Multimodal Maximum Entropy Dynamic Games

Oswin So*, Kyle Stachowicz and Evangelos A. Theodorou Autonomous Control and Decision Systems Lab Georgia Institute of Technology, Atlanta, Georgia *Correspondence to: oswinso@gatech.edu

Abstract-Environments with multi-agent interactions often result a rich set of modalities of behavior between agents due to the inherent suboptimality of decision making processes when agents settle for satisfactory decisions. However, existing algorithms for solving these dynamic games are strictly unimodal and fail to capture the intricate multimodal behaviors of the agents. In this paper, we propose MMELQGames (Multimodal Maximum-Entropy Linear Quadratic Games), a novel constrained multimodal maximum entropy formulation of the Differential Dynamic Programming algorithm for solving generalized Nash equilibria. By formulating the problem as a certain dynamic game with incomplete and asymmetric information where agents are uncertain about the cost and dynamics of the game itself, the proposed method is able to reason about multiple local generalized Nash equilibria, enforce constraints with the Augmented Lagrangian framework and also perform Bayesian inference on the latent mode from past observations. We assess the efficacy of the proposed algorithm on two illustrative examples: multi-agent collision avoidance and autonomous racing. In particular, we show that only MMELQGames is able to effectively block a rear vehicle when given a speed disadvantage and the rear vehicle can overtake from multiple positions.

I. INTRODUCTION

Planning in multi-agent scenarios is a challenging task – actions taken by any robot cannot be considered in isolation and must take the response from other agents into account. Classical approaches to tackling this problem adopt a predict-then-plan architecture, where the actions of other agents are assumed to be independent of the ego-agent's actions. However, this assumption breaks down in the context of multi-agent games such as racing, surveillance or autonomous driving where agents may have conflicting interests and it is advantageous to manipulate the behavior of other agents.

Recent works have introduced game-theoretic formulations to address this problem [4, 12, 16, 20, 28, 32, 33]. By assuming that agents act rationally, the optimal actions of other agents can be predicted by solving for the Nash equilibria of the non-cooperative game assuming that the true objectives of other agents are known. However, in practice humans face cognitive limitations and seek decisions that are satisfactory but often suboptimal, a phenomenon described by the concept of "bounded rationality" [29].

This disconnect is often modeled using the the Maximum Entropy (MaxEnt) framework which models actions of agents as being stochastic in nature. MaxEnt has been successfully applied to diverse areas including inverse reinforcement learning [20, 37], forecasting [9], and biology [6]. In particular,



Fig. 1: Two-agent autonomous racing scenario where the rear agent (green) has a higher maximum velocity than the lead agent (blue). Different highlighted colors representing different planned modes. The planned trajectory for the ego agent is shown as a thin blue line. The lead agent is able to reason about and correctly infer the different modes that the rear agent could take. Furthermore, when the ego agent is uncertain about which mode the rear agent will choose, it hedges against both possibilities by planning a path down the middle.

Mehr et al. [20] propose a game-theoretic maximum entropy algorithm for finding nash equilibria policies to dynamic games via an extension of the iLQGames method [12].

However, these methods consider only a single local minimum. As a result, the resulting policies are *unimodal* and fail to capture the *multimodal* nature of solutions under the MaxEnt framework. In the presence of multiple local minima, the true maximum-entropy policy will be multimodal, with one mode corresponding to each local minimum. However, computing expectations over multimodal policies in multi-agent settings quickly leads to combinatorial explosion and is computationally intractable when solving for multimodal Nash equilibria.

By considering a novel MaxEnt dynamic game with incomplete information and information asymmetry, we convert the previous challenge into that of solving a POMDP for the ego agent where the discrete latent variable corresponds to different local generalized MaxEnt Nash equilibria found by running a constrained version of Maximum Entropy Linear Quadratic GAMES (MELQGames) in parallel. This approach further allows for Bayesian inference over the discrete latet variable by considering past observations of controls from non-ego agents.

The main contributions of our work are threefold:

- We propose the constrained MaxEnt dynamic game setting to handle bounded rationality in decison making under inequality constraints and extend MELQGames to solve generalized Nash equilibrium.
- We consider a novel MaxEnt dynamic game with incomplete and asymmetric information. We provide a a computationally efficient solution that allows for Bayesian inference of the latent mode when the underlying control policies are multimodal in nature.
- We showcase the benefits of the proposed algorithm in simulation, including an autonomous racing example. The results demonstrate the superiority of Multimodal Maximum Entropy Linear Quadratic GAMES (MMELQGames) against other game theoretic formulations in terms of successfully predicting agents behavior in dynamic settings.

II. PROBLEM SETUP

In this section, we introduce the constrained MaxEnt dynamic game and the corresponding Generalized Nash Equilibrium Problem (GNEP).

