 \subset \Pi_\infty^1$. This establishes relation (c).

The gist of the above exercise is that the value $V_\infty^*(\boldsymbol{\mu}_0, -G_C)$ is sandwiched between $V_N^*(\mathbf{x}_0^N, 0) + G_R$, and $V_N^*(\mathbf{x}_0^N, -2G_C) - G_R$. To prove (14), one needs to bound the distance between $V_N^*(\mathbf{x}_0^N, 0)$, and $V_N^*(\mathbf{x}_0^N, -2G_C)$. This can be accomplished utilizing the concavity property of $V_N^*(\mathbf{x}_0^N, \cdot)$ in the domain $[-\zeta_0, 0]$ where ζ_0 is given in Assumption 3. The concavity property is a consequence of Proposition 1 of (Paternain et al., 2019). If $N \geq N_0$ where N_0 is a sufficiently large value, then $2G_C < \zeta_0$. In such a case,

$$V_N^*(\mathbf{x}_0^N, -2G_C) \geq \left(1 - \frac{2G_C}{\zeta_0}\right) V_N^*(\mathbf{x}, 0) + \frac{2G_C}{\zeta_0} V_N^*(\mathbf{x}_0^N, -\zeta_0)$$

$$\text{Equivalently, } V_N^*(\mathbf{x}_0^N, 0) - V_N^*(\mathbf{x}, -2G_C) \leq \frac{2G_C}{\zeta_0} [V_N^*(\mathbf{x}_0^N, 0) - V_N^*(\mathbf{x}_0^N, -\zeta_0)]$$

Using Assumption 1(a), one can trivially deduce, $|V_N^*(\mathbf{x}_0^N, \zeta)| \leq M_R/(1 - \gamma)$, $\forall \zeta \geq -\zeta_0$. More specifically, this is true for $\zeta = -\zeta_0, 0$. Hence,

$$V_N^*(\mathbf{x}_0^N, 0) - V_N^*(\mathbf{x}_0^N, -2G_C) \leq G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1 - \gamma} \right) \right] \quad (23)$$

Using (22), (23), we obtain,

$$V_N^*(\mathbf{x}_0^N, 0) - V_\infty^*(\boldsymbol{\mu}_0, -G_C) \leq G_R + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1 - \gamma} \right) \right] \quad (24)$$

Combining with (21), we conclude the result.

6. Improvement of Optimality Gap in a Special Case

In this section, we shall impose the following additional restriction on the structure of reward, r , cost, c , and transition function, P to improve the approximation bound of Lemma 1.

Assumption 4 *The functions r , c , and P are independent of the action distribution of the population. Mathematically, $\forall x \in \mathcal{X}, \forall u \in \mathcal{U}, \forall \boldsymbol{\mu} \in \mathcal{P}(\mathcal{X})$, and $\forall \boldsymbol{\nu} \in \mathcal{P}(\mathcal{U})$,*

$$(a) \ r(x, u, \boldsymbol{\mu}, \boldsymbol{\nu}) = r(x, u, \boldsymbol{\mu}), (b) \ c(x, u, \boldsymbol{\mu}, \boldsymbol{\nu}) = c(x, u, \boldsymbol{\mu}), (c) \ P(x, u, \boldsymbol{\mu}, \boldsymbol{\nu}) = P(x, u, \boldsymbol{\mu})$$

Assumption 4 removes the dependence of r, c , and P on the action distribution. However, for each agent, those functions still take the action executed by the same agent as an argument. Below we present our improved approximation result.

Theorem 4 *Let \mathbf{x}_0^N denote an initial joint state in an N -agent system, and $\boldsymbol{\mu}_0$ its empirical distribution. If Assumptions 1 – 4 are true, then there exists a sufficiently large $N_0 > 0$ such that $\forall N \geq N_0$ the following inequality holds whenever $\gamma S_P < 1$.*

$$|V_N^*(\mathbf{x}_0^N, 0) - V_\infty^*(\boldsymbol{\mu}_0, -G_C^0)| \leq G_R^0 + G_C^0 \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1 - \gamma} \right) \right] \quad (25)$$

The terms G_R^0, G_C^0 are defined as shown below.

$$G_J^0 \triangleq \frac{M_J}{\sqrt{N}} \left(\frac{1}{1-\gamma} \right) + \frac{\sqrt{|\mathcal{X}|}}{\sqrt{N}} \left(\frac{2S_J}{S_P-1} \right) \left[\frac{1}{1-\gamma S_P} - \frac{1}{1-\gamma} \right] \quad (26)$$

where $J \in \{R, C\}$, and the terms $M_R, M_C, S_R, S_C, S_P, \zeta_0$ are given in Lemma 1. Additionally, if π_∞^* is a solution of (G-CMFC) with $\zeta = -G_C^0$, then the following inequalities hold $\forall N \geq N_0$ whenever $\gamma S_P < 1$.

$$V_N^*(\mathbf{x}_0^N, 0) - V_N^R(\mathbf{x}_0^N, \pi_\infty^*) \leq 2G_R^0 + G_C^0 \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1-\gamma} \right) \right], \quad (27)$$

$$\text{and } V_N^C(\mathbf{x}_0^N, \pi_\infty^*) \leq 0 \quad (28)$$

where $V_N^R(\cdot, \cdot)$, and $V_N^C(\cdot, \cdot)$ are defined in (6), and (7) respectively.

Inequality (25) states that the optimal value function obtained by solving (CMARL) is approximated by the optimal value function obtained by solving (G-CMFC) with $\zeta = -G_C$ within an error of $\mathcal{O}(\sqrt{|\mathcal{X}|}/\sqrt{N})$. On the other hand, (27), (28) suggest that if the optimal policy sequence obtained by solving (G-CMFC), with $\zeta = -G_C$ is adopted in an N -agent system, then the value generated in such a system is at most $\mathcal{O}(\sqrt{|\mathcal{X}|}/\sqrt{N})$ distance away from the optimal N -agent value function. Moreover, the said policy does not violate the constraint of (CMARL).

Note that, the dependence of the approximation error on N is still $\mathcal{O}(1/\sqrt{N})$. However, its dependence on the sizes of state, and action spaces has been reduced to $\mathcal{O}(\sqrt{|\mathcal{X}|})$ from $\mathcal{O}(\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|})$ stated in Lemma 1. Therefore, the stated approximation result may be useful in those situations where reward, cost, and transition functions are independent of the action distribution, and the size of action space of individual agents is large. Interestingly, we could not derive an approximation error that is independent of the size of the state space, $|\mathcal{X}|$, by imposing the restriction that r, c , and P are independent of the state distribution. This indicates an inherent asymmetry between the roles played by state, and action spaces in mean-field approximation.

7. Natural Policy Gradient Algorithm to Solve CMFC Problem

In section 4, we demonstrated that the N -agent problem (CMARL) is well-approximated by the mean-field problem (G-CMFC) with appropriate choice of ζ . In this section, we shall discuss how one can employ a Natural Policy Gradient (NPG) based algorithm to approximately solve (G-CMFC). Recall that in a mean-field setup, it is sufficient to track only one representative agent. At time t , the representative chooses an action $u_t \in \mathcal{U}$ based upon its observation of its own state, $x_t \in \mathcal{X}$, and the mean-field state distribution, $\mu_t^\infty \in \mathcal{P}(\mathcal{X})$. Thus, (G-CMFC) can be described as a constrained single agent problem with state space $\mathcal{X} \times \mathcal{P}(\mathcal{X})$, and action space, \mathcal{U} . Without loss of generality, we can therefore assume the policy-sequences to be stationary (Dolgov and Durfee, 2005). With slight abuse of notations, we denote both an arbitrary policy, and its associated stationary sequence by the same notation, π . The class of all admissible policies is, Π . Let, the elements of Π be parameterized by the parameter, $\Phi \in \mathbb{R}^d$. For a given policy, π_Φ , the Q -function, $Q_\Phi^R(\cdot, \cdot, \cdot)$, value function,

$V_{\Phi}^R(\cdot, \cdot)$, and the advantage function, $A_{\Phi}^R(\cdot, \cdot, \cdot)$ are defined as follows.

$$Q_{\Phi}^R(x, \boldsymbol{\mu}, u) \triangleq \mathbb{E} \left[\sum_{t=0}^{\infty} \gamma^t r(x_t, u_t, \boldsymbol{\mu}_t^{\infty}, \boldsymbol{\nu}_t^{\infty}) \middle| x_0 = x, \boldsymbol{\mu}_0 = \boldsymbol{\mu}, u_0 = u \right], \quad (29)$$

$$V_{\Phi}^R(x, \boldsymbol{\mu}) \triangleq \mathbb{E} [Q_{\Phi}^R(x, \boldsymbol{\mu}, u)], \quad A_{\Phi}^R(x, \boldsymbol{\mu}, u) = Q_{\Phi}^R(x, \boldsymbol{\mu}, u) - V_{\Phi}^R(x, \boldsymbol{\mu}) \quad (30)$$

The expectation in (29) is obtained over $x_{t+1} \sim P(x_t, u_t, \boldsymbol{\mu}_t, \boldsymbol{\nu}_t)$, $u_t \sim \pi_{\Phi}(x_t, \boldsymbol{\mu}_t)$, $\forall t \in \{0, 1, \dots\}$. The deterministic quantities, $\{\boldsymbol{\mu}_t^{\infty}, \boldsymbol{\nu}_t^{\infty}\}_{t \in \{0, 1, \dots\}}$ are evaluated using relations (9), and (8) respectively. On the other hand, the expectation in (30) is computed over, $u \sim \pi_{\Phi}(x, \boldsymbol{\mu})$. The functions, Q_{Φ}^C , V_{Φ}^C , and A_{Φ}^C are defined similarly for the cost function, c . In the parametric form, (G-CMFC) can be rewritten as the following primal problem.

$$\sup_{\Phi \in \mathbb{R}^d} V_{\infty}^R(\boldsymbol{\mu}_0, \pi_{\Phi}) \quad \text{subject to: } V_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi}) \leq \zeta \quad (\text{PRIMAL})$$

Its corresponding dual problem is as follows.

