



Mean-Field Games Among Teams

Jayakumar Subramanian, Akshat Kumar and Aditya Mahajan

EasyChair preprints are intended for rapid dissemination of research results and are integrated with the rest of EasyChair.

October 25, 2023

Mean-field games among teams

Jayakumar Subramanian · Akshat
Kumar · Aditya Mahajan

Received: date / Accepted: date

Abstract In this paper, we present a model of a game among teams. Each team consists of a homogeneous population of agents. Agents within a team are cooperative while the teams compete with other teams. The dynamics and the costs are coupled through the empirical distribution (or the mean field) of the state of agents in each team. This mean-field is assumed to be observed by all agents. Agents have asymmetric information (also called a non-classical information structure). We propose a mean-field based refinement of the Team-Nash equilibrium of the game, which we call mean-field Markov perfect equilibrium (MF-MPE). We identify a dynamic programming decomposition to characterize MF-MPE. We then consider the case where each team has a large number of players and present a mean-field approximation which approximates the game among large-population teams as a game among infinite-population teams. We show that MF-MPE of the game among teams of infinite population is easier to compute and is an ε -approximate MF-MPE of the game among teams of finite population.

Keywords Mean-field games among teams, Team-Nash equilibrium, Markov perfect equilibrium, large population games among teams

The work of AM was supported by Natural Sciences and Engineering Research Council of Canada, Discovery Grant RGPIN-2021-03511.

J. Subramanian
Media and Data Science Research Lab, Digital Experience Cloud, Adobe Inc.,
Noida, Uttar Pradesh, India, E-mail: jasubram@adobe.com

A. Kumar
School of Computing and Information Systems at the Singapore Management University,
Singapore, E-mail: akshatkumar@smu.edu.sg

A. Mahajan
Department of Electrical and Computer Engineering,
McGill University, Montreal, Canada, E-mail: aditya.mahajan@mcgill.ca

1 Introduction

Traditionally, agents in a multi-agent system are modeled either as cooperative agents who optimize a common system-wide objective or as strategic agents who optimize individual objectives. This difference in the behavior of agents leads to different conceptual difficulties and different solution concepts. As a result, the two settings are studied as separate sub-disciplines of decision theory: models with cooperative agents are studied under the heading of team theory (Marschak and Radner 1972) and models with strategic agents are studied under the heading of game theory (von Neumann and Morgenstern 1944). For the most part, these two research subdisciplines have evolved independently.

However, in recent years, many applications have emerged which may be considered as games among teams. Examples include: multiple demand aggregators competing in the same electricity markets, multiple ride-sharing companies competing in the same city, multiple firms competing in the same industry with different levels of competitive advantages (Weintraub et al. 2008), and the DARPA Spectrum Sharing Challenge, where teams of multiple radio units compete with other such teams in the same geographic area (Tilghman 2019).

In such applications, teams of agents compete with other teams of agents. These models are different from traditional team theory models because agents belonging to different teams have separate objectives and are, therefore, not cooperative. These models are also different from traditional game theory models because agents belonging to the same team can enter into pregame agreements; therefore, the notion of equilibrium in games among teams must account for the possibility of simultaneous and coordinated deviations by all agents belonging to the same team. Such an equilibrium is called a *Team-Nash equilibrium* (Tang et al. 2022).

Games among teams are also different from cooperative games (Shapley 1953). In a game among teams, the team structure (i.e, which agent belongs to which team) is pre-specified and unlike cooperative game theory, the process of team formation and how to distribute rewards among members of the team is not investigated.

	L	R		L	R
T	3, 3, 1	0, 0, 0	T	0, 0, 0	1, 1, 5
B	0, 0, 0	1, 1, 3	B	2, 2, 2	0, 0, 0
	I			II	

Fig. 1 A game between the team of P_1 (who chooses the row) and P_2 (who chooses the column) versus P_3 (who chooses the matrix I or II).

To illustrate the difference between Nash equilibrium (NE) and Team-Nash equilibrium (TNE), we consider a static game among two teams shown in Fig. 1. If we view the above as multiplayer game with three players, then the game has four NE in pure strategies: $\{(T, L, I), (B, R, I), (T, R, II), (B, L, II)\}$.

Of these, only $\{(T, L, I), (B, R, II)\}$ are TNE. (B, R, I) is not TNE because the team of P_1 and P_2 can deviate to (T, L) and obtain a higher payoff. Thus, the key difference between NE and TNE is that in TNE players belonging to the same team can deviate together.

When agents in a team have symmetric information (as is the case in the example above), the game among teams can be reduced to a regular game by considering a single player with vector-valued actions. However, such a reduction is not possible when players belonging to the same team have *asymmetric information*. The situation is even more complicated for multi-stage (or dynamic) games, where games among teams inherit all the conceptual challenges of dynamic games with asymmetric information.

In many of the motivating applications described above, each team has a large number of agents. So, we investigate games among teams where each team has a large number of agents and call such systems *mean-field games among teams* (or MFGT for short). In the last decade, various *mean-field approximations* of teams and games with a large number of agents have been proposed in the literature. The general flavor of the results are as follows. For games, it is shown that an appropriate refinement of a Nash equilibrium of a game with a large number of players can be approximated by an equilibrium solution of a game with an infinite number of players [Huang et al. \(2007, 2012\)](#); [Lasry and Lions \(2007\)](#); [Weintraub et al. \(2008\)](#) (and follow up literature). Similarly, for teams, it is shown that a globally optimal solution of a team with a large number of players can be approximated by an optimal solution of a team with infinite number of players [Arabneydi and Mahajan \(2014, 2016\)](#); [Bauerle \(2023\)](#); [Elliott et al. \(2013\)](#).

Our main contribution is to show that a similar high-level idea works for games among teams. In particular:

- We present a model of multi-stage games among teams where each team has a large number of agents.
- We propose a mean-field based Markov perfect equilibrium (MF-MPE) for the game among teams and present a dynamic programming decomposition to compute MF-MPE.
- When each team has a large number of players, we approximate the system with a game among teams with infinite players and show that any MF-MPE of the infinite population game among teams is an ε -approximate MF-MPE of the large population game among teams, where we provide an upper bound on ε . Our approximation results are different from typical infinite-population approximations in mean-field games (MFG). In MFG, as the number of players becomes large, each player has negligible influence on the dynamics and payoffs of other agents. However, in the game among teams, the number of teams remains fixed, so a different approach is required to establish the approximation results.

The rest of the paper is organized as follows. In Sect. 2, we review the relevant literature. In Sect. 3, we present the system model and problem formulation for mean-field game among teams. In Sect. 4, we present an equivalent game among

coordinators for teams and show that any MPE (Markov perfect equilibrium) of the game among coordinators is a Team-Nash equilibrium of the original game, which we call MF-MPE (mean-field MPE). We also present a dynamic program to characterize MF-MPE. In Sec. 5, we present a mean-field approximation of MF-MPE when each team has a large number of players and we conclude in Sect. 6

2 Literature Overview

There has been some recent interest in modeling and analyzing games among teams. A dynamic game among teams with delayed intra-agent information sharing is considered in [Tang et al. \(2022\)](#), where common-information based refinements for Team-Nash equilibrium are presented. The results of [Tang et al. \(2022\)](#) consider teams with general heterogeneous agents and no simplifications due to homogeneity and large number of agents are considered. Such mean-field approximations for games among teams are considered in [Pedram and Tanaka \(2019\)](#); [Yu et al. \(2021\)](#); [Sanjari et al. \(2023\)](#); [Guan et al. \(2023\)](#) and we summarize their results below.