A. Discrete Dynamic Games and Generalized Nash Equilibria

We consider a discrete dynamic game with N players with joint controls $u_t = [u_t^1, \ldots, u_t^N] = [u_t^i, u_t^{\neg i}] \in \mathbb{R}^{n_u}$, where u_t^i denotes the control input of player *i* and we use $u_t^{\neg i}$ to denote the controls excluding the *i*th agent. We denote $x_t \in \mathbb{R}^{n_x}$ the joint state of the system at timestep *t* which evolves under the discrete-time dynamics

$$\boldsymbol{x}_{t+1} = f(\boldsymbol{x}_t, u_t^1, \dots, u_t^N) = f(\boldsymbol{x}_t, \boldsymbol{u}_t).$$
(1)

Each agent's objective is to minimize a corresponding cost function J^i in finite-horizon T with running cost l^i and terminal cost Φ^i , where $\mathbf{U} = [\boldsymbol{u}_1, \dots, \boldsymbol{u}_{T-1}]$ denotes a trajectory

$$J^{i}(\mathbf{U}) = \Phi^{i}(\boldsymbol{x}_{T}) + \sum_{t=1}^{T-1} l^{i}(\boldsymbol{x}_{t}, \boldsymbol{u}_{t}).$$
(2)

The goal of agent *i* is to choose controls U^i that optimize J^i while respecting inequality constraints h^i :

$$\min_{u^i} \quad J^i(\mathbf{U}), \qquad \text{s.t.} \quad h^i(\mathbf{U}) \le 0. \tag{3}$$

Note that both the objective function and the constraints are a function of U^i and $\mathbf{U}^{\neg i}$. Due to the presence of constraints h^i which may couple agent *i*'s feasible control set to other agent's controls $\mathbf{U}^{\neg i}$, (3) is a GNEP [10]. Let

$$\mathcal{U}^{i}(\mathbf{U}^{\neg i}) = \{ U^{i} | h^{i}(\mathbf{U}) \le 0 \},$$
(4)

denote the feasible control set for agent *i* given the controls of other agents $U^{\neg i}$. Solving an (open loop) GNEP amounts to finding controls U^* such that for each agent *i*,

$$J^{i}(U^{i*}, \mathbf{U}^{\neg i*}) \leq J^{i}(U^{i}, \mathbf{U}^{\neg i*}), \quad \forall U^{i} \in \mathcal{U}^{i}(\mathbf{U}^{\neg i}).$$
(5)

However, since solving for the global GNEP is in general intractable, we instead look for local GNEP where (5) is satisfied for all U^i within a local neighborhood of the optimal U^* .

B. Maximum Entropy Dynamic Games

Although the (global) Nash equilibria is a powerful concept, achieving Nash equilibria requires the assumption that all cost functions are known exactly and that each agent acts rationally by solving the GNEP exactly. However, often times we only have an approximation of the cost function for the non-ego agents and other agents may only act with approximate solutions of the GNEP. To take these stochasticities into account, we take a *relaxed control* approach and consider a stochastic control policy $\pi^i(u^i|\mathbf{x})$ with the same deterministic dynamics as in (1). Let $\mathbb{E}_p[\cdot]$ denote the expectation with respect to a distribution p. We introduce an entropy term to the original objective (2) and consider the expected cost under all agent's policies π^i :

$$J^{i}(\pi^{i}) = \mathbb{E}_{\pi} \left[\Phi^{i}(\boldsymbol{x}_{T}) + \sum_{t=1}^{T-1} \left(l^{i}(\boldsymbol{x}_{t}, \boldsymbol{u}_{t}) - \alpha H[\pi^{i}(\cdot | \boldsymbol{x}_{t})] \right) \right],$$
(6)

where $\alpha > 0$ is a temperature term and $H[\pi]$ is the Shannon entropy of π defined as

$$H[\pi] = -\mathbb{E}_{\pi}[\log \pi] = -\int \pi(u) \log \pi(u) \, \mathrm{d}u.$$
 (7)

The resulting GNEP is then formulated similarly to the deterministic case, except that we only require that the constraints h^i hold for the mean controls:

$$\min_{\pi^i} \quad J^i(\mathbf{U}), \qquad \text{s.t.} \quad h^i(\mathbb{E}_{\pi}[\mathbf{U}]) \le 0. \tag{8}$$

With (8), the task is now to find a stochastic policy π^{i*} such that the following holds for all π^i within a neighborhood of π^{i*} :

$$J^{i}(\pi^{i*}, \pi^{\neg i*}) \le J^{i}(\pi^{i}, \pi^{\neg i*}), \quad \pi^{i} \in \Pi(\pi^{\neg i*}),$$
(9)

where $\Pi(\pi^{\neg i*}) \coloneqq \{\pi^i | h^i(\mathbb{E}_{\pi}[\mathbf{U}]) \le 0\}.$

III. ALGORITHMS FOR SOLVING MAXIMUM ENTROPY DYNAMIC GAMES

In this section, we now derive algorithms for solving the MaxEnt Dynamics Games introduced in the previous section.

A. Unconstrained Dynamic Programming

We first consider the unconstrained case for simplicity. Defining the value function for agent *i* given the policies of other agents π^{-i} to be

$$V^{i}(\boldsymbol{x}) = \inf_{\pi^{i}} \left\{ J^{i}(\boldsymbol{x}, \pi^{i}, \pi^{\neg i}) \right\},$$
(10)

applying dynamic programming results in Bellman's equation:

$$V^{i}(\boldsymbol{x}) = \inf_{\pi^{i}} \Big\{ \mathbb{E}_{\pi} \Big[l^{i}(\boldsymbol{x}, \boldsymbol{u}) + V^{i'}(f(\boldsymbol{x}, \boldsymbol{u}) \Big] - \alpha H[\pi^{i}(\cdot | \boldsymbol{x})] \Big\}.$$
(11)

In (11) and below we omit the time index t for nonterminal times for simplicity and use $V^{i'}(f(x, u))$ to denote the value function for agent i at the next timestep.