$$\sup_{\Phi \in \mathbb{R}^d} \inf_{\lambda \geq 0} V_{\infty}^R(\boldsymbol{\mu}_0, \pi_{\Phi}) + \lambda [\zeta - V_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi})] \quad (\text{DUAL})$$

Let $\Phi^* \in \mathbb{R}^d$ be a solution of (DUAL). Ideally, Φ^* should be obtained via the following primal-dual natural policy gradient (NPG) updates that start from Φ_0 , $\lambda_0 = 0$ and utilize $\eta_1, \eta_2 > 0$ as learning parameters (Ding et al., 2020).

$$\Phi_{j+1} = \Phi_j + \frac{\eta_1}{1 - \gamma} \mathbf{w}_j, \quad \mathbf{w}_j \triangleq \arg \min_{\mathbf{w} \in \mathbb{R}^d} E_{\xi_{\boldsymbol{\mu}_0}^{\Phi_j}}(\mathbf{w}, \Phi_j, \lambda_j), \quad (31)$$

$$\lambda_{j+1} = P_{[0, \infty)} [\lambda_j + \eta_2 [V_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi_j}) - \zeta]] \quad (32)$$

The function, $P_{[0, \infty)}$ projects its argument onto the set $[0, \infty)$. The term $\xi_{\boldsymbol{\mu}_0}^{\Phi_j}$ indicates the occupancy measure induced by policy, π_{Φ_j} from the initial distribution, $\boldsymbol{\mu}_0$. The function, $E_{\xi_{\boldsymbol{\mu}_0}^{\Phi_j}}(\cdot, \cdot, \cdot)$ and the occupancy measure $\xi_{\boldsymbol{\mu}_0}^{\Phi_j}$ are mathematically defined below.

$$\xi_{\boldsymbol{\mu}_0}^{\Phi_j} \triangleq \sum_{\tau=0}^{\infty} \gamma^{\tau} (x, \boldsymbol{\mu}, u) \mathbb{P}(x_{\tau} = x, \boldsymbol{\mu}_{\tau} = \boldsymbol{\mu}, u_{\tau} = u | \boldsymbol{\mu}_0 = \boldsymbol{\mu}, \pi_{\Phi_j})(1 - \gamma), \quad (33)$$

$$E_{\xi_{\boldsymbol{\mu}_0}^{\Phi_j}}(\mathbf{w}, \Phi_j, \lambda_j) \triangleq \mathbb{E}_{(x, \boldsymbol{\mu}, u) \sim \xi_{\boldsymbol{\mu}_0}^{\Phi_j}} \left[\left(A_{\Phi_j}^{\lambda_j}(x, \boldsymbol{\mu}, u) - \mathbf{w}^T \nabla_{\Phi_j} \log \pi_{\Phi_j}(x, \boldsymbol{\mu})(u) \right)^2 \right] \quad (34)$$

where $A_{\Phi_j}^{\lambda_j}(\cdot, \cdot, \cdot)$ is defined as $A_{\Phi_j}^{\lambda_j} \triangleq A_{\Phi_j}^R - \lambda_j A_{\Phi_j}^C$. Due to the sampling-based nature of the algorithm, we focus on the following approximate updates.

$$\Phi_{j+1} = \Phi_j + \frac{\eta_1}{1 - \gamma} \hat{\mathbf{w}}_j, \quad (35)$$

$$\lambda_{j+1} = P_{[0, \infty)} \left[\lambda_j + \eta_2 [\hat{V}_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi_j}) - \zeta] \right] \quad (36)$$

where $\hat{V}_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi_j})$ is an unbiased estimate of $V_{\infty}^C(\boldsymbol{\mu}_0, \pi_{\Phi_j})$, which is obtained via Algorithm 2 (Appendix O). To obtain $\hat{\mathbf{w}}_j$ (approximation of \mathbf{w}_j), we need to solve another minimization problem. We apply the stochastic gradient descent (SGD) updates described as follows to

solve this sub-problem: $\mathbf{w}_{j,l+1} = \mathbf{w}_{j,l} - \alpha \mathbf{h}_{j,l}$ where the gradient $\mathbf{h}_{j,l}$ is given below (Ding et al., 2020), and $\alpha > 0$ is the learning parameter.

$$\mathbf{h}_{j,l} \triangleq \left(\mathbf{w}_{j,l}^\top \nabla_{\Phi_j} \log \pi_{\Phi_j}(x, \boldsymbol{\mu})(u) - \hat{A}_{\Phi_j}^{\lambda_j}(x, \boldsymbol{\mu}, u) \right) \nabla_{\Phi_j} \log \pi_{\Phi_j}(x, \boldsymbol{\mu})(u) \quad (37)$$

where $(x, \boldsymbol{\mu}, u) \sim \zeta_{\boldsymbol{\mu}_0}^{\Phi_j}$, and $\hat{A}_{\Phi_j}^{\lambda_j}$ is an unbiased estimator of $A_{\Phi_j}^{\lambda_j} \triangleq A_{\Phi_j}^R - \lambda_j A_{\Phi_j}^C$. The sampling process is detailed in Algorithm 2 (Appendix O). Algorithm 1 summarizes the NPG process.

Algorithm 1 Natural Policy Gradient Algorithm to solve the Dual Problem

Input: η_1, η_2, α : Learning rates, J, L : Number of execution steps

$\mathbf{w}_0, \Phi_0, \lambda_0 = 0$: Initial parameters, $\boldsymbol{\mu}_0$: Initial state distribution, Constraint bound: ζ

- 1: **for** $j \in \{0, 1, \dots, J-1\}$ **do**
- 2: $\mathbf{w}_{j,0} \leftarrow \mathbf{w}_0$
- 3: **for** $l \in \{0, 1, \dots, L-1\}$ **do**
- 4: Sample $(x, \boldsymbol{\mu}, u) \sim \zeta_{\boldsymbol{\mu}_0}^{\Phi_j}$ and $\hat{A}_{\Phi_j}^{\lambda_j}(x, \boldsymbol{\mu}, u)$ using Algorithm 2
- 5: Compute $\mathbf{h}_{j,l}$ using (37)
- 6: $\mathbf{w}_{j,l+1} \leftarrow \mathbf{w}_{j,l} - \alpha \mathbf{h}_{j,l}$
- 7: **end for**
- 8: $\hat{\mathbf{w}}_j \leftarrow \frac{1}{L} \sum_{l=1}^L \mathbf{w}_{j,l}$
- 9: $\Phi_{j+1} \leftarrow \Phi_j + \frac{\eta_1}{1-\gamma} \hat{\mathbf{w}}_j$
- 10: Sample $(x, \boldsymbol{\mu}, u) \sim \zeta_{\boldsymbol{\mu}_0}^{\Phi_j}$ and $\hat{V}_{\Phi_j}^C(x, \boldsymbol{\mu})$ using Algorithm 2
- 11: $\lambda_{j+1} \leftarrow \text{P}_{[0,\infty)} \left[\lambda_j + \eta_2 [\hat{V}_{\Phi_j}^C(x, \boldsymbol{\mu}) - \zeta] \right]$
- 12: **end for**

Output: $\{\Phi_1, \dots, \Phi_J\}$: Policy parameters

Recall that Lemma 1 proves that the optimal solution of (G-CMFC) with appropriate choice of ζ approximately solves the N -agent problem (CMARL). However, Algorithm 1 can only approximately solve (G-CMFC). One, therefore, naturally asks whether the solution of Algorithm 1 is close to the optimal solution of (CMARL). Theorem 5 (stated below) provides an affirmative answer to this question. The proof of Theorem 5 is relegated to Appendix N. The following assumptions are needed to establish this result.

Assumption 5 *There exists $\zeta_1 < 0$ such that the (G-CMFC) problem associated with $\zeta = \zeta_1$ has a feasible solution.*

Assumption 6 $\forall \Phi \in \mathbb{R}^d, \forall \boldsymbol{\mu} \in \mathcal{P}(\mathcal{X}), \forall x \in \mathcal{X}, \forall u \in \mathcal{U},$

$$|\nabla_{\Phi} \log \pi_{\Phi}(x, \boldsymbol{\mu})(u)|_1 \leq G$$

for some positive constant G .

Assumption 7 $\forall \Phi_1, \Phi_2 \in \mathbb{R}^d, \forall \boldsymbol{\mu} \in \mathcal{P}(\mathcal{X}), \forall x \in \mathcal{X}, \forall u \in \mathcal{U},$

$$|\nabla_{\Phi_1} \log \pi_{\Phi_1}(x, \boldsymbol{\mu})(u) - \nabla_{\Phi_2} \log \pi_{\Phi_2}(x, \boldsymbol{\mu})(u)|_1 \leq M |\Phi_1 - \Phi_2|_1$$

for some positive constant M .

Assumption 8 $\forall \Phi \in \mathbb{R}^d, \forall \mu_0 \in \mathcal{P}(\mathcal{X})$,

$$E_{\xi_{\mu_0}^*}(\mathbf{w}_{\Phi}^*, \Phi) \leq \epsilon_{\text{bias}}, \quad \mathbf{w}_{\Phi}^* \triangleq \arg \min_{\mathbf{w} \in \mathbb{R}^d} E_{\xi_{\mu_0}^*}(\mathbf{w}, \Phi)$$

where Φ^* is the parameter of the optimal policy.

Assumption 9 The gradient iterates $\{\mathbf{w}_{j,l}\}_{j \in \{0, \dots, J-1\}, l \in \{0, \dots, L-1\}}$, and $\{\mathbf{w}_j\}_{j \in \{0, \dots, J-1\}}$ of Algorithm 1 are such that,

$$\mathbb{E} [|\mathbf{w}_{j,l}|_1^2] \leq W_0^2, \text{ and } \mathbb{E} [|\mathbf{w}_j|_1^2] \leq W_1^2$$

for some constants W_0, W_1 .