[Pedram and Tanaka \(2019\)](#) are motivated by transportation networks and focus on designing incentive mechanisms to mitigate congestion in routing games over graphs. Their main result is proposing a toll mechanism, establishing a mean-field limit of the resulting large population game among teams, and showing that the resulting mean-field equilibrium can be computed efficiently. Note that the dynamics and reward models considered in our setup are different from [Pedram and Tanaka \(2019\)](#).

[Yu et al. \(2021\)](#) are motivated large firms aiming to develop new products or technologies with a rank-based rewards, where each team member contributes to the jump intensity of a common Poisson process, and the reward received by each team depends on its ranking. Different settings are considered for determining the team size: (i) by a central planner; (ii) by a team manager; and (iii) by team members voting in a partnership. Their main result is to propose a relative performance criteria which enables an explicit computation of the equilibrium solution. They also establish that the equilibrium eliminates the free riding problem. Note that the dynamics and reward models considered in our setup are different from [Yu et al. \(2021\)](#).

[Sanjari et al. \(2023\)](#) consider a generalization of Witsenhausen’s intrinsic model ([Witsenhausen 1975](#)) in where there are multiple agents that act sequentially. Each agent acts only once. Agents belong to one of finite number of teams and all agents within a team are exchangeable. They establish that the setting with large number of agents in each team can be approximated by an infinite population mean-field limit. Moreover, there exists a Nash equilibrium for the infinite population limit which is symmetric (i.e., each agent in the team considers identical strategies) and independently randomized. Note that in contrast to [Sanjari et al. \(2023\)](#) we consider dynamic games (where each agent acts more than once) and consider a refinement of Markov perfect equilibrium.

Guan et al. (2023) zero-sum games between two teams is considered. Under some technical assumptions, it is established that the optimal strategies of each team can be computed via a common-information based dynamic programming decomposition. It is then established that such games among teams with large number of agents can be approximated by their mean-field limit. Note that in contrast to Guan et al. (2023), we consider general sum games.

3 Model and problem formulation

3.1 Model of mean-field games among teams

3.1.1 State and action spaces

Consider a multi-agent system with K teams of homogeneous agents. Team $k \in \mathcal{K} := \{1, \dots, K\}$ consists of $N^{(k)}$ agents denoted by the set $\mathcal{N}^{(k)}$, with state space $\mathcal{S}^{(k)}$ and action space $\mathcal{A}^{(k)}$. We assume that $\mathcal{S}^{(k)}$ and $\mathcal{A}^{(k)}$ are finite sets. At time t , the state and action of a generic agent i in the team k are denoted by $S_t^i \in \mathcal{S}^{(k)}$ and $A_t^i \in \mathcal{A}^{(k)}$, respectively. Moreover, let

$$S_t^{(k)} = (S_t^i)_{i \in \mathcal{N}^{(k)}} \quad \text{and} \quad A_t^{(k)} = (A_t^i)_{i \in \mathcal{N}^{(k)}}$$

denote the states and actions of all agents in team k and

$$S_t = (S_t^{(k)})_{k \in \mathcal{K}} \quad \text{and} \quad A_t = (A_t^{(k)})_{k \in \mathcal{K}}$$

denote the global state and actions of the entire system.

Given $s^{(k)} = (s^i)_{i \in \mathcal{N}^{(k)}}$, $s^i \in \mathcal{S}^{(k)}$, we use $\xi(s^{(k)})$ to denote the mean field (or empirical distribution) of $s^{(k)}$, i.e.,

$$\xi(s^{(k)}) = \frac{1}{N^{(k)}} \sum_{i \in \mathcal{N}^{(k)}} \delta_{s^i},$$

where δ_s is a Dirac delta measure centered at s . We use $Z_t^{(k)} = \xi(S_t^{(k)})$ to denote the mean field of the team k , $\mathcal{Z}^{(k)}$ to denote the space of realizations of $Z_t^{(k)}$, $Z_t = (Z_t^{(k)})_{k \in \mathcal{K}}$ to denote the mean field of the entire system, and \mathcal{Z}^* to denote the space of realizations of Z_t . With a slight abuse of notation, we use $Z_t = \xi(S_t)$ to denote the mean-field of the entire system corresponding to the global state S_t .

3.1.2 System dynamics

We use $(s_{1:T}, a_{1:T})$ to denote a realization of $(S_{1:T}, A_{1:T})$ and $z_t^{(k)} = \xi(s_t^{(k)})$ to denote the mean field at time t . We assume that the initial states of all agents are independent, i.e.,

$$\mathbb{P}(S_1 = s_1) = \prod_{k \in \mathcal{K}} \prod_{i \in \mathcal{N}^{(k)}} \mathbb{P}(S_1^i = s_1^i) =: \prod_{k \in \mathcal{K}} \prod_{i \in \mathcal{N}^{(k)}} P_0^{(k)}(s_1^i),$$

where $P_0^{(k)}$ denotes the initial state distribution of agents in team k . The global state of the system evolves independently across agents in a controlled Markov manner, i.e.,

$$\begin{aligned} \mathbb{P}(S_{t+1} = s_{t+1} \mid S_{1:t} = s_{1:t}, A_{1:t} = a_{1:t}) \\ = \prod_{k \in \mathcal{K}} \prod_{i \in \mathcal{N}^{(k)}} \mathbb{P}(S_{t+1}^i = s_{t+1}^i \mid S_t = s_t, A_t = a_t). \end{aligned}$$

Finally, all agents within a team are exchangeable. So the state evolution of a generic agent depends on the states and actions of other agents only through the mean-fields of the states of the teams, i.e., for agent i in team k ,

$$\begin{aligned} \mathbb{P}(S_{t+1}^i = s_{t+1}^i \mid S_t = s_t, A_t = a_t) &= \mathbb{P}(S_{t+1}^i = s_{t+1}^i \mid S_t^i = s_t^i, A_t^i = a_t^i, Z_t = z_t) \\ &=: P^{(k)}(s_{t+1}^i \mid s_t^i, a_t^i, z_t), \end{aligned}$$

where $P^{(k)}$ denotes the controlled transition matrix for team k .

Combining all of the above, we have

$$\mathbb{P}(S_{t+1} = s_{t+1} \mid S_{1:t} = s_{1:t}, A_{1:t} = a_{1:t}) = \prod_{k \in \mathcal{K}} \prod_{i \in \mathcal{N}^{(k)}} P^{(k)}(s_{t+1}^i \mid s_t^i, a_t^i, z_t). \quad (1)$$

3.1.3 Cost function

There is a cost function: $c_t^{(k)} : \mathcal{S}^{(k)} \times \mathcal{A}^{(k)} \times \mathcal{Z}^* \rightarrow \mathbb{R}$ associated with each agent in team k . The per-step cost incurred by team k is the average of the cost associated with all agents in the team, i.e.,

$$C_t^{(k)} = \frac{1}{N^{(k)}} \sum_{i \in \mathcal{N}^{(k)}} c_t^{(k)}(S_t^i, A_t^i, Z_t). \quad (2)$$