It is well known that the optimal policy π^{i*} that solves the infimum in (11) is the Gibbs distribution [34, 15]

$$\pi^{i*}(u^i|\boldsymbol{x}) = \frac{1}{Z^i} \exp\left(-\frac{1}{\alpha} \mathbb{E}_{\pi^{\neg i}} \left[V^{i'}(f(\boldsymbol{x}, \boldsymbol{u})) + l^i(\boldsymbol{x}, \boldsymbol{u}) \right] \right),$$
(12)

where Z^i denotes the partition function

$$Z^{i} \coloneqq \int \exp\left(-\frac{1}{\alpha} \mathbb{E}_{\pi^{-i}} \left[V^{i'}(f(\boldsymbol{x}, \boldsymbol{u})) + l^{i}(\boldsymbol{x}, \boldsymbol{u}) \right] \right) \, \mathrm{d}u^{i}.$$
(13)

Although we have obtained a closed-form expression for π^{i*} in (12), it is defined in terms of $\pi^{\neg i}$ which is unknown. Hence, unlike the optimal control case, we must solve a system of equations for each agent to find π^{i*} and $\pi^{\neg i*}$.

B. Unconstrained MELQGames

In this section, we propose using DDP to solve for Nash Equilibria of MaxEnt dynamic games and derive the unconstrained MELQGames algorithm similar to iLQGames in [12]. For notational simplicity, we will drop the second-order approximation of the dynamics as in iterative Linear Quadratic Regulator (iLQR) in our description of Differential Dynamic Programming (DDP). The dropped second-order dynamics terms can easily be added back in the derivations below. We refer readers to [19, 17, 12] for a detailed overview of the vanilla DDP, iLQR and iLQGames algorithms.

The DDP algorithm consists of a forward pass and a backward pass. The forward pass simulates the dynamics forward in time obtaining a set of nominal state and control trajectories $(\bar{x}_{0:T}, \bar{u}_{0:T-1})$, while the backward pass solves the Bellman equation with a 2nd order approximation of the costs and dynamics equations around the nominal trajectories. The boundary conditions for the value functions V^i for each agent are obtained by performing a 2nd order Taylor expansion of the terminal costs Φ^i :

$$V^i_{\boldsymbol{x}\boldsymbol{x},T} = \Phi^i_{\boldsymbol{x}\boldsymbol{x}}, \quad V^i_{\boldsymbol{x},T} = \Phi^i_{\boldsymbol{x}}, \quad V_T = = \Phi^i,$$
(14)

where we follow the notation in DDP literature by denoting partial derivatives via subscripts.

Additionally, we perform a quadratic approximation of the costs and linear approximation of the dynamics around a

nominal trajectory $(\bar{\mathbf{X}}, \bar{\mathbf{U}})$, where $\mathbf{X} = [\boldsymbol{x}_0, \dots, \boldsymbol{x}_T]$ denotes a trajectory of states

$$l^{i}(\boldsymbol{x},\boldsymbol{u}) \approx l^{i} + \begin{bmatrix} l_{\boldsymbol{x}}^{i} \\ l_{\boldsymbol{u}}^{i} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} l_{\boldsymbol{x}\boldsymbol{x}}^{i} & l_{\boldsymbol{x}\boldsymbol{u}}^{i} \\ l_{\boldsymbol{u}\boldsymbol{x}}^{i} & l_{\boldsymbol{u}\boldsymbol{u}}^{i} \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix},$$

$$f(\boldsymbol{x},\boldsymbol{u}) \approx f + f_{\boldsymbol{x}} \delta \boldsymbol{x} + f_{\boldsymbol{u}} \delta \boldsymbol{u}, \qquad (15)$$

where we have defined $\delta x = x - \bar{x}$ and $\delta u = u - \bar{u}$.

Let Q^i denote the terms inside the expectation in Bellman's equation (11)

$$Q^{i}(\boldsymbol{x},\boldsymbol{u}) \coloneqq l^{i}(\boldsymbol{x},\boldsymbol{u}) + V^{i'}(f(\boldsymbol{x},\boldsymbol{u})).$$
(16)

Then, Q^i for the approximated system is quadratic:

$$Q^{i}(\boldsymbol{x}, \boldsymbol{u}) = \bar{V}^{i} + \delta Q^{i}, \qquad (17)$$

$$\delta Q^{i} \coloneqq \begin{bmatrix} Q_{\boldsymbol{x}}^{i} \\ Q_{\boldsymbol{u}}^{i} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} Q_{\boldsymbol{x}\boldsymbol{x}}^{i} & Q_{\boldsymbol{x}\boldsymbol{u}}^{i} \\ Q_{\boldsymbol{u}\boldsymbol{x}}^{i} & Q_{\boldsymbol{u}\boldsymbol{u}}^{i} \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}.$$
(18)

with $\bar{V}^i := l^i(\bar{x}, \bar{u}) + V^{i'}(f(x, u))$. The derivation and full expressions for the partial derivatives of Q^i are included in Appendix A. Performing a change of variables $\delta u = u - \bar{u}$ on π^i and taking out terms that not functions of δu^i , the infimum in the Bellman equation (11) simplifies to

$$\inf_{\pi^{i}} \int \pi^{i} (\delta u^{i}) \left\{ \left(Q_{u^{i}}^{i} + Q_{u^{i}x}^{i} \delta \boldsymbol{x} + Q_{u^{i}u^{\neg i}}^{i} \mathbb{E}_{\pi^{\neg i}} \left[\delta u^{\neg i} \right] \right)^{\mathrm{T}} \delta u^{i} + \frac{1}{2} \delta u^{i^{\mathrm{T}}} Q_{u^{i}u^{i}}^{i} \delta u^{i} + \alpha \log \pi^{i} (\delta u^{i}) \right\} d\delta u^{i},$$
(19)

Since the expression above is now quadratic in δu^i for each agent *i*, the coupled systems of equations can be solved as shown in the following lemma.