Assumption 5 expresses Slater’s condition for the (G-CMFC) problem. Assumption 6, and 7 ensure that the log-likelihood function is Lipschitz and smooth with respect to the parameters. Assumption 8 says that the expressivity error of the parameterized policy class is bounded by a term ϵ_{bias} . If the policy class is complete (e.g., in softmax parameterization), then $\epsilon_{\text{bias}} = 0$. Although, in general, $\epsilon_{\text{bias}} > 0$, for rich policy classes (e.g., where policies are represented by dense/wide neural networks), its value is negligibly small. Finally, Assumption 9 dictates that the gradient estimates used in our algorithm are bounded. All of the above assumptions are standard in the policy gradient literature. For a detailed discussion on their validity, see Ding et al. (2020).

Theorem 5 Let \mathbf{x}_0^N be the initial joint state in an N -agent system, and μ_0 its empirical distribution. Let $\{\Phi_j\}_{j=1}^J$ denote the policy parameters yielded from Algorithm 1 for an initial parameter, Φ_0 , and $\zeta = -2G_C$. If Assumptions 1 – 3, and 5 – 9 hold, then for appropriate choices of $\eta_1, \eta_2, \alpha, J, L$, and sufficiently large N , the following relations hold when $\gamma S_P < 1$.

$$\left| V_N^*(\mathbf{x}_0^N, 0) - \frac{1}{J} \sum_{j=1}^J V_{\infty}^R(\mu_0, \pi_{\Phi_j}) \right| \leq K_1 e + K_2 \sqrt{\epsilon_{\text{bias}}}, \quad (38)$$

$$\frac{1}{J} \sum_{j=1}^J V_N^C(\mathbf{x}_0^N, \pi_{\Phi_j}) \leq 0 \quad (39)$$

where ϵ_{bias} is defined in Assumption 8, K_1, K_2 are constants, and $e \triangleq \frac{1}{\sqrt{N}}[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|}]$. The sample complexity of the process is $\mathcal{O}(e^{-6})$.

Theorem 5 states that the solution of (G-CMFC) with $\zeta = -2G_C$, obtained via Algorithm 1 closely (within an error of $\mathcal{O}(e)$) approximates the optimal objective value obtained by solving the N -agent problem (CMARL). Moreover, the obtained policy also satisfies the constraint of (CMARL). The sample complexity of the whole process is $\mathcal{O}(e^{-6})$. Since the submission of this paper, some new developments have occurred in the field of algorithm design for CMDPs (Bai et al., 2023; Mondal and Aggarwal, 2024b). In particular, Mondal and Aggarwal (2024b) have recently proposed an algorithm that can achieve $e + \sqrt{\epsilon_{\text{bias}}}$ optimality error and zero constraint violation with $\tilde{\mathcal{O}}(e^{-2})$ sample complexity. We believe that the result of Theorem 5 can be improved to $\tilde{\mathcal{O}}(e^{-2})$ if one utilizes the algorithm suggested by (Mondal and Aggarwal, 2024b) for solving our CMARL problem. However, the verification of such a possibility is left for future research.

8. Experimental Results

We consider the following setting (taken from (Subramanian and Mahajan, 2019) with slight modifications) for our numerical experiment. Consider a network of N firms that produce the same product but with different qualities. At time $t \in \{0, 1, \dots\}$, the product quality of i -th, $i \in \{1, 2, \dots, N\}$, firm is denoted as x_t^i which can assume value from the set $\mathcal{Q} \triangleq \{0, 1, \dots, Q-1\}$. Each firm has two options. It can either remain unresponsive or invest some money to improve the quality of its product. These two possibilities are symbolised as the elements of the action set, $\mathcal{U} \triangleq \{0, 1\}$ where 0 indicates unresponsiveness and 1 denotes investment. If at time t , the i -th firm chooses the action, $u_t^i \in \mathcal{U}$, then its state in the next time step changes according to the following transition law.

$$x_{t+1}^i = \begin{cases} x_t^i + \left[\chi \left(1 - \frac{\bar{\boldsymbol{\mu}}_t^N}{Q-1} \right) (Q-1 - x_t^i) \right] & \text{if } u_t^i = 1, \\ x_t^i & \text{otherwise} \end{cases}$$

where χ is a uniform random variable in $[0, 1]$, and $\bar{\boldsymbol{\mu}}_t^N$ is empirical average product quality defined as the mean of the distribution, $\boldsymbol{\mu}_t^N$ defined in (4). Hence, if $u_t^i = 0$ (unresponsiveness), then the product quality does not change. On the other hand, if the firm invests, i.e., $u_t^i = 1$, then its product quality increases. The increase in the product quality, however, is dependent on the average product quality, $\bar{\boldsymbol{\mu}}_t^N$, in the economy. If $\bar{\boldsymbol{\mu}}_t^N$ is high, then it is difficult to improve the product quality of any individual firm. The reward, and cost received by the i -th firm at time t are given as follows.

$$\begin{aligned} r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) &= \alpha_R x_t^i - \beta_R \bar{\boldsymbol{\mu}}_t^N - \lambda_R u_t^i \\ c(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) &= \lambda_C u_t^i \end{aligned}$$

The reward consists of three parts. The first part, $\alpha_R x_t^i$, is earned as revenue; the second part, $\beta_R \bar{\boldsymbol{\mu}}_t^N$ expresses dependence on the whole population; the third part, $\lambda_R u_t^i$, is due to the investment. The cost is a constant, λ_C , if money is invested, otherwise it is zero. The objective of this N -agent RL problem is to maximize the expected time-discounted sum of rewards while ensuring that the cumulative time-discounted cost is bounded above by a certain constant, ζ . Let π_∞^* be the optimal policy of its associated constraint mean-field control (CMFC) problem. In Fig. 1a, we demonstrate how the following error changes as a function of N .

$$\text{Error} = \left| \frac{V_N^R(\boldsymbol{x}_0^N, \pi_\infty^*) - V_\infty^R(\boldsymbol{\mu}_0, \pi_\infty^*)}{V_\infty^R(\boldsymbol{\mu}_0, \pi_\infty^*)} \right| \times 100\% \quad (40)$$

The value functions V_R^N and V_∞^R are defined in (6) and (12), respectively. The initial state distribution, $\boldsymbol{\mu}_0$, is taken to be a uniform distribution over \mathcal{Q} , and \boldsymbol{x}_0^N is obtained by taking N -independent samples from $\boldsymbol{\mu}_0$. The policy, π_∞^* , is obtained via Algorithm 1. Values of other relevant parameters are stated in Fig. 1. We observe that the error decreases as a function of N . Essentially, Fig. 1a shows that if N is large, then the N -agent cumulative average reward generated by π_∞^* is well approximated by the optimal mean-field value. In Fig. 1b, we exhibit that the N -agent and mean-field cost values generated by π_∞^* are close for large N , and both of them lie below the specified upper bound, ζ .

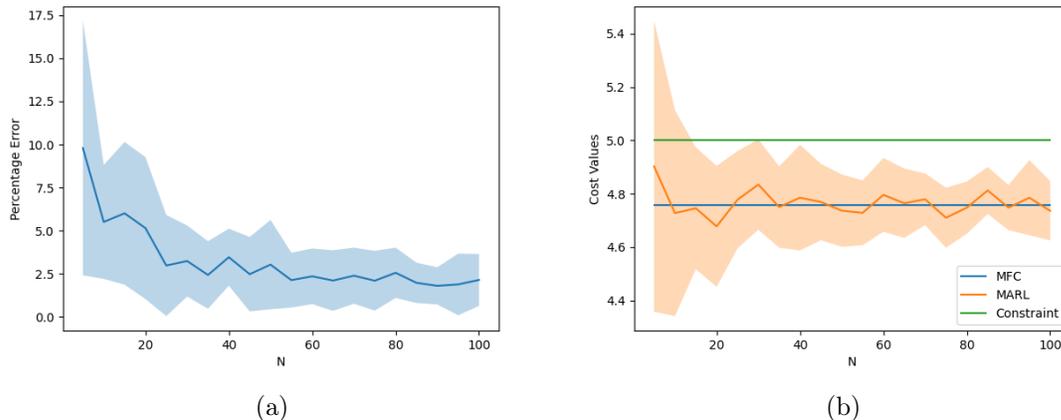


Figure 1: Fig. (1a) portrays the percentage error (defined by (40)) as a function of N . On the other hand, Fig. (1b) plots the N -agent (orange), and infinite-agent (blue) cost-values corresponding to the optimal mean-field policy as a function of N . It also shows that both of these values lie below the specified upper bound, ζ (green). The values of different system parameters are given as: $\alpha_R = 1$, $\beta_R = 0.5$, $\lambda_R = 0.5$, $\lambda_C = 1$, $\zeta = 5$, $\gamma = 0.9$, and $Q = 10$. The hyperparameters used in Algorithm 1 are chosen as follows: $\eta_1 = \eta_2 = \alpha = 10^{-3}$, $J = L = 10^2$. The bold lines, and the half-width of the shaded regions respectively denote the mean values, and the standard deviation values obtained over 25 random seeds. The experiments were performed on a 1.8 GHz Dual-Core Intel *i5* processor with 8 GB 1600 MHz DDR3 memory.

9. Conclusions

This paper shows that a constrained multi-agent reinforcement learning (CMARL) problem can be well-approximated via a constrained mean-field control (CMFC) problem with a suitably adjusted constraint bound. We have characterized the approximation error as a function of population size, and the sizes of state, and action spaces respectively. We also state an algorithm to solve the CMFC problem and analyze its sample complexity. One limitation of this study is that it only considers constraints in the form of time-discounted cumulative costs. Studying the same problem with other forms of constraints such as instantaneous cost, average cost, etc. could be a potential future direction.

Appendix A. Proof of Theorem 2

Using Lemma 3, we obtain,

$$V_N^C(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \leq V_\infty^C(\boldsymbol{\mu}_0, \boldsymbol{\pi}_\infty^*) + G_C \stackrel{(a)}{\leq} 0 \quad (41)$$

Inequality (a) follows from the fact that $\boldsymbol{\pi}_\infty^*$ is a feasible solution of (G-CMFC) with $\zeta = -G_C$. Due to (41), $\boldsymbol{\pi}_\infty^*$ can be treated as a feasible solution of (CMARL). Therefore,

$$\begin{aligned} & V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_N^*) - V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \\ &= V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_N^*) - V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_\infty^*) + V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_\infty^*) - V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \\ &\stackrel{(a)}{=} V_N^*(\mathbf{x}_0^N, 0) - V_\infty^*(\boldsymbol{\mu}_0, -G_C) + V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_\infty^*) - V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_\infty^*) \\ &\stackrel{(b)}{\leq} 2G_R + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1-\gamma} \right) \right] \end{aligned}$$

Relation (a) follows from the fact that $V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}_N^*) = V_N^*(\mathbf{x}_0^N, 0)$, and $V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi}_\infty^*) = V_\infty^*(\boldsymbol{\mu}_0, -G_C)$. Please refer to Table 1 for a detailed explanation. Inequality (b) follows from Lemma 1, and Lemma 3. This concludes the result.