3.1.4 Information structure and control laws

We assume that the system has mean-field sharing information-structure ([Arabneydi and Mahajan 2014](#)), i.e., each agent has access to its local state S_t^i and the history of mean-field $Z_{1:t}$ of all teams. Thus, the information available to agent i is given by:

$$I_t^i = \{S_t^i, Z_{1:t}\}. \quad (3)$$

In addition, we assume that all agents in team k use identical¹ (stochastic) control laws: $\pi_t^{(k)} : \mathcal{S}^{(k)} \times \mathcal{Z}^{*t} \rightarrow \Delta(\mathcal{A}^{(k)})$ to choose the action at time t , i.e.,

$$A_t^i \sim \pi_t^{(k)}(S_t^i, Z_{1:t}), \quad (4)$$

¹ Restricting attention to identical strategies for all agents in a team may result in a loss of optimality for that team. However, as argued in [Arabneydi and Mahajan \(2014\)](#), such a restriction is often justifiable due to other considerations such as simplicity, robustness and fairness.

where each agent randomizes independently. Let $\pi^{(k)} := (\pi_1^{(k)}, \dots, \pi_T^{(k)})$ denote the strategy of team k . We use $\pi^{(-k)}$ to denote the strategy of all teams other than k . Given a strategy $\pi = (\pi^{(k)}, \pi^{(-k)})$ for all teams, the expected total cost incurred by team k is given by the following:

$$J^{(k)}(\pi^{(k)}, \pi^{(-k)}) = \mathbb{E}^{\pi^{(k)}, \pi^{(-k)}} \left[\sum_{t=1}^T C_t^{(k)} \right]. \quad (5)$$

It is assumed that the system dynamics $(P_0^{(k)})_{k \in \mathcal{K}}$, $(P^{(k)})_{k \in \mathcal{K}}$, the cost functions $(c^{(k)})_{k \in \mathcal{K}}$ and the information structure are common knowledge for all agents. Each team is interested in minimizing the total expected cost incurred over a finite horizon. Following Tang et al. (2022), we say that a strategy profile $\pi^* = (\pi^{*(k)})_{k \in \mathcal{K}}$ is a **Team-Nash equilibrium** if for all teams $k \in \mathcal{K}$ and all other strategy profiles $\pi^{(k)}$ for team k , we have:

$$J^{(k)}(\pi^{*(k)}, \pi^{*(-k)}) \leq J^{(k)}(\pi^{(k)}, \pi^{*(-k)}). \quad (6)$$

In the sequel, we refer to the model defined above as mean-field game among teams (MFGT). We are interested in the following:

Game 1 *Identify a Team-Nash equilibrium of the mean-field game among teams (MFGT) model defined above.*

3.2 Salient features of the model

Some salient features of the MFGT model are as follows:

3.2.1 Dynamic Game with asymmetric information

MFGT is a dynamic game (also called stochastic game), where there is a state space model which describes the evolution of the state of the environment. Agents have imperfect and asymmetric information about the state of the environment.

3.2.2 All agents in a team are homogeneous

We have assumed that all agents in a team have homogeneous dynamics and cost functions. This assumption is made only for ease of exposition. It is conceptually straightforward to extend the discussion to models with heterogeneous agents where the agents have multiple types. In fact, such a model can be converted into a model with homogeneous agents by augmenting the state space and considering the type of each agent to be a (static) component of its state.

3.2.3 All agents in a team play identical strategies

In a general Team-Nash equilibrium (Tang et al. 2022), all agents in a team are allowed to deviate together and in a correlated manner. However, we have imposed an additional assumption that all agents in a team play identical strategies. Under this assumption, when agents in a team consider deviations, they only consider deviations in which all agents of that team are playing identical (randomized) strategies.

3.2.4 Agents in a team randomize independently

We have assumed that all agents randomize independently. In principle, at each time, agents in the same team could randomize in a correlated manner, but we do not consider that setup in this paper. Correlated randomizations can be obtained by augmenting the state space of all agents in a team with a common correlating random variable, which is independent over time.

4 Mean-field based Markov perfect equilibrium (MF-MPE)

4.1 Road map of the solution approach

Our main conceptual idea is as follows. First, we start by an alternative, but equivalent representation of the mean-field in terms of the state counts. Borrowing the idea of prescriptions (i.e., partially evaluated strategies) from decentralized stochastic control (Nayyar et al. 2013), we show that state-counts (and, therefore, the mean-field) is a controlled Markov process controlled by prescriptions. Using these results, we show that for any team k and any arbitrary but fixed strategy π^{-k} of all teams other than k , the mean-field $\{Z_t\}_{t \geq 1}$ is an *information state* for players of team k in the sense of Subramanian et al. (2022) (see Proposition 1). Therefore, we can follow the idea of game between virtual players introduced by Nayyar et al. (2014), to propose a game between K virtual players, where virtual player $k \in \mathcal{K}$ observes the mean-field Z_t and chooses the prescription for all agents in team k . We call this game, Game 2, and show that Game 2 is equivalent to Game 1. In particular, any Nash equilibrium of Game 2 is a Team-Nash equilibrium of Game 1 (see Theorem 1).

Game 2 is a dynamic game with perfect information. Following Maskin and Tirole (1988a,b), we can identify a dynamic program to characterize the Markov perfect equilibria (i.e., a Nash equilibria where all players play Markov strategies) of Game 2 (see Theorem 2). By Theorem 1, any MPE identified by the dynamic program of Theorem 2 is a Team-Nash equilibrium of Game 1. We call such Team-Nash equilibria as MF-MPE (mean-field Markov perfect equilibria).

4.2 A count-based representation of the mean-field

We start by an alternative, but equivalent, representation of the mean-field in terms of state counts. Count-based representation has been explored earlier in collective decentralized POMDPs (Nguyen et al. 2017) for systems with a single team. We generalize these ideas to mean-field systems with multiple teams. We start by defining different types of counts:

- *State counts*, denoted by $M_t^{(k)}$, which count the number of agents of team k in each state and is given by

$$M_t^{(k)}(s) = \sum_{i \in N^{(k)}} \mathbb{1}\{S_t^i = s\}, \quad \forall s \in \mathcal{S}^{(k)}.$$

- *State-action counts*, denoted by $\bar{M}_t^{(k)}$, count the number of agents of team k in each state-action pair and is given by

$$\bar{M}_t^{(k)}(s, a) = \sum_{i \in N^{(k)}} \mathbb{1}\{S_t^i = s, A_t^i = a\}, \quad \forall s, a \in \mathcal{S}^{(k)} \times \mathcal{A}^{(k)}.$$

- *State-action-next state counts*, denoted by $\widehat{M}_t^{(k)}$, count the number of agents of team k in each state-action-next-state tuples, and is given by:

$$\widehat{M}_t^{(k)}(s, a, s') = \sum_{i \in N^{(k)}} \mathbb{1}\{S_t^i = s, A_t^i = a, S_{t+1}^i = s'\}, \quad \forall s, a, s' \in \mathcal{S}^{(k)} \times \mathcal{A}^{(k)} \times \mathcal{S}^{(k)}.$$

Similar to mean-field, we use $M_t = (M_t^{(k)})_{k \in \mathcal{K}}$, $\bar{M}_t = (\bar{M}_t^{(k)})_{k \in \mathcal{K}}$, and $\widehat{M}_t = (\widehat{M}_t^{(k)})_{k \in \mathcal{K}}$ to denote the counts for the entire system.

Note that the state counts are equivalent to the empirical mean field, i.e., $M_t^{(k)} = N^{(k)} Z_t^{(k)}$, or equivalently, $Z_t^{(k)} = M_t^{(k)} / N^{(k)}$. We use $Z_t = \mu(M_t)$ to denote the vector of mean-fields equivalent to the vector M_t of counts.