Lemma 1 (Optimal MELQGames Policy). The optimal policy π^{i*} (in the Nash equilibria sense) which solves (19) for each agent *i* has the form

$$\pi^{i*} = Z^{i^{-1}} \exp\left(-\frac{\alpha}{2} (\delta u^i - \delta u^{i*})^{\mathrm{T}} Q^i_{u^i u^i} (\delta u^i - \delta u^{i*})\right),$$
(20)

where δu^{i*} is the solution to the following system of equations

$$0 = \hat{Q}_{\boldsymbol{u}\boldsymbol{u}}\delta\boldsymbol{u}^* + Q_{\boldsymbol{u}\boldsymbol{x}}\delta\boldsymbol{x} + Q_{\boldsymbol{u}}, \quad \delta\boldsymbol{u}^* = \boldsymbol{k} + \mathbf{K}\delta\boldsymbol{x}, \quad (21)$$

and the matrix $\hat{Q}_{uu} \in \mathbb{R}^{n_u \times n_u}$ is obtained by vertically stacking the row vectors $Q^i_{u^i u}$ for each $i \in \{1, \ldots, N\}$:

$$(\hat{Q}_{\boldsymbol{u}\boldsymbol{u}})^{\mathrm{T}} \coloneqq \begin{bmatrix} (Q_{u^{1}\boldsymbol{u}}^{1})^{\mathrm{T}} & \dots & (Q_{u^{N}\boldsymbol{u}}^{N})^{\mathrm{T}} \end{bmatrix}^{\mathrm{T}}.$$
 (22)

We defer the proof of Lemma 1 to the appendix in Appendix B. Note that the optimal policy π^{i*} is Gaussian distribution with mean δu^{i*} and covariance matrix $\Sigma^i = \alpha (Q_{u^i u^i})^{-1}$, where δu^{i*} has the same expression as in the iLQGames case. Additionally, as $\alpha \to 0$, π^{i*} converges to the delta distribution centered on δu^{i*} . Information from the value function is incorporated in the policy in both the mean and the covariance as the covariance is higher along eigenvectors of $Q_{u^i u^i}^i$ which have small eigenvalues. Fig. 2 illustrates sampled trajectories from the optimal MaxEnt Nash equilibria for a single agent with triple integrator dynamics. The trajectories have a tighter distribution in the middle where the cost function,



Fig. 2: Sampled trajectories from the stochastic policy obtained from solving the MaxEnt Nash equilibria (20) for a single agent. The covariance of the MaxEnt policy reflects the curvature of the cost landscape with a tighter distribution of position in the middle than at the ends.

shown in the contour plot, has a much higher curvature. Also, the terminal quadratic cost is higher along the direction of the x-axis, causing the trajectories near the target (yellow star) to have a lower variance along the x-axis.

We next substitute equation (20) from Lemma 1 into the Bellman equation (19) to derive the update equations for the value function, shown in the following lemma.

Lemma 2 (MELQGames Value Function Update). Suppose the infimum in the Bellman equation (11) is solved with policy π^{i*} with mean $\delta u = k + K \delta x$ according to Lemma 1. Then, the value function using for agent *i* has the form

$$V^{i}(\boldsymbol{x}) = \mathbb{E}_{\delta \boldsymbol{u} \sim \tilde{\pi}} \left[l^{i}(\boldsymbol{x}, \boldsymbol{u}) + V^{i'}(f(\boldsymbol{x}, \boldsymbol{u}) \right] - \alpha H[\tilde{\pi}^{i}], \quad (23)$$

$$= \left(V^{i} + V_{H}^{i}\right) + V_{\boldsymbol{x}}^{i^{\mathrm{T}}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}V_{\boldsymbol{x}\boldsymbol{x}}^{i}\delta\boldsymbol{x}, \qquad (24)$$

where the terms V^i , V^i_H , V^i_x and V^i_{xx} have the form

$$V^{i} = \bar{V}^{i} + Q^{i}_{u}\boldsymbol{k} + \frac{1}{2}\boldsymbol{k}^{\mathrm{T}}Q^{i}_{\boldsymbol{u}\boldsymbol{u}}\boldsymbol{k}, \qquad (25)$$

$$V_{H}^{i} = \frac{\alpha}{2} \Big(\log |Q_{u^{i}u^{i}}^{i}| - n_{u^{i}} \log(2\pi\alpha) \Big) + \alpha \sum_{j=1, j\neq i}^{N} \operatorname{tr} \Big[(Q_{u^{j}, u^{j}}^{j})^{-1} Q_{u^{j}, u^{j}}^{i} \Big],$$
(26)

$$V_{\boldsymbol{x}}^{i} = Q_{\boldsymbol{x}}^{i} + \mathbf{K}^{\mathrm{T}} Q_{\boldsymbol{u}\boldsymbol{u}}^{i} \boldsymbol{k} + \mathbf{K}^{\mathrm{T}} Q_{\boldsymbol{u}\boldsymbol{u}}^{i} + Q_{\boldsymbol{x}\boldsymbol{u}}^{i} \boldsymbol{k}, \qquad (27)$$

$$V_{\boldsymbol{x}\boldsymbol{x}}^{i} = Q_{\boldsymbol{x}\boldsymbol{x}}^{i} + \mathbf{K}^{\mathrm{T}}Q_{\boldsymbol{u}\boldsymbol{u}}^{i}\mathbf{K} + \mathbf{K}^{\mathrm{T}}Q_{\boldsymbol{u}\boldsymbol{x}}^{i} + Q_{\boldsymbol{x}\boldsymbol{u}}^{i}\mathbf{K}.$$
 (28)

Again, we defer the proof of Lemma 2 to the appendix in Appendix C.