Appendix B. Proof of Lemma 3

In order to establish Lemma 3, the following results are necessary. We use the notation π_t to indicate the t -th component of the policy sequence, $\boldsymbol{\pi}$. Moreover, the $\boldsymbol{\pi}$ -induced empirical N -agent state, and action distributions at time t are denoted as $\boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N$ respectively. Their counterparts for infinite agent system are $\boldsymbol{\mu}_t^\infty$, and $\boldsymbol{\nu}_t^\infty$. Following the notation of section 2, the ($\boldsymbol{\pi}$ -induced) joint state, and action at time t are denoted as $\mathbf{x}_t^N \triangleq \{x_t^i\}_{i \in \{1, \dots, N\}}$, $\mathbf{u}_t^N \triangleq \{u_t^i\}_{i \in \{1, \dots, N\}}$ respectively.

B.1 Lipschitz Continuity Lemmas

Lemma 6 *If $\nu^{\text{MF}}(\cdot, \cdot)$ is defined by (8), then the following holds $\forall t \in \{0, 1, \dots\}$.*

$$|\nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \leq (1 + L_Q) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \quad (42)$$

Lemma 7 *If $P^{\text{MF}}(\cdot, \cdot)$ is defined by (9), then the following holds $\forall t \in \{0, 1, \dots\}$.*

$$|P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - P^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \leq S_P |\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}|_1 \quad (43)$$

where $S_P \triangleq 1 + 2L_P + L_Q(1 + L_P)$.

Lemma 8 *If $r^{\text{MF}}(\cdot, \cdot)$ is defined by (10), then the following holds $\forall t \in \{0, 1, \dots\}$.*

$$|r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)| \leq S_R |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \quad (44)$$

where $S_R \triangleq M_R + 2L_R + L_Q(M_R + L_R)$.

Lemma 9 *If $c^{\text{MF}}(\cdot, \cdot)$ is defined by (11), then the following holds $\forall t \in \{0, 1, \dots\}$.*

$$|c^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - c^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)| \leq S_C |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \quad (45)$$

where $S_C \triangleq M_C + 2L_C + L_Q(M_C + L_C)$.

Lemma 6 – 9 shows that the state and action evolution functions, average reward, and cost functions of an infinite agent system are Lipschitz continuous with respect to the state distribution argument. Lemma 6 is an essential ingredient in the proof of Lemma 7, and 8. The proofs of Lemma 6, 7, and 8 are presented in Appendix D, E, and F respectively. The proof of Lemma 9 is identical to that of Lemma 8.

B.2 Large- N Approximation Lemmas

Lemma 10 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} |\boldsymbol{\nu}_t^N - \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)|_1 \leq \frac{1}{\sqrt{N}} \sqrt{|\mathcal{U}|}$$

Lemma 11 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} |\boldsymbol{\mu}_{t+1}^N - P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)|_1 \leq \frac{C_P}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right]$$

where $C_P \triangleq 2 + L_P$.

Lemma 12 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \leq \frac{M_R}{\sqrt{N}} + \frac{L_R}{\sqrt{N}} \sqrt{|\mathcal{U}|}$$

Lemma 13 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N c(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - c^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \leq \frac{M_C}{\sqrt{N}} + \frac{L_C}{\sqrt{N}} \sqrt{|\mathcal{U}|}$$

The proofs of Lemma 10, 11, and 12 are given in Appendix G, H, and I, respectively. The proof of Lemma 13 can be obtained by replacing r, M_R, L_R with c, M_C, L_C respectively in the proof of Lemma 12. Finally, invoking Lemma 7, and 11, we obtain the following.

Lemma 14 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty| \leq \frac{C_P}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] \left(\frac{S_P^t - 1}{S_P - 1} \right)$$

where S_P is defined in Lemma 7 while C_P is given in Lemma 11.

The proof of Lemma 14 is given in Appendix J.

B.3 Proof of the Lemma

We shall establish (18) only for $J = R$. The proof for $J = C$ is identical. Consider the following difference,

$$\begin{aligned}
 & |V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}) - V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi})| \\
 & \stackrel{(a)}{\leq} \sum_{t=0}^{\infty} \gamma^t \left| \frac{1}{N} \sum_{i=1}^N \mathbb{E} [r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N)] - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t) \right| \\
 & \leq \underbrace{\sum_{t=0}^{\infty} \gamma^t \mathbb{E} \left| \frac{1}{N} r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|}_{\triangleq J_1} + \underbrace{\sum_{t=0}^{\infty} \gamma^t \mathbb{E} |r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|}_{\triangleq J_2}
 \end{aligned}$$

Inequality (a) follows from the definition of the functions $V_N^R(\cdot, \cdot)$, $V_\infty^R(\cdot, \cdot)$ given in (6), and (12) respectively. The first term, J_1 can be bounded using Lemma 12 as follows.

$$J_1 \leq \left(\frac{1}{1-\gamma} \right) \left[\frac{M_R}{\sqrt{N}} + \frac{L_R}{\sqrt{N}} \sqrt{|\mathcal{U}|} \right]$$

The second term, J_2 , can be bounded as follows.

$$\begin{aligned}
 J_2 & \triangleq \sum_{t=0}^{\infty} \gamma^t \mathbb{E} |r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)| \\
 & \stackrel{(a)}{\leq} S_R \sum_{t=0}^{\infty} \gamma^t \mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty| \\
 & \stackrel{(b)}{\leq} \frac{1}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] \left(\frac{S_R C_P}{S_P - 1} \right) \left[\frac{1}{1-\gamma S_P} - \frac{1}{1-\gamma} \right]
 \end{aligned}$$

Inequality (a) follows from Lemma 8, whereas (b) is a consequence of Lemma 14.

Appendix C. Proof of Theorem 4

The proof of Theorem 4 hinges on the following Lemma.

Lemma 15 *Let \mathbf{x}_0^N be the initial joint state in an N -agent system, and $\boldsymbol{\mu}_0$ its empirical distribution. If Assumption 1 – 3 hold, then there exists $N_0 > 0$ such that $\forall N \geq N_0$, and $\forall \boldsymbol{\pi} \in \Pi_\infty$, the following inequalities hold whenever $\gamma S_P < 1$.*

$$|V_N^J(\mathbf{x}_0^N, \boldsymbol{\pi}) - V_\infty^J(\boldsymbol{\mu}_0, \boldsymbol{\pi})| \leq G_J^0 \quad (46)$$

where $J \in \{R, C\}$. The value functions V_N^R, V_N^C are given by (6), (7) respectively, and the terms G_R^0, G_C^0, S_P are defined in Theorem 4.

The role of Lemma 15 in establishing Theorem 4 is analogous to the role of Lemma 3 in proving Lemma 1. In this section, we shall primarily focus on proving Lemma 15. Once it is established, Theorem 4 can be proven following an argument similar to that is used in section 5.

C.1 Auxiliary Lemmas

To prove Lemma 15, the following results are necessary. The notations are the same as used in Appendix B.

Lemma 16 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \boldsymbol{\mu}_{t+1}^N - P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \leq \frac{2}{\sqrt{N}} \sqrt{|\mathcal{X}|} \quad (47)$$

Lemma 17 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \leq \frac{M_R}{\sqrt{N}}$$

Lemma 18 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N c(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - c^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \leq \frac{M_C}{\sqrt{N}}$$

Lemma 19 *The following inequality holds $\forall t \in \{0, 1, \dots\}$.*

$$\mathbb{E} \left| \boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty \right| \leq \frac{2}{\sqrt{N}} \sqrt{|\mathcal{X}|} \left(\frac{S_P^t - 1}{S_P - 1} \right)$$

The term S_P is defined in Lemma 7.

The proofs of Lemma 16, 17, 19 are relegated to Appendix *K*, *L*, and *M*, respectively. The proof of Lemma 18 can be obtained by replacing r with c , and M_R with M_C in the proof of Lemma 17.

C.2 Proof of Lemma 15

We use the same notations as in Appendix B.3. Consider the following difference,

$$\begin{aligned} & |V_N^R(\mathbf{x}_0^N, \boldsymbol{\pi}) - V_\infty^R(\boldsymbol{\mu}_0, \boldsymbol{\pi})| \\ & \stackrel{(a)}{\leq} \sum_{t=0}^{\infty} \gamma^t \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) \right| - \mathbb{E} [r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)] \\ & \leq \underbrace{\sum_{t=0}^{\infty} \gamma^t \mathbb{E} \left| \frac{1}{N} r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|}_{\triangleq J_1} + \underbrace{\sum_{t=0}^{\infty} \gamma^t \mathbb{E} |r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|}_{\triangleq J_2} \end{aligned}$$

Inequality (a) follows from the definition of the value functions $V_N^R(\cdot, \cdot)$, $V_\infty^C(\cdot, \cdot)$ given in (6), and (12) respectively. The first term, J_1 can be bounded using Lemma 12 as follows.

$$J_1 \leq \frac{M_R}{\sqrt{N}} \left(\frac{1}{1 - \gamma} \right)$$

The second term, J_2 , can be bounded as follows.

$$\begin{aligned} J_2 &\triangleq \sum_{t=0}^{\infty} \gamma^t \mathbb{E} |r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)| \\ &\stackrel{(a)}{\leq} S_R \sum_{t=0}^{\infty} \gamma^t \mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty| \stackrel{(b)}{\leq} \frac{1}{\sqrt{N}} \sqrt{|\mathcal{X}|} \left(\frac{2S_R}{S_P - 1} \right) \left[\frac{1}{1 - \gamma S_P} - \frac{1}{1 - \gamma} \right] \end{aligned}$$

Inequality (a) follows from Lemma 8, whereas (b) is a consequence of Lemma 19. This establishes (46) for $J = R$. The other case, $J = C$, can be proven similarly.