The main advantage of a count-based representation is that it captures the inherent symmetry in the model due to the homogeneity of agents. For example, the state dynamics (1) may be written as

$$\begin{aligned} \mathbb{P}(S_{t+1} = s_{t+1} \mid S_{1:t} = s_{1:t}, A_{1:t} = a_{1:t}) \\ = \prod_{k \in \mathcal{K}} \prod_{\substack{(s^i, a^i, s_+^i) \in \\ \mathcal{S}^{(k)} \times \mathcal{A}^{(k)} \times \mathcal{S}^{(k)}}} P^{(k)}(s_+^i \mid s, a^i, \xi(s_t)) \widehat{m}_t^{(k)}(s^i, a^i, s_+^i), \end{aligned} \quad (7)$$

where $\widehat{m}_t^{(k)}$ denotes the realization of $\widehat{M}_t^{(k)}$ along the sample path $(s_{1:t+1}^i, a_{1:t}^i)$. Furthermore, the average cost $C^{(k)}$ (which is given by (2)) may be written as

$$C_t^{(k)} = \sum_{\substack{s^i \in \mathcal{S}^{(k)} \\ a^i \in \mathcal{A}^{(k)}}} c^{(k)}(s^i, a^i, \xi(s_t)) \bar{m}_t^{(k)}(s^i, a^i), \quad (8)$$

where $\bar{m}_t^{(k)}$ denotes the realization of $\bar{M}_t^{(k)}$ along the sample path $(s_{1:t}^i, a_{1:t}^i)$. Notice that computing the right-hand-side expressions in (7) and (8) require only aggregate information, namely counts \widehat{m}_t^k and \bar{m}_t^k for each team k . We will exploit such symmetries to present a simpler and equivalent representation of Game 1.

4.3 Dynamics of the counts

Given any strategy $\pi = (\pi^{(1)}, \dots, \pi^{(K)})$ for the system and any realization $z_{1:T}$ of the mean field $Z_{1:T}$ (or equivalently, for any realization m_t of the state counts M_t and $z_t = \mu(m_t)$), we define the following partial functions, which we call *prescriptions* following the terminology of [Nayyar et al. \(2013\)](#):

$$\gamma_t^{(k)} = \pi_t^{(k)}(\cdot, z_{1:t}), \quad \forall k \in \mathcal{K}. \quad (9)$$

When the realization $z_{1:t}$ is given, $\gamma_t^{(k)}$ is a function from $\mathcal{S}^{(k)}$ to $\Delta(\mathcal{A}^{(k)})$. When $Z_{1:t}$ is a random variable, $\pi^{(k)}(\cdot, Z_{1:t})$ is a random function from $\mathcal{S}^{(k)}$ to $\Delta(\mathcal{A}^{(k)})$ and we denote this random function by $\Gamma_t^{(k)}$. We use γ_t to denote $(\gamma_t^{(1)}, \dots, \gamma_t^{(K)})$ and Γ_t to denote $(\Gamma_t^{(1)}, \dots, \Gamma_t^{(K)})$.

Now we describe the dynamics of the state counts given the current state count m_t and prescription $\gamma_t^{(k)}$.

4.3.1 From state-counts to state-action counts

Let $m_t^{(k)}$ and $\bar{m}_t^{(k)}$ be consistent values of state counts and state-action counts, i.e.,

$$m_t^{(k)}(s) = \sum_{a \in \mathcal{A}^{(k)}} \bar{m}_t^{(k)}(s, a), \quad \forall s \in \mathcal{S}^{(k)}.$$

Then, from a basic combinatorial counting argument, we get

$$\begin{aligned} \mathbb{P}(\bar{M}_t^{(k)} = \bar{m}_t^{(k)} \mid M_t = m_t, \Gamma_t^{(k)} = \gamma_t^{(k)}) \\ = \prod_{s \in \mathcal{S}^{(k)}} \left[\frac{m_t^{(k)}(s)!}{\prod_{a \in \mathcal{A}^{(k)}} \bar{m}_t^{(k)}(s, a)} \prod_{a \in \mathcal{A}^{(k)}} \gamma_t^{(k)}(a|s)^{\bar{m}_t^{(k)}(s, a)} \right], \end{aligned} \quad (10)$$

which is a product of multinomial distributions.

4.3.2 From state-action counts to state-action-state counts

Let $m_t^{(k)}$, $\bar{m}_t^{(k)}$ and $\hat{m}_t^{(k)}$ be consistent values of state counts, state-action counts and state-action-next state counts, i.e.,

$$\bar{m}_t^{(k)}(s, a) = \sum_{s' \in \mathcal{S}^{(k)}} \hat{m}_t^{(k)}(s, a, s'), \quad \forall s, a \in \mathcal{S}^{(k)} \times \mathcal{A}^{(k)}.$$

Let z_t be the mean-field corresponding to $(m_t^{(1)}, \dots, m_t^{(K)})$. Then, from a basic combinatorial counting argument, we get

$$\begin{aligned} \mathbb{P}(\widehat{M}_t^{(k)} = \hat{m}_t^{(k)} \mid \bar{M}_t^{(k)} = \bar{m}_t^{(k)}, M_t = m_t) \\ = \prod_{\substack{s \in \mathcal{S}^{(k)} \\ a \in \mathcal{A}^{(k)}}} \left[\frac{\bar{m}_t^{(k)}(s, a)!}{\prod_{s'} \hat{m}_t^{(k)}(s, a, s')!} \prod_{s' \in \mathcal{S}^{(k)}} P^{(k)}(s' \mid s, a, z_t)^{\hat{m}_t^{(k)}(s, a, s')} \right], \end{aligned} \quad (11)$$

which is also a product of multinomial distributions.

4.3.3 From state-action-state counts to updated state counts

Let $\widehat{m}_t^{(k)}$ and $m_{t+1}^{(k)}$ be consistent values of state-action-state and state counts, i.e.,

$$m_{t+1}^{(k)}(s') = \sum_{s \in \mathcal{S}^{(k)}} \sum_{a \in \mathcal{A}^{(k)}} \widehat{m}_t^{(k)}(s, a, s'), \quad \forall s' \in \mathcal{S}^{(k)}. \quad (12)$$

Thus, if we “marginalize” the sampled state-action-next state count $\widehat{m}_t^{(k)}$, we will obtain the state count $m_{t+1}^{(k)}$.

4.4 A game between virtual players

We start by establishing that the mean-field Z_t is an information state (in the sense of [Subramanian et al. \(2022\)](#)) for every team.

Proposition 1 *For any strategy $\pi = (\pi^{(1)}, \dots, \pi^{(K)})$ and any time $t \in \{1, \dots, T\}$, the mean field $\{Z_t\}_{t \geq 1}$ is an information state for every team k , i.e., the following properties hold for any realization $(z_{1:t}, \gamma_{1:t})$ of $(Z_{1:t}, \Gamma_{1:t})$:*

(P1) **Sufficient for performance evaluation:**

$$\begin{aligned} \mathbb{E}[C_t^{(k)} \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}] &= \mathbb{E}[C_t^{(k)} \mid Z_t = z_t, \Gamma_t^{(k)} = \gamma_t^{(k)}] \\ &=: \ell_t^{(k)}(z_t, \gamma_t^{(k)}) \end{aligned} \quad (13)$$

(P2) **Sufficient for predicting itself:** for any $z = (z^{(1)}, \dots, z^{(K)}) \in \mathcal{Z}^*$, we have

$$\begin{aligned} &\mathbb{P}(Z_{t+1} = z \mid Z_{1:t} = z_{1:t}, \Gamma_{1:t} = \gamma_{1:t}) \\ &= \prod_{k \in \mathcal{K}} \mathbb{P}(Z_{t+1}^{(k)} = Z^{(k)} \mid Z_t = z_t, \Gamma_t^{(k)} = \gamma_t^{(k)}) \\ &=: \prod_{k \in \mathcal{K}} Q^{(k)}(z^{(k)} \mid z_t, \gamma_t^{(k)}), \\ &=: Q(z \mid z_t, \gamma_t), \end{aligned} \quad (14)$$

where $Q^{(k)}(z^{(k)} \mid z_t, \gamma_t^{(k)})$ may be computed by combining (10)–(12).