Note that the update rules for V^i, V^i_x and V^i_{xx} are exactly the same as in iLQGames [12], with the only difference being the addition of the V^i_H term resulting from the maximum entropy term. This term only depends on $Q^i_{u^i u^i}$ and approaches 0 as $\alpha \to 0$. Since we take $Q^i_{u^i u^i}$ to be constant during each iteration of MELQGames, V^i_H is not a function of x. Consequently, the backward pass of MELQGames can be performed by additionally computing V^i_H and Σ^i in the backward pass of iLQGames. Also, note that unlike the optimal control case where we can simplify the value function update (25)–(28) by substituting the definition of k and K to cancel out terms, this does not hold in the dynamic games case. This is because

 Q_{uu}^i (Hessian of Q^i) does not equal to \hat{Q}_{uu} (stacked rows of $Q_{u^i u}^i$) in general.

C. Constrained MELQGames via Augmented Lagrangian

To incorporate constraints, we use the augmented Lagrangian framework similar to the approach in ALGames [4], except we choose to use the DDP style unconstrained optimizer as opposed to Newton's method. Let us denote by \mathcal{L}^i the augmented Lagrangian for agent *i*. Then,

$$\mathcal{L}_{A}^{i}(\pi,\lambda;\rho) = J^{i}(\pi) + \frac{\rho}{2} \sum_{j} \max\left(0, (h^{i})_{j}(\mathbb{E}[\mathbf{U}]) + \frac{\lambda_{j}^{i}}{\rho}\right)^{2}.$$
(29)

where $\rho > 0$ denotes the penalty parameter and $\lambda_j^i > 0$ denotes the Lagrange multiplier corresponding to the *j*th constraint $(h^i)_j$ for agent *i*. Since each agent may in general have a different number of constraints, the Lagrange multipliers are specific to each agent's constraints.

If the optimal Lagrange multipliers λ_j^{i*} were known, then for a sufficiently large value of ρ the minimizer of the augmented Lagrangian \mathcal{L}_A^i would be a local minimum of the constrained problem [21, 25]. To update the dual parameters λ_j^i , we perform the following dual-ascent step

$$(\lambda_j^i)^+ \leftarrow \max(0, \lambda_j^i + \rho(h^i)_j(\mathbb{E}[\mathbf{U}]).$$
 (30)

where we use the + to denote the new λ . Intuitively, the above update can be seen as an approximation to the optimal Lagrange multipliers λ_j^{i*} . Let \mathcal{I}_A denote the set of active constraints, and suppose that the Lagrange multipliers for all inactive constraints are zero such that

$$j \notin \mathcal{I}_A \implies (h^i)_j + \frac{\lambda_j^i}{\rho} < 0.$$
 (31)

Then, if π is a minimizer of \mathcal{L}_A^i then

$$0 = \nabla_{u^i} \mathcal{L}_A^i = \nabla_{u^i} J^i + \sum_{j \in \mathcal{I}_A} \left[\rho(h^i)_j + \lambda_j^i \right] \nabla_{u^i}(h^i)_j, \quad (32)$$

and hence $(\lambda_j^i)^+ = \rho(h^i)_j + \lambda_j^i$ is a good approximation to the Lagrange multiplier of the original constrained problem. To maintain dual feasibility of the Lagrange multipliers, we project them by taking the positive part.

Alternatively, the Augmented Lagrangian method can be viewed as projected steepest ascent applied to the dual problem of (5) where ρ plays the role of the step size in (30). Indeed, convergence rate results for Augmented Lagrangian methods show that under suitable conditions (ex. Linear Independent Constraint Qualification (LICQ) and Second-order Sufficient Optimality Condition (SOSC) are satisfied at the local minima), the Lagrange multipliers converge linearly [21, 25] with rate

$$\frac{\|(\lambda_j^i)^+ - \lambda_j^{i*}\|}{\|\lambda_j^i - \lambda_j^{i*}\|} \le \frac{M}{\rho},\tag{33}$$

for some constant M. Hence, larger values of ρ should theoretically result in faster convergence of the Lagrange multipliers. However, values of ρ that are too large may make the unconstrained problem ill-conditioned and slow down convergence [21, 25]. Hence, Augmented Lagrangian methods usually advocate for increasing ρ by some factor $\gamma > 0$ if a sufficient decrease condition on some feasibility metric such as

$$\mathcal{V} \coloneqq \left\| \min\left\{ -(h^i)_j, \frac{\lambda_j^i}{\rho} \right\} \right\|,\tag{34}$$

has not been met [2].