Appendix D. Proof of Lemma 6

Note the following inequalities.

$$\begin{aligned} &|\nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \\ &\stackrel{(a)}{=} \left| \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^N) \boldsymbol{\mu}_t^N(x) - \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^\infty) \boldsymbol{\mu}_t^\infty(x) \right|_1 \\ &= \sum_{u \in \mathcal{U}} \left| \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) - \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \right| \\ &\leq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} |\pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x)| \\ &\leq \sum_{x \in \mathcal{X}} |\boldsymbol{\mu}_t^N(x) - \boldsymbol{\mu}_t^\infty(x)| \underbrace{\sum_{u \in \mathcal{U}} \pi_t(x, \boldsymbol{\mu}_t^N)(u)}_{=1} + \sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x) \sum_{u \in \mathcal{U}} |\pi_t(x, \boldsymbol{\mu}_t^N)(u) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u)| \\ &\stackrel{(b)}{\leq} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 + L_Q |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty| \underbrace{\sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x)}_{=1} \stackrel{(c)}{=} (1 + L_Q) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \end{aligned}$$

Equality (a) follows from the definition of $\nu^{\text{MF}}(\cdot, \cdot)$ as given in (8). On the other hand, inequality (b) is a consequence of Assumption 2, and the fact that $\pi_t(x, \boldsymbol{\mu}_t^N)$ is probability distribution $\forall x \in \mathcal{X}$. Finally, (c) follows because $\boldsymbol{\mu}_t^\infty$ is a probability distribution. This concludes the result.

Appendix E. Proof of Lemma 7

Note that,

$$\begin{aligned} &|P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - P^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \\ &\stackrel{(a)}{=} \left| \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} P(x, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) \right. \\ &\quad \left. - P(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \right|_1 \leq J_1 + J_2 \end{aligned}$$

Equality (a) follows from the definition of $P^{\text{MF}}(\cdot, \cdot)$ as stated in (9). The first term can be bounded as follows.

$$\begin{aligned}
 J_1 &\triangleq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} \left| P(x, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) - P(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \right|_1 \times \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) \\
 &\stackrel{(a)}{\leq} L_P \left[|\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 + |\nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \right] \times \underbrace{\sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^N(x) \sum_{u \in \mathcal{U}} \pi_t(x, \boldsymbol{\mu}_t^N)(u)}_{=1} \\
 &\stackrel{(b)}{\leq} L_P(2 + L_Q) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1
 \end{aligned}$$

Inequality (a) is a result of Assumption 1(e) while (b) follows from Lemma 6, and the fact that $\pi_t(x, \boldsymbol{\mu}_t^N)$, $\boldsymbol{\mu}_t^N$ are probability distributions $\forall x \in \mathcal{X}$. The second term, J_2 can be bounded as follows.

$$\begin{aligned}
 J_2 &\triangleq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} \underbrace{|P(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t))|_1}_{=1} \times |\pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x)| \\
 &\leq \sum_{x \in \mathcal{X}} |\boldsymbol{\mu}_t^N(x) - \boldsymbol{\mu}_t^\infty(x)| \underbrace{\sum_{u \in \mathcal{U}} \pi_t(x, \boldsymbol{\mu}_t^N)(u)}_{=1} \\
 &\quad + \sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x) \sum_{u \in \mathcal{U}} |\pi_t(x, \boldsymbol{\mu}_t^N)(u) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u)| \\
 &\stackrel{(a)}{\leq} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 + L_P |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \underbrace{\sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x)}_{=1} \\
 &= (1 + L_P) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1
 \end{aligned}$$

Inequality (a) follows from the fact that $\pi_t(x, \boldsymbol{\mu}_t^N)$ is a probability distribution $\forall x \in \mathcal{X}$ while (b) uses Assumption 2, and the fact that $\boldsymbol{\mu}_t^\infty$ is a probability distribution. This concludes the result.

Appendix F. Proof of Lemma 8

Notice the following inequality,

$$\begin{aligned}
 &|r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - r^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)| \\
 &\stackrel{(a)}{=} \left| \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} r(x, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) \right. \\
 &\quad \left. - r(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x) \right| \leq J_1 + J_2
 \end{aligned}$$

Equality (a) follows from the definition of $r^{\text{MF}}(\cdot, \cdot)$ as depicted in (10). The first term can be bounded as follows.

$$\begin{aligned}
 J_1 &\triangleq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} \left| r(x, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) - r(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)) \right| \times \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) \\
 &\stackrel{(a)}{\leq} L_R \left[|\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 + |\nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) - \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t)|_1 \right] \times \underbrace{\sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^N(x) \sum_{u \in \mathcal{U}} \pi_t(x, \boldsymbol{\mu}_t^N)(u)}_{=1} \\
 &\stackrel{(b)}{\leq} L_R(2 + L_Q) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1
 \end{aligned}$$

Inequality (a) is a result of Assumption 1(c) while relation (b) follows from Lemma 6, and the fact that $\pi_t(x, \boldsymbol{\mu}_t^\infty), \boldsymbol{\mu}_t^N$ are probability distributions $\forall x \in \mathcal{X}$. The term, J_2 can be bounded as follows.

$$\begin{aligned}
 J_2 &\triangleq \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} |r(x, u, \boldsymbol{\mu}_t^\infty, \nu^{\text{MF}}(\boldsymbol{\mu}_t^\infty, \pi_t))| \times |\pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x)| \\
 &\stackrel{(a)}{\leq} M_R \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} |\pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u) \boldsymbol{\mu}_t^\infty(x)| \\
 &\leq M_R \sum_{x \in \mathcal{X}} |\boldsymbol{\mu}_t^N(x) - \boldsymbol{\mu}_t^\infty(x)| \underbrace{\sum_{u \in \mathcal{U}} \pi_t(x, \boldsymbol{\mu}_t^N)(u)}_{=1} \\
 &\quad + M_R \sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x) \sum_{u \in \mathcal{U}} |\pi_t(x, \boldsymbol{\mu}_t^N)(u) - \pi_t(x, \boldsymbol{\mu}_t^\infty)(u)| \\
 &\stackrel{(b)}{\leq} M_R |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 + M_R L_Q |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \underbrace{\sum_{x \in \mathcal{X}} \boldsymbol{\mu}_t^\infty(x)}_{=1} \stackrel{(c)}{=} M_R(1 + L_Q) |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1
 \end{aligned}$$

Inequality (a) follows from Assumption 1(a). On the other hand, (b) is a consequence of Assumption 2. Finally, (c) follows from the fact that $\boldsymbol{\mu}_t^\infty$ is a probability distribution. This concludes the result.

Appendix G. Proof of Lemma 10

The following Lemma is required to prove the result.

Lemma 20 *If $\forall m \in \{1, \dots, M\}$, $\{X_{mn}\}_{n \in \{1, \dots, N\}}$ are independent random variables that lie in $[0, 1]$, and satisfy $\sum_{m \in \{1, \dots, M\}} \mathbb{E}[X_{mn}] \leq 1$, $\forall n \in \{1, \dots, N\}$, then the following holds,*

$$\sum_{m=1}^M \mathbb{E} \left| \sum_{n=1}^N (X_{mn} - \mathbb{E}[X_{mn}]) \right| \leq \sqrt{MN} \quad (48)$$

Lemma 20 is adapted from Lemma 13 of (Mondal et al., 2022a).

Observe the following inequalities.

$$\begin{aligned}
 & \mathbb{E} \left| \boldsymbol{\nu}_t^N - \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \\
 &= \mathbb{E} \left[\mathbb{E} \left[\left| \boldsymbol{\nu}_t^N - \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \middle| \mathbf{x}_t^N \right] \right] \\
 &\stackrel{(a)}{=} \mathbb{E} \left[\mathbb{E} \left[\left| \boldsymbol{\nu}_t^N - \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^N) \boldsymbol{\mu}_t^N(x) \right|_1 \middle| \mathbf{x}_t^N \right] \right] \\
 &= \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{u \in \mathcal{U}} \boldsymbol{\nu}_t^N(u) - \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^N)(u) \boldsymbol{\mu}_t^N(x) \right| \middle| \mathbf{x}_t^N \right] \right] \\
 &\stackrel{(b)}{=} \mathbb{E} \left[\sum_{u \in \mathcal{U}} \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \delta(u_t^i = u) - \frac{1}{N} \sum_{x \in \mathcal{X}} \pi_t(x, \boldsymbol{\mu}_t^N)(u) \sum_{i=1}^N \delta(x_t^i = x) \right| \middle| \mathbf{x}_t^N \right] \right] \\
 &= \mathbb{E} \left[\sum_{u \in \mathcal{U}} \mathbb{E} \left[\left| \frac{1}{N} \sum_{i=1}^N \delta(u_t^i = u) - \frac{1}{N} \sum_{i=1}^N \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \middle| \mathbf{x}_t^N \right] \right] \stackrel{(c)}{\leq} \frac{1}{\sqrt{N}} \sqrt{|\mathcal{U}|}
 \end{aligned}$$

Equality (a) follows from the definition of $\nu^{\text{MF}}(\cdot, \cdot)$ as depicted in (8) while (b) is a consequence of the definitions of $\boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N$. Finally, (c) is an application of Lemma 20. Specifically, it uses the facts that, $\{u_t^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given \mathbf{x}_t^N , and

$$\mathbb{E} \left[\delta(u_t^i = u) \middle| \mathbf{x}_t^N \right] = \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u), \quad \sum_{u \in \mathcal{U}} \mathbb{E} \left[\delta(u_t^i = u) \middle| \mathbf{x}_t^N \right] = 1$$

$\forall i \in \{1, \dots, N\}, \forall u \in \mathcal{U}$. This concludes the lemma.