Proof Property (P1) follows from (8) and (10). Property (P2) follows from (10)–(12) and the fact that there is a one-to-one relationship between the mean field Z_t and the state counts M_t .

Following [Nayyar et al. \(2014\)](#), we consider a stochastic game between K virtual players. At time t , the state is $Z_t = (Z_t^{(1)}, \dots, Z_t^{(K)}) \in \mathcal{Z}^*$; virtual player $k \in \mathcal{K}$ observes Z_t , chooses the prescription $\gamma_t^{(k)} : \mathcal{S}^{(k)} \rightarrow \Delta(\mathcal{A}^{(k)})$, and

incurs a per-step cost $\ell^{(k)}(Z_t, \gamma_t^{(k)})$ given by (13). The initial state Z_1 has a probability mass function given by:

$$\begin{aligned} \mathbb{P}(Z_1 = z) &= \prod_{k \in \mathcal{K}} \mathbb{P}(Z_1^{(k)} = z^{(k)}) \\ &= \prod_{k \in \mathcal{K}} \sum_{s^{(k)} \in (\mathcal{S}^{(k)})^{N^{(k)}}} \prod_{i \in N^{(k)}} P_0^{(k)}(s_1^i). \end{aligned} \quad (15)$$

The state Z_t evolves in a controlled Markov manner according to (14).

The information available to the virtual player at time t is the history of mean-fields $Z_{1:t}$. Virtual player k selects the prescription according to a strategy $\varphi^{(k)}$, i.e.,

$$\Gamma_t^{(k)} = \varphi^{(k)}(Z_{1:t}).$$

Let $\varphi^{(k)} = (\varphi_1^{(k)}, \dots, \varphi_T^{(k)})$ denote the strategy of virtual player k . Then, the total cost incurred by virtual player k is given by:

$$L^{(k)}(\varphi^{(k)}, \varphi^{(-k)}) = \mathbb{E} \left[\sum_{t=1}^T \ell_t^{(k)}(Z_t, \Gamma_t^{(k)}) \right]. \quad (16)$$

We are interested in the following:

Game 2 *Given the system model described above, identify a Nash equilibrium strategy $\varphi^* = (\varphi^{*(k)})_{k \in \mathcal{K}}$, i.e., $\varphi_t^{*(k)}: Z_t \mapsto \Gamma_t^{(k)}$, i.e., for any other strategy $\varphi = (\varphi^{(k)})_{k \in \mathcal{K}}$, we have*

$$L^{(k)}(\varphi^{*(k)}, \varphi^{*(-k)}) \leq L^{(k)}(\varphi^{(k)}, \varphi^{*(-k)}), \quad \forall k \in \mathcal{K}. \quad (17)$$

4.5 Relationship between Games 1 and 2

We have the following result that connects the solutions of Game 1 and Game 2.

Theorem 1 *Let $\pi = (\pi^{(1)}, \dots, \pi^{(K)})$ be a Team-Nash equilibrium of Game 1. Define a strategy $\varphi = (\varphi^{(1)}, \dots, \varphi^{(K)})$ for Game 2 as follows: for any $z_{1:t} \in \mathcal{Z}^{*t}$:*

$$\varphi_t^{(k)}(z_{1:t}) = \pi_t^{(k)}(\cdot, z_{1:t}). \quad (18)$$

Then φ is a Nash equilibrium for Game 2.

Conversely, let $\varphi = (\varphi^{(1)}, \dots, \varphi^{(K)})$ be any Nash equilibrium for Game 2. Define a strategy $\pi = (\pi^{(1)}, \dots, \pi^{(K)})$ for Game 1 as follows: for any $s \in \mathcal{S}^{(k)}$ and $z \in \mathcal{Z}^$,*

$$\pi_t^{(k)}(s, z_{1:t}) = \varphi_t^{(k)}(z_{1:t})(s). \quad (19)$$

Then π is a Team-Nash equilibrium of Game 1.

See Appendix A

4.6 Markov perfect equilibrium for Game 2

Game 2 among virtual players is a game with perfect information since all players choose prescriptions based on the history $Z_{1:t}$ of mean-field which is common knowledge between the players. Proposition 1 implies that we can view the current mean field Z_t as the “state” of the system. Following Maskin and Tirole (1988a,b), we restrict our attention to the Markov perfect equilibrium for Game 2, which can be thought of as a subgame perfect equilibrium of Game 2 where all virtual players are playing Markov strategies which map current state to prescriptions. Such a Markov perfect equilibrium can be obtained using dynamic programming as follows (Maskin and Tirole 1988a,b).

Theorem 2 Consider a strategy profile $\psi = (\psi^{(k)})_{k \in \mathcal{K}}$, where each virtual player is playing a Markov strategy, i.e., $\psi_t^{(k)}: Z_t \mapsto \Gamma_t^{(k)}$.

A necessary and sufficient condition for ψ to be a Markov perfect equilibrium for Game 2 is that it satisfy the following conditions:

1. For each possible realization z_T of Z_T , define the value function for virtual player k :

$$V_T^{(k)}(z_T) = \min_{\gamma_T^{(k)}} \ell_T^{(k)}(z_T, \gamma_T^{(k)}). \quad (20)$$

Then, $\psi_T^{(k)}(z_T)$ must be a minimizing $\gamma_T^{(k)}$ in (20).

2. For $t \in \{T-1, \dots, 1\}$ and for each possible realization z_t of Z_t , recursively define the value function for virtual player k :

$$V_t^{(k)}(z_t) = \sum_{i \in N^{(k)}} \mathbb{E}[\ell_t^{(k)}(z_t, \gamma_t^{(k)}) + V_{t+1}^{(k)}(Z_{t+1}) \mid \mathcal{F}^{(k)}] \quad (21)$$

where the event $\mathcal{F}^{(k)} = \{Z_t = z_t, \Gamma_t^{(k)} = \gamma_t^{(k)}, \Gamma_t^{(-k)} = \psi_t^{(-k)}(z_t)\}$ and the expectation is with respect to the distribution (14). Then, $\psi_t^{(k)}(z_t)$ must be a minimizing $\gamma_t^{(k)}$ in (21).

Remark 1 Theorem 2 states that the Markov perfect equilibrium for the virtual players can be obtained by dynamic programming. Let ψ be such a Markov perfect equilibrium. Let π be the policy obtained by (19). Then, by Theorem 1, π is a Team-Nash equilibrium of Game 1, which we call *mean-field based Markov perfect equilibrium* (MF-MPE).

Remark 2 In general, solving the dynamic program of Theorem 2 suffers from the curse of dimensionality. The space $\mathcal{Z}^{(k)}$ has at most $(N^{(k)} + 1)^{|S^{(k)}|}$ elements, which increase polynomially with the number $N^{(k)}$ of agents. Another added challenge is that it is complicated to explicitly construct the conditional distribution $Q^{(k)}$ used in property (P2) of Proposition 1. An approach to bypass both difficulties is to use sampling-based reinforcement learning techniques Chang et al. (2007) which do not explicitly construct the action-value function. Sampling based techniques are particularly efficient in our setting because the condition distribution $Q^{(k)}$ are defined via a series of products of multinomial distributions (10) and (11), so we can efficiently sample from $Q^{(k)}$ without constructing it explicitly.