The constrained MELQGames algorithm consists of an iterative scheme where we alternate approximately solving the (unconstrained) Nash equilibria for cost functions \mathcal{L}_A^i and updating the Lagrange multipliers λ_j^i with (30). If the metric \mathcal{V} has not decreased by more than some $\tau \in (0, 1)$ since the last dual update, we increase ρ .

We now show in the following lemma that if the augmented Lagrangian MELQGames converges, the solution is a local GNEP.

Lemma 3. Suppose that the augmented Lagrangian *MELQGames converges to a tuple* $(\pi^*, \lambda_j^i, \rho)$ for some $\rho > 0$ which satisfies dual feasibility of λ , *LICQ and SOSC. Then,* π^* is a Generalized Nash Equilibrium (GNE) for (5).

We defer the proof of Lemma 3 to Appendix D.

IV. EXTENSIONS TO MODELING MULTIMODALITY

The vast majority of existing works that investigate game theoretic multi-agent interactions consider unimodal behaviors [4, 11, 12, 20, 31, 36]. However, this may not be enough when the true value function is not unimodal, a common situation that can arise from nonconvex dynamics, cost functions or constraints. In this section, we first outline a computational challenge for tackling the problem of multimodality in dynamic games. We then propose a computationally efficient multimodal extension to MELQGames by reformulating the problem as an incomplete information games with information asymmetry and show how our formulation can additionally be used for Bayesian inference of the latent mode.

In this paper, we focus our attention on multi-agent interactions where multimodality is induced by uncertainty about multimodal behavior of non-ego agents despite having unimodal best-responses. We leave extensions to situations where the ego agent has a multimodal best-response by considering approaches such as compositionality [30] as future work.

One of the biggest challenges for handling multimodality with multiple agents is that computational costs can quickly become unfeasible due to the combinatorial explosion when accounting for the interactions between all agents. To see this, consider a game with N agents where each agent has a policy with A different modes. Evaluating the Nash equilibrium requires evaluating the expectation over all policies for all timesteps $t = 0, \ldots, T - 1$, resulting in a total of $O(T A^N)$ evaluations of the cost function and dynamics. This combinatorial explosion is computationally intractable for any kind of realtime planning.

A. Multimodal Dynamic Games via Information Asymmetry

We now propose a method for tackling the challenges mentioned above. To start, we refine our problem setup and identify one agent as the "ego" agent. Given that many applications for planning have the goal of controlling an agent, this is not an unreasonable assumption. From hereon after, we denote the ego agent with index 1 and use $\neg 1$ to refer to non-ego agents.

Consider a set of A different local (generalized) Nash equilibria with $\{(\mathbf{X}^a, \mathbf{U}^a)\}_{a=1}^A$ and let $a \in \{1, \ldots, A\}$ be a discrete latent random variable that determines the nominal trajectory around which the cost functions and dynamics are approximated. We now consider an extension of the maximum entropy dynamic game (8), where all agents except for the ego agent have full knowledge of the value of a and are playing the corresponding optimal policy π^a . However, the non-ego agents incorrectly believes that the ego agent knows what the true mode is. This is now a **dynamic game with incomplete information** since the ego agent only has a *belief* of what the true game is but does not know what the true dynamics nor cost are. However, since the non-ego agents know what the true mode is, their actions act as a *signal* for the ego agent and allow the ego-agent to update its belief using this information.

In the theory of games with incomplete information, a *state* of the world ω fully defines a possibility of the true cost function and dynamics of the game [18]. Uncertainty over the true state of the world ω_* implies uncertainty over what game is being played. Let Y be a finite set containing all possible states of the worlds and $p_i: Y \to \Delta(Y)$ be the *belief* of each agent mapping each state of the world $\omega \in Y$ to a probability distribution over Y.

For our problem, let $\omega_{a,s}$ and $\omega_{a,m}$ for $a = 1, \ldots, A$ denote different states of the world such that when $\omega_* \in \{\omega_{a,s}, \omega_{a,m}\}$, the true cost and dynamics that correspond to the approximated costs and dynamics around the local (generalized) Nash equilibrium (\bar{x}^a, \bar{u}^a). The "s" and "m" here can be taken to mean "single" and "multi". We now define our set Y as

$$Y \coloneqq \{\omega_{1,s}, \dots, \omega_{A,s}, \omega_{1,m}, \dots, \omega_{A,m}\}.$$
 (35)

Let 1 denote the indicator function and p_a denote the ego agent's prior over the modes a with support on the set $\{\omega_{a,m}\}_{a=1}^{A}$. We define the beliefs of the ego agent p_1 and non-ego agents $p_{\neg 1}$ as

$$p_1(\omega) = \begin{cases} \mathbb{1}_{\omega_{a,s}}(\omega), & \omega = \omega_{a,s}, \ a \in \{1, \dots, A\}\\ p_a(\omega), & \omega = \omega_{a,m}, \ a \in \{1, \dots, A\} \end{cases},$$
(36)

$$p_{\neg 1}(\omega) = \begin{cases} \mathbb{1}_{\omega_{a,s}}(\omega), & \omega = \omega_{a,s}, \ a \in \{1, \dots, A\} \\ \mathbb{1}_{\omega_{a,s}}(\omega), & \omega = \omega_{a,m}, \ a \in \{1, \dots, A\} \end{cases},$$
(37)

To give intuition to the above game setup, note that each mode a is associated with two world states: $\omega_{a,s}$ and $\omega_{a,m}$. When the true state ω_* is at $\omega_{a,s}$, all agents correctly believe that the state is $\omega_{a,u}$. However, when $\omega_* = \omega_{a,m}$, the non-ego agents incorrectly believe that the true state is $\omega_{a,s}$ while the ego agent correctly believes that the true state lies is in the set $\{\omega_{a,m}\}$ but is unsure which is the correct one.