Appendix H. Proof of Lemma 11

Note that,

$$\begin{aligned}
 & \mathbb{E} \left| \boldsymbol{\mu}_{t+1}^N - P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \stackrel{(a)}{=} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N \delta(x_{t+1}^i = x) \right. \\
 & \quad \left. - \sum_{x' \in \mathcal{X}} \sum_{u \in \mathcal{U}} P(x', u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \pi_t(x', \boldsymbol{\mu}_t^N)(u) \frac{1}{N} \sum_{i=1}^N \delta(x_t^i = x') \right| \\
 &= \sum_{x \in \mathcal{X}} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N \delta(x_{t+1}^i = x) - \frac{1}{N} \sum_{i=1}^N \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \\
 &\leq J_1 + J_2 + J_3
 \end{aligned}$$

Equality (a) follows from the definition of $P^{\text{MF}}(\cdot, \cdot)$ as given in (9). The first term, J_1 , can be upper bounded as follows.

$$\begin{aligned}
 J_1 &\triangleq \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \sum_{i=1}^N \delta(x_{t+1}^i = x) - \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N)(x) \right| \\
 &= \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N \delta(x_{t+1}^i = x) - \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N)(x) \right| \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] \right] \stackrel{(a)}{\leq} \frac{1}{\sqrt{N}} \sqrt{|\mathcal{X}|}
 \end{aligned}$$

Inequality (a) can be derived using Lemma 20, and the facts that $\{x_{t+1}^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given $\{\mathbf{x}_t^N, \mathbf{u}_t^N\}$, and,

$$\mathbb{E} \left[\delta(x_{t+1}^i = x) \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] = P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N)(x), \quad \sum_{x \in \mathcal{X}} \mathbb{E} \left[\delta(x_{t+1}^i = x) \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] = 1$$

$\forall i \in \{1, \dots, N\}$, and $\forall x \in \mathcal{X}$. The second term can be upper bounded as follows.

$$\begin{aligned} J_2 &\triangleq \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N)(x) - \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \right| \\ &\leq \frac{1}{N} \sum_{i=1}^N \mathbb{E} \left| P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \right|_1 \\ &\stackrel{(a)}{\leq} L_P \mathbb{E} \left| \boldsymbol{\nu}_t^N - \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \stackrel{(b)}{\leq} \frac{L_P}{\sqrt{N}} \sqrt{|\mathcal{U}|} \end{aligned}$$

Inequality (a) follows from Assumption 3 while (b) follows from Lemma 10. Finally, the third term can be bounded as follows.

$$\begin{aligned} J_3 &\triangleq \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \right. \\ &\quad \left. - \sum_{i=1}^N \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \stackrel{(a)}{\leq} \frac{1}{\sqrt{N}} \sqrt{|\mathcal{X}|} \end{aligned}$$

Inequality (a) is a consequence of Lemma 20. Specifically, it uses the facts that, $\{u_t^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given \mathbf{x}_t^N , and

$$\begin{aligned} \mathbb{E} \left[P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \middle| \mathbf{x}_t^N \right] &= \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u), \\ \sum_{x \in \mathcal{X}} \mathbb{E} \left[P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))(x) \middle| \mathbf{x}_t^N \right] &= 1 \end{aligned}$$

$\forall i \in \{1, \dots, N\}$, and $\forall x \in \mathcal{X}$. This concludes the Lemma.

Appendix I. Proof of Lemma 12

Note that,

$$\begin{aligned} &\mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \\ &\stackrel{(a)}{=} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} r(x, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x, \boldsymbol{\mu}_t^N)(u) \frac{1}{N} \sum_{i=1}^N \delta(x_t^i = x) \right| \\ &= \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - \frac{1}{N} \sum_{i=1}^N \sum_{u \in \mathcal{U}} r(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \\ &\leq J_1 + J_2 \end{aligned}$$

Equation (a) is a result of the definition of $r^{\text{MF}}(\cdot, \cdot)$ as given in (10). The first term, J_1 , can be bounded as follows.

$$\begin{aligned} J_1 &\triangleq \frac{1}{N} \mathbb{E} \left| \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \right| \\ &\leq \frac{1}{N} \mathbb{E} \sum_{i=1}^N |r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \boldsymbol{\nu}_t^N) - r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))| \\ &\stackrel{(a)}{\leq} L_R \mathbb{E} |\boldsymbol{\nu}_t^N - \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)|_1 \stackrel{(b)}{\leq} \frac{L_R}{\sqrt{N}} \sqrt{|\mathcal{U}|} \end{aligned}$$

Inequality (a) follows from Assumption 1, and (b) is a consequence of Lemma 10. The second term, J_2 , can be bounded as follows.

$$\begin{aligned} J_2 &\triangleq \frac{1}{N} \mathbb{E} \left| \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} r(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \\ &= \frac{1}{N} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} r(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \middle| \boldsymbol{x}_t^N \right] \right] \\ &= \frac{M_R}{N} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} r_0(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \middle| \boldsymbol{x}_t^N \right] \right] \\ &\stackrel{(a)}{\leq} \frac{M_R}{\sqrt{N}} \end{aligned}$$

where $r_0(\cdot, \cdot, \cdot, \cdot) \triangleq r(\cdot, \cdot, \cdot, \cdot)/M_R$. Inequality (a) follows from Lemma 20. Specifically, it uses the fact that $\{u_t^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given \boldsymbol{x}_t , and

$$\begin{aligned} |r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t))| &\leq 1, \\ \mathbb{E} \left[r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \middle| \boldsymbol{x}_t^N \right] &= \sum_{u \in \mathcal{U}} r_0(x_t^i, u, \boldsymbol{\mu}_t^N, \nu^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t)) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \end{aligned}$$

$\forall i \in \{1, \dots, N\}, \forall u \in \mathcal{U}$.

Appendix J. Proof of Lemma 14

Observe that,

$$\begin{aligned} \mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 &\leq \mathbb{E} |\boldsymbol{\mu}_t^N - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1})|_1 + \mathbb{E} |P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - \boldsymbol{\mu}_t^\infty|_1 \\ &\stackrel{(a)}{\leq} \frac{C_P}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] + \mathbb{E} |P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^\infty, \pi_{t-1})|_1 \end{aligned}$$

Inequality (a) follows from Lemma 11, and relation (9). Using Lemma 7, we get

$$|P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^\infty, \pi_{t-1})|_1 \stackrel{(a)}{\leq} S_P |\boldsymbol{\mu}_{t-1}^N - \boldsymbol{\mu}_{t-1}^\infty|_1$$

Combining, we get,

$$\mathbb{E}|\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \leq \frac{C_P}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] + S_P \mathbb{E}|\boldsymbol{\mu}_{t-1}^N - \boldsymbol{\mu}_{t-1}^\infty|_1 \quad (49)$$

Recursively applying the above inequality, we finally obtain,

$$\mathbb{E}|\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \leq \frac{C_P}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] \left(\frac{S_P^t - 1}{S_P - 1} \right)$$

Appendix K. Proof of Lemma 16

Note that,

$$\begin{aligned} & \mathbb{E} \left| \boldsymbol{\mu}_{t+1}^N - P^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right|_1 \\ & \stackrel{(a)}{=} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N \delta(x_{t+1}^i = x) - \sum_{x' \in \mathcal{X}} \sum_{u \in \mathcal{U}} P(x', u, \boldsymbol{\mu}_t^N)(x) \pi_t(x', \boldsymbol{\mu}_t^N)(u) \frac{1}{N} \sum_{i=1}^N \delta(x_t^i = x') \right| \\ & = \sum_{x \in \mathcal{X}} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N \delta(x_{t+1}^i = x) - \frac{1}{N} \sum_{i=1}^N \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N)(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \leq J_1 + J_2 \end{aligned}$$

Equality (a) follows from the definition of $P^{\text{MF}}(\cdot, \cdot)$ as depicted in (9). The first term, J_1 , can be upper bounded as follows.

$$\begin{aligned} J_1 & \triangleq \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \sum_{i=1}^N \delta(x_{t+1}^i = x) - \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x) \right| \\ & = \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N \delta(x_{t+1}^i = x) - \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x) \right| \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] \right] \stackrel{(a)}{\leq} \frac{1}{\sqrt{N}} \sqrt{|\mathcal{X}|} \end{aligned}$$

Inequality (a) can be derived using Lemma 20, and the facts that $\{x_{t+1}^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given $\{\mathbf{x}_t^N, \mathbf{u}_t^N\}$, and,

$$\mathbb{E} \left[\delta(x_{t+1}^i = x) \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] = P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x), \quad \sum_{x \in \mathcal{X}} \mathbb{E} \left[\delta(x_{t+1}^i = x) \middle| \mathbf{x}_t^N, \mathbf{u}_t^N \right] = 1$$

$\forall i \in \{1, \dots, N\}$, and $\forall x \in \mathcal{X}$. The second term can be bounded as follows.

$$J_2 \triangleq \frac{1}{N} \sum_{x \in \mathcal{X}} \mathbb{E} \left| \sum_{i=1}^N P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N)(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \stackrel{(a)}{\leq} \frac{\sqrt{|\mathcal{X}|}}{\sqrt{N}}$$

Inequality (a) is a consequence of Lemma 20. Specifically, it uses the facts that, $\{u_t^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given \mathbf{x}_t^N , and

$$\begin{aligned} \mathbb{E} \left[P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x) \middle| \mathbf{x}_t^N \right] & = \sum_{u \in \mathcal{U}} P(x_t^i, u, \boldsymbol{\mu}_t^N)(x) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u), \\ \sum_{x \in \mathcal{X}} \mathbb{E} \left[P(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)(x) \middle| \mathbf{x}_t^N \right] & = 1 \end{aligned}$$

$\forall i \in \{1, \dots, N\}$, and $\forall x \in \mathcal{X}$. This concludes the Lemma.