Remark 3 Sampling from the multinomial distribution is still computationally involved when the number $N^{(k)}$ of agents in each team is large. In such a setting, one option is to approximate the dynamics $Q^{(k)}$ by using the multivariate Gaussian approximation to the multinomial distribution.

5 Mean-field approximation

In this section, we consider the setting in which each team has a large number of players. We approximate the system by an infinite population mean field limit and show that any MPE of the infinite population game is an ε -Team-Nash equilibrium of the original game, where $\varepsilon = (\varepsilon_k)_{k \in \mathcal{K}}$ and $\varepsilon_k = \mathcal{O}(1/\sqrt{N^{(k)}})$.

The infinite population approximation provides a drastic simplification of the dynamic program of Theorem 2 because in the infinite population mean-field system, the mean-field is equivalent to the statistical distribution of the agents and, therefore, evolves in a deterministic manner.

5.1 Regularity conditions on the system

When there are $N^{(k)}$ agents in team k , the mean-field $Z_t^{(k)}$ takes values in $\mathcal{Z}^{(k)}$, which is the set of all probability mass functions (PMFs) $z^{(k)}$ on $\mathcal{S}^{(k)}$ such that $N^{(k)}z^{(k)}$ is a vector of non-negative integers. Note that $\mathcal{Z}^{(k)} \subset \Delta(\mathcal{S}^{(k)})$, the set of all PMFs on $\mathcal{S}^{(k)}$. To understand the behavior of the finite population system as the number of agents become large, we first extend the domain of $z^{(k)}$ in the cost function and the transition functions from $\mathcal{Z}^{(k)}$ to $\Delta(\mathcal{S}^{(k)})$. In particular, we let $\tilde{\mathcal{Z}}^*$ denote the set $\prod_{k \in \mathcal{K}} \Delta(\mathcal{S}^{(k)})$ and assume that the per-step cost function $c_t^{(k)}$ is a function from $\mathcal{S}^{(k)} \times \mathcal{A}^{(k)} \times \tilde{\mathcal{Z}}^*$ to \mathbb{R} and the controlled transition matrix $P^{(k)}$ is a transition matrix from $\mathcal{S}^{(k)} \times \mathcal{A}^{(k)} \times \tilde{\mathcal{Z}}^*$ to $\mathcal{S}^{(k)}$.

We start with some formal definitions needed to define the Lipschitz continuity of $c_t^{(k)}$ and $P^{(k)}$. Let $d_s^{(k)}$ be a metric on the state space $\mathcal{S}^{(k)}$, $k \in \mathcal{K}$. Then, based on this metric, let $d_w^{(k)}$ be the Kantorovich metric (also called Wasserstein metric) on $\Delta(\mathcal{S}^{(k)})$ (i.e., the space of mean-fields for each team) $k \in \mathcal{K}$. Define a metric d_W on the set of mean-fields for all teams $\tilde{\mathcal{Z}}^*$ as:

$$d_W(z, \hat{z}) = \sum_{k \in \mathcal{K}} \left(d_w^{(k)}(z^{(k)}, \hat{z}^{(k)}) \right), \quad z, \hat{z} \in \tilde{\mathcal{Z}}^*. \quad (22)$$

5.2 Infinite population mean-field approximation

We now consider an infinite population approximation of Game 2, where we approximate the per-step cost $\ell_t^{(k)}$ and the dynamics $Q^{(k)}$, defined in

Proposition 1 by $\bar{\ell}^{(k)}$ and $\bar{Q}^{(k)}$ defined as follows:

$$\bar{\ell}_t^{(k)}(\bar{z}, \gamma^{(k)}) = \sum_{s \in \mathcal{S}^{(k)}} \bar{z}^{(k)}(s) \ell_t^{(k)}(s, \gamma^{(k)}(s), \bar{z}), \quad (23)$$

$$\bar{Q}^{(k)}(\bar{z}^{(k)} \mid \bar{z}, \gamma^{(k)}) = \mathbf{1}\{\bar{z}^{(k)} = \bar{q}^{(k)}(\bar{z}, \gamma^{(k)})\}, \quad (24)$$

where

$$\bar{q}^{(k)}(\bar{z}, \gamma^{(k)}) = \sum_{s \in \mathcal{S}^{(k)}} \bar{z}^{(k)}(s) P^{(k)}(s' \mid s, \gamma_t^{(k)}(s), \bar{z}_t), \quad (25)$$

and under the mean-field dynamics

$$\mathbb{P}(Z_{t+1} \mid Z_t = z_t, \Gamma_t = \gamma_t) =: \bar{q}(\bar{z}_t, \gamma_t) = \prod_{k \in \mathcal{K}} \bar{q}^{(k)}(\bar{z}, \gamma^{(k)}). \quad (26)$$

Lemma 1 *The infinite population game is an approximation of the finite population game in the following sense.*

1. For any $k \in \mathcal{K}$, we have

$$|\ell_t^{(k)}(z, \gamma^{(k)}) - \bar{\ell}_t^{(k)}(z, \gamma^{(k)})| = 0, \forall z \in \mathcal{Z}^*, \gamma^{(k)}$$

2. For any $k \in \mathcal{K}$, we have

$$d_{\mathcal{W}}(\bar{q}(z_t, \gamma_t), Q(\cdot \mid z_t, \gamma_t)) \leq \sum_{k \in \mathcal{K}} \frac{\kappa}{\sqrt{N^{(k)}}} =: \delta_t$$

where κ is a constant that depends on the state spaces $\mathcal{S}^{(k)}$ and the metric d_s .

Proof The first part of the lemma follows from the definitions of $\ell^{(k)}$ and $\bar{\ell}^{(k)}$. For the second point, we first note that from (Sinha and Mahajan 2023, Lemma 4) we have:

$$d_{\mathcal{W}}(\bar{q}(z_t, \gamma_t), Q(\cdot \mid z_t, \gamma_t)) \leq \sum_{k \in \mathcal{K}} d_{\mathcal{W}}(Q^{(k)}(\cdot \mid z_t, \gamma_t^{(k)}), \bar{q}^{(k)}(z_t, \gamma_t^{(k)}),$$

Furthermore, concentration of empirical measure to statistical measure with respect to the Wasserstein distance (Sommerfeld et al. 2018) implies that

$$d_{\mathcal{W}}(\bar{q}^{(k)}(z_t, \gamma_t^{(k)}), Q^{(k)}(\cdot \mid z_t, \gamma_t^{(k)})) \leq \frac{\kappa}{\sqrt{N^{(k)}}},$$

where κ is a constant that depends on the state spaces $\mathcal{S}^{(k)}$ and the metric d_s . Combining the above two equations implies the second result.

Since the infinite population approximation is a Markov game, its Markov perfect equilibrium is characterized as follows (Maskin and Tirole 1988a,b).

Theorem 3 Consider a strategy profile $\bar{\psi} = (\bar{\psi}^{(k)})_{k \in \mathcal{K}}$, where each virtual player is playing a Markov strategy.

A necessary and sufficient condition for $\bar{\psi}$ to be a Markov perfect equilibrium for the mean-field limit of Game 2 is that it satisfy the following conditions:

1. For each possible realization \bar{z}_T of \bar{Z}_T , define the value function for virtual player k :

$$\bar{V}_T^{(k)}(\bar{z}_T) = \min_{\gamma_T^{(k)}} \bar{\ell}_T^{(k)}(\bar{z}_T, \gamma_T^{(k)}). \quad (27)$$

Then, $\bar{\psi}_T^{(k)}(\bar{z}_T)$ must be a minimizing $\gamma_T^{(k)}$ in (27).