With the above setup, we can use tools from game theory to gain intuition into the *belief-structure* of this game, namely what each agent believes is true, what each agent believes other agents believe is true and so on [18].

Definition 1 (Belief Operator). Let B_i denote the belief operator such that for any event $A \subseteq Y$

$$B_i A \coloneqq \{ \omega \in Y : p_i(A|\omega) = 1 \}, \tag{38}$$

and the conditional belief $p_i(A|\omega)$ is defined as

$$p_i(A|\omega) \coloneqq \mathbb{E}_{p_i(\omega)}[\mathbb{1}_A]. \tag{39}$$

In other words, $\omega \in B_i A$ means that at ω , event A obtains according to agent *i*'s belief in the sense that at least one of $\omega \in A$ corresponds to the true state. We next define the concept of common belief [18]:

Definition 2 (Common Belief). Let $A \subseteq Y$ be an event and $\omega \in Y$. The event A is common belief at state ω if every agent believes that A obtains, every agent believes that every agent believes that A obtains, and so on. In other words, for every finite sequence i_1, \ldots, i_L of agents:

$$\omega \in B_{i_1} \dots B_{i_L} A,\tag{40}$$

In particular, the event Y is common belief among the agents, i.e., each agent knows that the true world state ω must be an element of Y.

Using Definition 1 and Definition 2, we now precisely define the key property of the dynamic game — it is common knowledge that the non-ego agents believes that the ego agent knows the mode.

Lemma 4. Let $E = \{\omega_{a,s}\}_{a=1}^{A}$ denote the event that the ego agent knows the correct mode with probability 1. Then, for all $i \neq 1$, $\omega \in B_i E$, i.e. all non-ego agents believe that the ego agent knows the mode perfectly. Furthermore, it is common belief that $B_i E_a$ for all $i \neq 1$. In particular, the ego agent (correctly) believes that all non-ego agents (incorrectly) believes that it knows the correct mode.

We delegate the proof of Lemma 4 to Appendix E. Furthermore, we refer the interested reader to Maschler et al. [18] for more information on belief spaces and games with incomplete information. The above lemma shows how this choice of information asymmetry simplifies computation. Namely, since each non-ego agent believes that the ego-agent knows the mode perfectly, each non-ego agent's control will exactly be the Nash-equilibria control in the unimodal case.

Finally, we introduce the concept of a (local) *Bayesian Nash equilibrium*. This is a generalization of the classical (generalized) Nash equilibrium to the incomplete information case, where the *expectation* of an agent's cost with respect to its own belief is used [18]:

Definition 3 (Bayesian Nash equilibrium). A policy π^* is a Bayesian Nash equilibrium if, for all agents *i*,

$$\mathbb{E}_{p_i} \left[J^i(\pi^{i*}, \pi^{\neg i*}) \right] \le \mathbb{E}_{p_i} \left[J^i(\pi^i, \pi^{\neg i*}) \right], \quad \pi^i \in \Pi(\pi^{\neg i*}),$$
(41)



Fig. 3: Graphical model of the POMDP for the ego agent with corresponding value function (42). The state x_t and non-ego agent controls $u_t^{\neg 1}$ are observed variables, while the computed set of non-ego policies $\pi_t^{\neg 1,a}$ is dependent on the partially observable discrete latent mode a. The latent mode a also affects the true cost function $l_t^{1,a}$ and the dynamics for the next state x_{t+1} .

where the expectation is taken with respect to agent *i*'s belief of the world state $\omega_* \in Y$.

As a result, when $\omega_* = \omega_{a,s}$, since all agents know the mode a, the optimal π^* simply corresponds to the optimal policy π^a for the generalized Nash equilibrium found previously. When $\omega_* = \omega_{a,m}$, the non-ego agents believe that the true state is still $\omega_{a,s}$ and hence the optimal policy is still $\pi^{\neg i,a}$. On the other hand, for $\omega_* = \omega_{a,m}$, although the ego-agent is uncertain about the true value of a, it knows via Lemma 4 that the nonego agents will play $\pi^{\neg i,a}$. Furthermore, the controls of the non-ego agents $u^{\neg i}$ give information about what the true value of a is. Hence, the problem of computing the Bayesian Nash equilibrium π^{i*} reduces to the problem of solving a Partially Observable Markov Decision Process (POMDP) where the belief space is over the latent variable a and the observations oare the observed controls of the non-ego agents. The conditional dependencies for the POMDP are shown in the graphical model in Fig. 3. By using the standard belief-space approach to solve the POMDP, the Bellman equation takes the form

$$V^{1}(\boldsymbol{x}, \boldsymbol{b}_{t}) = \inf_{\pi^{1}} \mathbb{E}_{\pi} \left[\mathbb{E}_{\boldsymbol{b}_{t}} l^{1, a}(\boldsymbol{x}, \boldsymbol{u}) + \mathbb{E}_{\pi^{-1}} V^{1'}(f(\boldsymbol{x}, \boldsymbol{u}), \boldsymbol{b}_{t+1}) \right],$$
(42)

where b_{t+1} refers to the updated belief for the next timestep via Bayesian filtering:

$$\boldsymbol{b}_{t+1}(a) \propto \pi^{-1}(\boldsymbol{u}^{-1}|a) \, \boldsymbol{b}_t(a). \tag{43}$$

While solving for general POMDPs is computationally untractable [23], there exist computationally efficient methods for solving POMDPs over discrete latent spaces by using a DDPbased approach [24]. However, to simplify the implementation and presentation of this work, we choose to bypass the problem of solving the full POMDP by making the following assumption. Assumption 1. The ego agent will be informed of the true mode a after one timestep at t = 1.