Appendix L. Proof of Lemma 17

Note that,

$$\begin{aligned}
 & \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - r^{\text{MF}}(\boldsymbol{\mu}_t^N, \pi_t) \right| \\
 & \stackrel{(a)}{=} \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - \sum_{x \in \mathcal{X}} \sum_{u \in \mathcal{U}} r(x, u, \boldsymbol{\mu}_t^N) \pi_t(x, \boldsymbol{\mu}_t^N)(u) \frac{1}{N} \sum_{i=1}^N \delta(x_t^i = x) \right| \\
 & = \mathbb{E} \left| \frac{1}{N} \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - \frac{1}{N} \sum_{i=1}^N \sum_{u \in \mathcal{U}} r(x_t^i, u, \boldsymbol{\mu}_t^N) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \\
 & = \frac{1}{N} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N r(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} r(x_t^i, u, \boldsymbol{\mu}_t^N) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \middle| \boldsymbol{x}_t^N \right] \right] \\
 & = \frac{M_R}{N} \mathbb{E} \left[\mathbb{E} \left[\left| \sum_{i=1}^N r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) - \sum_{i=1}^N \sum_{u \in \mathcal{U}} r_0(x_t^i, u, \boldsymbol{\mu}_t^N) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u) \right| \middle| \boldsymbol{x}_t^N \right] \right] \stackrel{(a)}{\leq} \frac{M_R}{\sqrt{N}}
 \end{aligned}$$

where $r_0(\cdot, \cdot, \cdot, \cdot) \triangleq r(\cdot, \cdot, \cdot, \cdot)/M_R$. Inequality (a) follows from Lemma 20. Specifically, it uses the fact that $\{u_t^i\}_{i \in \{1, \dots, N\}}$ are conditionally independent given \boldsymbol{x}_t^N , and

$$|r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N)| \leq 1, \quad \mathbb{E} \left[r_0(x_t^i, u_t^i, \boldsymbol{\mu}_t^N) \middle| \boldsymbol{x}_t^N \right] = \sum_{u \in \mathcal{U}} r_0(x_t^i, u, \boldsymbol{\mu}_t^N) \pi_t(x_t^i, \boldsymbol{\mu}_t^N)(u)$$

$\forall i \in \{1, \dots, N\}, \forall u \in \mathcal{U}$.

Appendix M. Proof of Lemma 19

Observe that,

$$\begin{aligned}
 \mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 & \leq \mathbb{E} |\boldsymbol{\mu}_t^N - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1})|_1 + \mathbb{E} |P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - \boldsymbol{\mu}_t^\infty|_1 \\
 & \stackrel{(a)}{\leq} \frac{2}{\sqrt{N}} \sqrt{|\mathcal{X}|} + \mathbb{E} |P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^\infty, \pi_{t-1})|_1
 \end{aligned}$$

Inequality (a) follows from Lemma 11, and relation (9). Using Lemma 7, we get

$$|P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^N, \pi_{t-1}) - P^{\text{MF}}(\boldsymbol{\mu}_{t-1}^\infty, \pi_{t-1})|_1 \stackrel{(a)}{\leq} S_P |\boldsymbol{\mu}_{t-1}^N - \boldsymbol{\mu}_{t-1}^\infty|_1$$

Combining, we get,

$$\mathbb{E} |\boldsymbol{\mu}_t^N - \boldsymbol{\mu}_t^\infty|_1 \leq \frac{2}{\sqrt{N}} \sqrt{|\mathcal{X}|} + S_P \mathbb{E} |\boldsymbol{\mu}_{t-1}^N - \boldsymbol{\mu}_{t-1}^\infty|_1 \leq \frac{2}{\sqrt{N}} \sqrt{|\mathcal{X}|} \left(\frac{S_P^t - 1}{S_P - 1} \right) \quad (50)$$

the last inequality is obtained via recursion.

Appendix N. Proof of Theorem 5

Fix J, L . Following Theorem 3 of (Ding et al., 2020), we can write,

$$\left| V_\infty^*(\boldsymbol{\mu}_0, -2G_C) - \frac{1}{J} \sum_{j=1}^J V_\infty^R(\boldsymbol{\mu}_0, \pi_{\Phi_j}) \right| \leq C_1 \frac{1}{\sqrt{J}} + C_2 \left(\sqrt{\epsilon_{\text{bias}}} + \frac{C_3}{\sqrt{L}} \right) \quad (51)$$

for some constants C_1, C_2, C_3 . On the other hand, Lemma 1 suggests,

$$|V_N^*(\mathbf{x}_0^N, 0) - V_\infty^*(\boldsymbol{\mu}_0, -G_C)| \leq G_R + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1-\gamma} \right) \right] \quad (52)$$

Recall the definition of ζ_1 from Assumption 5. If N is sufficiently large, then $2G_C < \zeta_1$. Using the concavity of $V_\infty^*(\boldsymbol{\mu}_0, \cdot)$ (Paternain et al., 2019), we obtain:

$$V_\infty^*(\boldsymbol{\mu}_0, -2G_C) \geq \left(\frac{\zeta_1 - 2G_C}{\zeta_1 - G_C} \right) V_\infty^*(\boldsymbol{\mu}_0, -G_C) + \left(\frac{G_C}{\zeta_1 - G_C} \right) V_\infty^*(\boldsymbol{\mu}_0, -\zeta_1),$$

$$\begin{aligned} \text{Equivalently, } V_\infty^*(\boldsymbol{\mu}_0, -G_C) - V_\infty^*(\boldsymbol{\mu}_0, -2G_C) &\leq \left(\frac{G_C}{\zeta_1 - G_C} \right) [V_\infty^*(\boldsymbol{\mu}_0, -G_C) - V_\infty^*(\boldsymbol{\mu}_0, -\zeta_1)] \\ &\leq \left(\frac{2G_C}{\zeta_1 - G_C} \right) \left(\frac{M_R}{1-\gamma} \right) \leq \left(\frac{2G_C}{\zeta_1} \right) \left(\frac{M_R}{1-\gamma} \right) \end{aligned}$$

Combining, we obtain,

$$\begin{aligned} \left| V_N^*(\mathbf{x}_0^N, 0) - \frac{1}{J} \sum_{j=1}^J V_\infty^R(\boldsymbol{\mu}_0, \pi_{\Phi_j}) \right| &\leq C_1 \frac{1}{\sqrt{J}} + C_2 \left(\sqrt{\epsilon_{\text{bias}}} + \frac{C_3}{\sqrt{L}} \right) + G_R \\ &\quad + G_C \left[\frac{4}{\zeta_0} \left(\frac{M_R}{1-\gamma} \right) \right] + \left(\frac{2G_C}{\zeta_1} \right) \left(\frac{M_R}{1-\gamma} \right) \end{aligned} \quad (53)$$

From Theorem 3 of (Ding et al., 2020), we also obtain the following.

$$\frac{1}{J} \sum_{j=1}^J V_\infty^C(\boldsymbol{\mu}_0, \pi_{\Phi_j}) + 2G_C \leq \frac{C_3}{J^{\frac{1}{4}}} + \frac{C_4}{J^{\frac{1}{4}}} \left(\epsilon_{\text{bias}}^{\frac{1}{4}} + \frac{C_5}{L^{\frac{1}{4}}} \right) \quad (54)$$

for some constants C_4, C_5 . Using Lemma 3, we obtain,

$$\frac{1}{J} \sum_{j=1}^J V_N^C(\mathbf{x}_0^N, \pi_{\Phi_j}) \leq \frac{1}{J} \sum_{j=1}^J V_\infty^C(\boldsymbol{\mu}_0, \pi_{\Phi_j}) + G_C \leq \frac{C_3}{J^{\frac{1}{4}}} + \frac{C_4}{J^{\frac{1}{4}}} \left(\epsilon_{\text{bias}}^{\frac{1}{4}} + \frac{C_5}{L^{\frac{1}{4}}} \right) - G_C \quad (55)$$

Note that, $G_R, G_C = \mathcal{O}(e)$ where $e \triangleq \left(\frac{1}{\sqrt{N}} \left[\sqrt{|\mathcal{X}|} + \sqrt{|\mathcal{U}|} \right] \right)$. To make the RHS of (55) negative, we must have $J = \mathcal{O}(G_C^{-4}) = \mathcal{O}(e^{-4})$. Moreover, if we choose $L = \mathcal{O}(e^{-2})$, then the RHS of (51) becomes $\mathcal{O}(e + \sqrt{\epsilon_{\text{bias}}})$. Hence, the desired sample complexity is $\mathcal{O}(e^{-6})$. As discussed earlier, this sample complexity might be improved to $\tilde{\mathcal{O}}(e^{-2})$ if the recently proposed algorithm of Mondal and Aggarwal (2024b) is used.

Appendix O. Sampling Procedure

Algorithm 2 Sampling Algorithm

Input: $\boldsymbol{\mu}_0, \pi_{\Phi_j}, P, r$

- 1: Sample $x_0 \sim \boldsymbol{\mu}_0$.
 - 2: Sample $u_0 \sim \pi_{\Phi_j}(x_0, \boldsymbol{\mu}_0)$
 - 3: $\boldsymbol{\nu}_0 \leftarrow \nu^{\text{MF}}(\boldsymbol{\mu}_0, \pi_{\Phi_j})$
 - 4: $t \leftarrow 0$
 - 5: FLAG \leftarrow FALSE
 - 6: **while** FLAG is FALSE **do**
 - 7: FLAG \leftarrow TRUE with probability $1 - \gamma$.
 - 8: Execute Update
 - 9: **end while**
 - 10: $T \leftarrow t$
 - 11: Accept $(x_T, \boldsymbol{\mu}_T, u_T)$ as a sample.
 - 12: $\hat{V}_{\Phi_j}^R \leftarrow 0, \hat{Q}_{\Phi_j}^R \leftarrow 0, \hat{V}_{\Phi_j}^C \leftarrow 0, \hat{Q}_{\Phi_j}^C \leftarrow 0$
 - 13: FLAG \leftarrow FALSE
 - 14: SumRewards $\leftarrow 0$
 - 15: SumCosts $\leftarrow 0$
 - 16: **while** FLAG is FALSE **do**
 - 17: FLAG \leftarrow TRUE with probability $1 - \gamma$.
 - 18: Execute Update
 - 19: SumRewards \leftarrow SumRewards + $r(x_t, u_t, \boldsymbol{\mu}_t, \boldsymbol{\nu}_t)$
 - 20: SumCosts \leftarrow SumCosts + $c(x_t, u_t, \boldsymbol{\mu}_t, \boldsymbol{\nu}_t)$
 - 21: **end while**
 - 22: With probability $\frac{1}{2}$, $\hat{V}_{\Phi_j}^R \leftarrow$ SumRewards. Otherwise $\hat{Q}_{\Phi_j}^R \leftarrow$ SumRewards.
 - 23: With probability $\frac{1}{2}$, $\hat{V}_{\Phi_j}^C \leftarrow$ SumCosts. Otherwise $\hat{Q}_{\Phi_j}^C \leftarrow$ SumCosts.
 - 24: $\hat{A}_{\Phi_j}^R(x_T, \boldsymbol{\mu}_T, u_T) \leftarrow 2(\hat{Q}_{\Phi_j}^R - \hat{V}_{\Phi_j}^R)$.
 - 25: $\hat{A}_{\Phi_j}^C(x_T, \boldsymbol{\mu}_T, u_T) \leftarrow 2(\hat{Q}_{\Phi_j}^C - \hat{V}_{\Phi_j}^C)$.
 - 26: $\hat{A}_{\Phi_j}^{\lambda_j}(x_T, \boldsymbol{\mu}_T, u_T) \leftarrow \hat{A}_{\Phi_j}^R(x_T, \boldsymbol{\mu}_T, u_T) - \lambda_j \hat{A}_{\Phi_j}^C(x_T, \boldsymbol{\mu}_T, u_T)$
- Output:** $(x_T, \boldsymbol{\mu}_T, u_T)$, $\hat{A}_{\Phi_j}^{\lambda_j}(x_T, \boldsymbol{\mu}_T, u_T)$, and $\hat{V}_{\Phi_j}^C$