2. For $t \in \{T-1, \dots, 1\}$ and for each possible realization \bar{z}_t of \bar{Z}_t , recursively define the value function for virtual player k :

$$\bar{V}_t^{(k)}(\bar{z}_t) = \min_{\gamma_t^{(k)}} \left\{ \bar{\ell}_t^{(k)}(\bar{z}_t, \gamma_t^{(k)}) + \bar{V}_{t+1}^{(k)}((q^{(k)}(\bar{z}_t, \gamma_t^{(k)}))_{k \in \mathcal{K}}) \right\}. \quad (28)$$

Then $\bar{\psi}_t^{(k)}(\bar{z}_t)$ must be a minimizing $\gamma_t^{(k)}$ in (28).

5.3 ϵ -Team-Nash equilibrium

Now we address the fundamental question of the mean-field approximation: is the infinite population approximation good and, if so, in what sense? We show that the any MPE of the infinite population limit is an approximate MPE of the original finite population game, where the approximation error scales as $\mathcal{O}(1/N)$.

Theorem 4 Suppose there exists an MPE $\bar{\psi}$ for the infinite population game such that the corresponding total costs $\bar{L}_t^{(k)}$ are $\mathcal{L}_t^{(k)}$ -Lipschitz, $k \in \mathcal{K}$, $t \in \text{time}$. Then, $\bar{\psi}$ is an $\mathcal{O}(1/\sqrt{N})$ -approximate MPE of the finite population game, where $N = \inf N^{(k)}$. In particular, for any other strategy $\psi^{(k)}$ of team k , we have

$$\bar{L}^{(k)}(\bar{\psi}^{(k)}, \bar{\psi}^{(-k)}) \leq \bar{L}^{(k)}(\psi^{(k)}, \bar{\psi}^{(-k)}) + 2 \sum_{t=1}^T \sum_{k \in \mathcal{K}} \frac{\kappa^{(k)} \mathcal{L}_t^{(k)}}{\sqrt{N^k}}$$

Proof See Appendix B

6 Conclusion

In this paper, we presented a model for mean-field games among teams and presented a common-information based refinement of the Team-Nash equilibrium for this game. This common-information based Markov perfect equilibrium can be obtained by solving coupled dynamic programs. These dynamic programs use the mean field of all teams as a state. In general, solving such dynamic programs suffers from the curse of dimensionality. To circumvent this curse of dimensionality, we use a mean-field limit to approximate finite population teams by an infinite population. We show that a Markov perfect equilibrium obtained using the mean-field approximations is an approximate Markov perfect equilibrium of the original game.

References

- J. Arabneydi and A. Mahajan. Team optimal control of coupled subsystems with mean-field sharing. In *IEEE Conference on Decision and Control*, pages 1669–1674. IEEE, 2014.
- J. Arabneydi and A. Mahajan. Linear quadratic mean field teams: Optimal and approximately optimal decentralized solutions. arXiv:1609.00056v2, 2016.
- N. Bäuerle. Mean field markov decision processes. *Applied Mathematics & Optimization*, 88(1):12, 2023.
- H. Chang, M. Fu, J. Hu, and S. Marcus. *Simulation-based Algorithms for Markov Decision Processes*. Communications and Control Engineering. Springer, 2007. ISBN 9781846286896.
- R. Elliott, X. Li, and Y.-H. Ni. Discrete time mean-field stochastic linear-quadratic optimal control problems. *Automatica*, 49(11):3222–3233, 2013.
- Y. Guan, M. Afshari, and P. Tsiotras. Zero-sum games between mean-field teams: A common information and reachability based analysis, 2023. arxiv:2303.12243.
- M. Huang, P. E. Caines, and R. P. Malhamé. Large-population cost-coupled LQG problems with nonuniform agents: individual-mass behavior and decentralized epsilon-Nash equilibria. *IEEE Trans. Autom. Control*, 52(9):1560–1571, 2007.
- M. Huang, P. E. Caines, and R. P. Malhamé. Social optima in mean field LQG control: centralized and decentralized strategies. *IEEE Trans. Autom. Control*, 57(7):1736–1751, 2012.
- J.-M. Lasry and P.-L. Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1):229–260, 2007.
- J. Marschak and R. Radner. *Economic theory of teams*. Yale University Press, 1972.
- E. Maskin and J. Tirole. A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs. *Econometrica: Journal of the Econometric Society*, pages 549–569, 1988a.
- E. Maskin and J. Tirole. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and edgeworth cycles. *Econometrica: Journal of the Econometric Society*, pages 571–599, 1988b.
- A. Nayyar, A. Mahajan, and D. Teneketzis. Decentralized stochastic control with partial history sharing: A common information approach. *IEEE Trans. Autom. Control*, 58(7):1644–1658, 2013.
- A. Nayyar, A. Gupta, C. Langbort, and T. Başar. Common information based Markov perfect equilibria for stochastic games with asymmetric information: Finite games. *IEEE Transactions on Automatic Control*, 59(3):555–570, 2014.
- D. T. Nguyen, A. Kumar, and H. C. Lau. Collective multiagent sequential decision making under uncertainty. In *AAAI Conference on Artificial Intelligence*, pages 3036–3043, 2017.
- A. R. Pedram and T. Tanaka. Linearly-solvable mean-field approximation for multi-team road traffic games. In *2019 IEEE 58th Conference on Decision and Control (CDC)*, pages 1243–1248, 2019.
- S. Sanjari, N. Saldi, and S. Yüksel. Nash equilibria for exchangeable team against team games, their mean field limit, and role of common randomness. In *American Control Conference*, 2023.
- L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28):307–317, 1953.
- A. Sinha and A. Mahajan. Sensitivity of Whittle index policy to model approximation. Under review, 2023. URL <https://cim.mcgill.ca/~adityam/projects/bandits/preprint/approx-Whittle.pdf>.
- M. Sommerfeld, J. Schrieber, Y. Zemel, and A. Munk. Optimal transport: Fast probabilistic approximation with exact solvers, 2018.
- J. Subramanian, A. Sinha, R. Seraj, and A. Mahajan. Approximate information state for approximate planning and reinforcement learning in partially observed systems. *Journal of Machine Learning Research*, 23(12):1–83, 2022. ISSN 1533-7928.
- D. Tang, H. Tavafoghi, V. Subramanian, A. Nayyar, and D. Teneketzis. Dynamic games among teams with delayed intra-team information Sharing. *Dynamic Games and*

Applications, Feb. 2022.

- P. Tilghman. Will rule the airwaves: A DARPA grand challenge seeks autonomous radios to manage the wireless spectrum. *IEEE Spectrum*, 56(6):28–33, 2019.
- J. von Neumann and O. Morgenstern. *Theory of games and economic behavior*. Princeton University Press, 1944.
- G. Y. Weintraub, C. L. Benkard, and B. Van Roy. Markov perfect industry dynamics with many firms. *Econometrica*, 76(6):1375–1411, 2008.
- H. S. Witsenhausen. The intrinsic model for discrete stochastic control: Some open problems. In A. Bensoussan and J. L. Lions, editors, *Control Theory, Numerical Methods and Computer Systems Modelling*, pages 322–335, Berlin, Heidelberg, 1975. Springer Berlin Heidelberg.
- X. Yu, Y. Zhang, and Z. Zhou. Teamwise mean field competitions. *Applied Mathematics & Optimization*, 84(1):903–942, Dec. 2021.

A Proof of Theorem 1

We first establish some intermediate results.

Lemma 2 *Given a strategy π of Game 1, let φ be a strategy of Game 2 constructed according to (18). Then, for any time t and any realizations $z_{1:t}$ and $\gamma_{1:t}$ of $Z_{1:t}$ and $\Gamma_{1:t}$, we have that*

$$\mathbb{P}^\pi(Z_{1:T} = z_{1:T}, \Gamma_{1:T} = \gamma_{1:t}) = \mathbb{P}^\varphi(Z_{1:T} = z_{1:T}, \Gamma_{1:T} = \gamma_{1:t})$$

Proof For simplicity of notation, we write $\mathbb{P}^\pi(z_{1:t}, \gamma_{1:t})$ instead of $\mathbb{P}^\pi(Z_{1:T} = z_{1:t}, \Gamma_{1:T} = \gamma_{1:t})$ and use similar shortcuts for other terms as well.

The proof proceeds by induction on T . For $t = 1$, both sides of the equation do not depend on the strategies, and the result is trivially true. This forms the basis of induction. Now, suppose that the result is true for t and consider the system at $t + 1$. From property (P2) of Proposition 1, we have

$$\mathbb{P}^\pi(z_{t+1} | z_{1:t}, \gamma_{1:t}) = Q(z_{t+1} | z_t, \gamma_t) \quad (29)$$

and similarly

$$\mathbb{P}^\varphi(z_{t+1} | z_{1:t}, \gamma_{1:t}) = Q(z_{t+1} | z_t, \gamma_t). \quad (30)$$

Furthermore, the construction of strategy φ implies that

$$\mathbb{P}^\pi(\gamma_{t+1} | z_{1:t+1}, \gamma_{1:t}) = \mathbb{P}^\varphi(\gamma_{t+1} | z_{1:t+1}, \gamma_{1:t}). \quad (31)$$

Properties (29)–(31) along with the induction hypothesis implies that

$$\begin{aligned} \mathbb{P}^\pi(z_{1:t+1}, \gamma_{1:t+1}) &= \mathbb{P}^\pi(\gamma_{t+1} | z_{1:t+1}, \gamma_{1:t}) \mathbb{P}^\pi(z_{t+1} | z_{1:t}, \gamma_{1:t}) \mathbb{P}^\pi(z_{1:t}, \gamma_{1:t}) \\ &= \mathbb{P}^\varphi(\gamma_{t+1} | z_{1:t+1}, \gamma_{1:t}) \mathbb{P}^\varphi(z_{t+1} | z_{1:t}, \gamma_{1:t}) \mathbb{P}^\varphi(z_{1:t}, \gamma_{1:t}) \\ &= \mathbb{P}^\varphi(z_{1:t+1}, \gamma_{1:t+1}). \end{aligned} \quad (32)$$

This completes the induction step.

Lemma 3 *Given a strategy φ of Game 2, let π be a strategy of Game 1 constructed according to (19). Then, for any time t and any realizations $z_{1:t}$ and $\gamma_{1:t}$ of $Z_{1:t}$ and $\Gamma_{1:t}$, we have that*

$$\mathbb{P}^\varphi(Z_{1:T} = z_{1:T}, \Gamma_{1:T} = \gamma_{1:t}) = \mathbb{P}^\pi(Z_{1:T} = z_{1:T}, \Gamma_{1:T} = \gamma_{1:t})$$

The proof is similar to the proof of Lemma 2 and is omitted.

Lemma 4 *Given a strategy π of Game 1, let φ be a strategy of Game 2 constructed according to (18). Then,*

$$J^{(k)}(\pi) = L^{(k)}(\varphi).$$

Proof Arbitrarily fix a team $k \in \mathcal{K}$ and consider

$$\mathbb{E}^\pi[C^{(k)}] \stackrel{(a)}{=} \mathbb{E}^\pi[\mathbb{E}[C^{(k)} \mid Z_{1:t}, \gamma_{1:t}]] \stackrel{(b)}{=} \mathbb{E}^\pi[\ell^{(k)}(Z_t, \gamma_t^{(k)})] \quad (33)$$

where (a) follows from the smoothing property of conditional expectation and (b) follows from (P1) in Proposition 1. Therefore,

$$\begin{aligned} J^{(k)}(\pi) &\stackrel{(c)}{=} \mathbb{E}^\pi \left[\sum_{t=1}^T \ell^{(k)}(Z_t, \gamma_t^{(k)}) \right] \\ &\stackrel{(d)}{=} \mathbb{E}^\varphi \left[\sum_{t=1}^T \ell^{(k)}(Z_t, \gamma_t^{(k)}) \right] \\ &= L^{(k)}(\varphi) \end{aligned} \quad (34)$$

where (c) follows from (33) and (d) follows from Lemma 2.

Lemma 5 *Given a strategy φ of Game 2, let π be a strategy of Game 1 constructed according to (19). Then,*

$$J^{(k)}(\pi) = L^{(k)}(\varphi).$$

The proof argument is similar to that of Lemma 4 and is omitted.

Now we are ready to prove the result of Theorem 1. Let π be a NE for Game 1 and let φ be a strategy for Game 2 constructed according to (18). Suppose φ is not a NE for Game 2. That is, there exists a team $k \in \mathcal{K}$ and a strategy $\psi^{(k)}$ for team k such that

$$L^{(k)}(\psi^{(k)}, \varphi^{(-k)}) < L^{(k)}(\varphi^{(k)}, \varphi^{(-k)}).$$

Let $\bar{\pi}^{(k)}$ be the strategy for Game 1 corresponding to $\psi^{(k)}$ constructed according to (19). Then, Lemmas 5 implies that

$$J^{(k)}(\bar{\pi}^{(k)}, \pi^{(-k)}) < J^{(k)}(\pi^{(k)}, \pi^{(-k)})$$

contradicting the fact that π is a NE for Game 1. Therefore, φ must be a NE of Game 2.

The second part of the theorem can be proved by an analogous argument.

B Proof of Theorem 4

Proof Fix a virtual player k and the strategy profile $\bar{\psi}^{(-k)}$ for the other virtual players and consider the best response dynamics at virtual player k given by the dynamic program in Thm. 2. The idea of the proof is to show that the history compression function $\nu_t(z_{1:t}, \gamma_{1:t}) = z_t$ dynamics $(q_t^{(k)})_{k \in \mathcal{K}}$ and the per-step cost $\bar{\ell}_t^{(k)}$ is an approximate information state (AIS) as defined in Subramanian et al. (2022). In particular, we observe that:

$$\mathbb{E}[\ell_t^{(k)}(z, \gamma^{(k)}) - \bar{\ell}_t^{(k)}(\nu_t(z_{1:t}, \gamma_{1:t}), \gamma^{(k)})] = 0 := \varepsilon_t, \quad (35)$$

and

$$d_{\mathcal{K}}(\mathbb{P}(Z_{t+1} \mid Z_t = z_t, \Gamma_t = \gamma_t), q_t(z_t, \gamma_t)) \leq \sum_{k \in \mathcal{K}} \frac{\kappa}{\sqrt{N^{(k)}}} := \delta_t, \quad (36)$$

which follow from Lemma 1. Equations (35), (36) show that $(\nu_t, (q_t^{(k)})_{k \in \mathcal{K}}, \bar{\ell}_t^{(k)})$ is an AIS. Then, the result follows from (Subramanian et al. 2022, Theorem 9).