Procedure Update:

- 1: $x_{t+1} \sim P(x_t, u_t, \boldsymbol{\mu}_t, \boldsymbol{\nu}_t)$.
- 2: $\boldsymbol{\mu}_{t+1} \leftarrow P^{\text{MF}}(\boldsymbol{\mu}_t, \pi_{\Phi_j})$
- 3: $u_{t+1} \sim \pi_{\Phi_j}(x_{t+1}, \boldsymbol{\mu}_{t+1})$
- 4: $\boldsymbol{\nu}_{t+1} \leftarrow \nu^{\text{MF}}(\boldsymbol{\mu}_{t+1}, \pi_{\Phi_j})$
- 5: $t \leftarrow t + 1$

EndProcedure

References

- Joshua Achiam, David Held, Aviv Tamar, and Pieter Abbeel. Constrained policy optimization. In *International conference on machine learning*, pages 22–31. PMLR, 2017.
- Abubakr O Al-Abbasi, Arnob Ghosh, and Vaneet Aggarwal. Deeppool: Distributed model-free algorithm for ride-sharing using deep reinforcement learning. *IEEE Transactions on Intelligent Transportation Systems*, 20(12):4714–4727, 2019.
- Andrea Angiuli, Jean-Pierre Fouque, and Mathieu Laurière. Unified reinforcement q-learning for mean field game and control problems. *Mathematics of Control, Signals, and Systems*, pages 1–55, 2022.
- Qinbo Bai, Amrit Singh Bedi, and Vaneet Aggarwal. Achieving zero constraint violation for constrained reinforcement learning via conservative natural policy gradient primal-dual algorithm. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 6737–6744, 2023.
- Shalabh Bhatnagar and K Lakshmanan. An online actor-critic algorithm with function approximation for constrained markov decision processes. *Journal of Optimization Theory and Applications*, 153(3):688–708, 2012.
- Chiara Buratti, Andrea Conti, Davide Dardari, and Roberto Verdone. An overview on wireless sensor networks technology and evolution. *Sensors*, 9(9):6869–6896, 2009.
- René Carmona, François Delarue, et al. *Probabilistic Theory of Mean Field Games with Applications I-II*. Springer, 2018.
- René Carmona, Mathieu Laurière, and Zongjun Tan. Linear-quadratic mean-field reinforcement learning: convergence of policy gradient methods. *arXiv preprint arXiv:1910.04295*, 2019.
- Yinlam Chow, Mohammad Ghavamzadeh, Lucas Janson, and Marco Pavone. Risk-constrained reinforcement learning with percentile risk criteria. *The Journal of Machine Learning Research*, 18(1):6070–6120, 2017.
- Dongsheng Ding, Kaiqing Zhang, Tamer Basar, and Mihailo Jovanovic. Natural policy gradient primal-dual method for constrained markov decision processes. *Advances in Neural Information Processing Systems*, 33:8378–8390, 2020.
- Dmitri A Dolgov and Edmund H Durfee. Stationary deterministic policies for constrained mdps with multiple rewards, costs, and discount factors. In *IJCAI*, volume 19, pages 1326–1331. Citeseer, 2005.
- Nan Geng, Tian Lan, Vaneet Aggarwal, Yuan Yang, and Mingwei Xu. A multi-agent reinforcement learning perspective on distributed traffic engineering. In *2020 IEEE 28th International Conference on Network Protocols (ICNP)*, pages 1–11. IEEE, 2020.
- Haotian Gu, Xin Guo, Xiaoli Wei, and Renyuan Xu. Mean-field controls with Q-learning for cooperative MARL: convergence and complexity analysis. *SIAM Journal on Mathematics of Data Science*, 3(4):1168–1196, 2021.

- Wonjun Lee, Siting Liu, Hamidou Tembine, Wuchen Li, and Stanley Osher. Controlling propagation of epidemics via mean-field control. *SIAM Journal on Applied Mathematics*, 81(1):190–207, 2021.
- Shihui Li, Yi Wu, Xinyue Cui, Honghua Dong, Fei Fang, and Stuart Russell. Robust multi-agent reinforcement learning via minimax deep deterministic policy gradient. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 4213–4220, 2019.
- Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *nature*, 518(7540):529–533, 2015.
- Washim Uddin Mondal and Vaneet Aggarwal. Last-iterate convergence of general parameterized policies in constrained mdps. *arXiv preprint arXiv:2408.11513*, 2024a.
- Washim Uddin Mondal and Vaneet Aggarwal. Sample-efficient constrained reinforcement learning with general parameterization. *arXiv preprint arXiv:2405.10624*, 2024b.
- Washim Uddin Mondal, Mridul Agarwal, Vaneet Aggarwal, and Satish V. Ukkusuri. On the approximation of cooperative heterogeneous multi-agent reinforcement learning (MARL) using mean field control (MFC). *Journal of Machine Learning Research*, 23(129):1–46, 2022a.
- Washim Uddin Mondal, Vaneet Aggarwal, and Satish Ukkusuri. On the near-optimality of local policies in large cooperative multi-agent reinforcement learning. *Transactions on Machine Learning Research*, 2022b. ISSN 2835-8856. URL <https://openreview.net/forum?id=t5HkgbxZp1>.
- Washim Uddin Mondal, Vaneet Aggarwal, and Satish V Ukkusuri. Can mean field control (mfc) approximate cooperative multi agent reinforcement learning (marl) with non-uniform interaction? In *Uncertainty in Artificial Intelligence*, pages 1371–1380. PMLR, 2022c.
- Barna Pásztor, Andreas Krause, and Ilija Bogunovic. Efficient model-based multi-agent mean-field reinforcement learning. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=gvcDSDYUZx>.
- Santiago Paternain, Luiz Chamon, Miguel Calvo-Fullana, and Alejandro Ribeiro. Constrained reinforcement learning has zero duality gap. *Advances in Neural Information Processing Systems*, 32, 2019.
- Santiago Paternain, Miguel Calvo-Fullana, Luiz FO Chamon, and Alejandro Ribeiro. Safe policies for reinforcement learning via primal-dual methods. *IEEE Transactions on Automatic Control*, 2022.
- Tabish Rashid, Mikayel Samvelyan, Christian Schroeder, Gregory Farquhar, Jakob Foerster, and Shimon Whiteson. Qmix: Monotonic value function factorisation for deep multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 4295–4304. PMLR, 2018.

- Tabish Rashid, Gregory Farquhar, Bei Peng, and Shimon Whiteson. Weighted qmix: Expanding monotonic value function factorisation for deep multi-agent reinforcement learning. *Advances in neural information processing systems*, 33:10199–10210, 2020.
- Gavin A Rummery and Mahesan Niranjana. *On-line Q-learning using connectionist systems*, volume 37. Citeseer, 1994.
- Kyunghwan Son, Daewoo Kim, Wan Ju Kang, David Earl Hostallero, and Yung Yi. Qtran: Learning to factorize with transformation for cooperative multi-agent reinforcement learning. In *International Conference on Machine Learning*, pages 5887–5896. PMLR, 2019.
- Jayakumar Subramanian and Aditya Mahajan. Reinforcement learning in stationary mean-field games. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, pages 251–259, 2019.
- Peter Sunehag, Guy Lever, Audrunas Gruslys, Wojciech Marian Czarnecki, Vinicius Zambaldi, Max Jaderberg, Marc Lanctot, Nicolas Sonnerat, Joel Z Leibo, Karl Tuyls, et al. Value-decomposition networks for cooperative multi-agent learning based on team reward. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems*, pages 2085–2087, 2018.
- Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.
- Xiaoqiang Wang, Liangjun Ke, Zhimin Qiao, and Xinghua Chai. Large-scale traffic signal control using a novel multiagent reinforcement learning. *IEEE transactions on cybernetics*, 51(1):174–187, 2020.
- Zhe Wang, Vaneet Aggarwal, and Xiaodong Wang. Iterative dynamic water-filling for fading multiple-access channels with energy harvesting. *IEEE Journal on Selected Areas in Communications*, 33(3):382–395, 2015.
- Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine learning*, 8(3):279–292, 1992.
- Nicholas J Watkins, Cameron Nowzari, Victor M Preciado, and George J Pappas. Optimal resource allocation for competitive spreading processes on bilayer networks. *IEEE Transactions on Control of Network Systems*, 5(1):298–307, 2016.
- Yu Xiang, Tian Lan, Vaneet Aggarwal, and Yih Farn R Chen. Joint latency and cost optimization for erasure coded data center storage. *ACM SIGMETRICS Performance Evaluation Review*, 42(2):3–14, 2014.
- Hanhan Zhou, Tian Lan, and Vaneet Aggarwal. Pac: Assisted value factorization with counterfactual predictions in multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 35:15757–15769, 2022.