With this assumption, the value functions for $t \ge 1$ are known and correspond to the value functions $V^{1,a}$ for each mode a of the found local Nash equilibrium. Hence, the Bellman equation at t = 0 reads

$$V^{1} = \inf_{\pi^{1}} \mathbb{E}_{\boldsymbol{b}_{0},\pi} \left[l^{1,a}(\boldsymbol{x},\boldsymbol{u}) + V^{1,a'}(f(\boldsymbol{x},\boldsymbol{u})) \right] - \alpha H[\pi^{1}],$$

$$= \inf_{\pi^{1}} \left\{ \mathbb{E}_{\pi^{i}} \left[\frac{1}{2} u^{1^{T}} \tilde{Q}_{u^{1}u^{1}}^{1} u^{1} + \left(\tilde{Q}_{u^{1}}^{1} + \tilde{Q}_{u^{1}}^{1} \boldsymbol{x} \right)^{T} u^{i} \right] + c - \alpha H[\pi^{1}] \right\},$$
(44)

where $b_0 \in \Delta(Y)$ denotes the ego agent's prior belief of the true mode, c encapsulates all terms that are constant with respect to u^1 , and

$$\tilde{Q}_{u^1u^1}^1 \coloneqq \mathbb{E}_{\boldsymbol{b}_0} \Big[Q_{u^1u^1}^{1,a} \Big], \qquad \tilde{Q}_{u^1\boldsymbol{x}}^1 \coloneqq \mathbb{E}_{\boldsymbol{b}_0} \Big[Q_{u^1\boldsymbol{x}}^{1,a} \Big], \quad (45)$$

$$\tilde{Q}_{u^{1}}^{1} \coloneqq \mathbb{E}_{\boldsymbol{b}_{0}} \Big[Q_{u^{1}}^{1,a} - Q_{u^{1}u^{1}}^{1,a} \bar{u}^{1,a} - Q_{u^{1}\boldsymbol{x}}^{1,a} \bar{\boldsymbol{x}}^{a} \Big].$$
(46)

Consequently, the optimal policy π^1 corresponds to the Gibbs distribution and is solved in the same way as in MELQGames Lemma 1.

Choices of the prior distribution: There are many valid choices for the prior distribution on the modes b_0 . If no prior information on the non-ego agents is known, one choice is to choose b_0 as the maximum entropy distribution that minimizes the sum of the value functions for each agent:

$$\boldsymbol{b}_0(a) = \inf_p \mathbb{E}_{a \sim p} \left[\sum_{i=1}^N V^{i,a}(\boldsymbol{x}_0) \right] - \alpha H[p].$$
(47)

With this prior, modes that result in lower costs for all agents will have higher probabilities over modes that have high costs for all agents. However, in the context of planning for an ego agent, this may favor local minima which are bad for the ego agent when better alternatives exist.

Alternatively, we can bias the selection of modes to ones which are favorable to the ego-agent by choosing b_0 as the maximum entropy distribution that minimizes the sum of the ego agent's value functions

$$\boldsymbol{b}_0(a) = p(a|\boldsymbol{x}_0) = \inf_p \mathbb{E}_{a \sim p} \left[V^{1,a}(\boldsymbol{x}_0) \right] - \alpha H[p].$$
(48)

B. Latent Mode Bayesian Inference via MaxEnt

In a MPC context, we can improve on our choices of the prior since the controls of the non-ego agents from previous timesteps act as signals for the true mode a. By observing them, the ego agent can form a more accurate estimate of b_0 . For simplicity, we assume that the observations of non-ego agent's controls are noiseless, though this can be generalized to the noisy case in a Bayesian fashion by incorporating an observation model. Suppose that the controls $\hat{u}^{\neg 1} \coloneqq u_{-k:-1}^{\neg 1}$ and optimal policies $\{\pi^{\neg 1,a}\}_{a=1}^{A}$ for all non-ego agents from

Algorithm 1 MMELQGames

1: Initialize $\{\pi\}_{a=1}^{\tilde{A}}$ with random policy.

- 2: Initialize buffer of non-ego control observations \mathcal{U} .
- 3: repeat
- $\{(\mathbf{X}^a, \pi^a)\}_{a=1}^{\tilde{A}} \leftarrow \text{Find Modes}$ $\pi^1 \leftarrow \text{Solve Ego Agent POMDP}$ 4:
- 5:
- Execute π^1 by sampling or taking the mean control 6:
- Observe non-ego controls $u^{\neg 1}$ and add to \mathcal{U} 7:
- Shift π by one timestep for warmstarting next iteration 8:

the past k timesteps are known. Then, the posterior on those previous controls can be used for our estimate of b_0 :

$$\boldsymbol{b}_0(a) \coloneqq p(a|\hat{\boldsymbol{u}}^{-1}) \propto p(a) \prod_{t=-k}^{-1} p(\boldsymbol{u}_t^{-1}|a), \qquad (49)$$

where p(a) is chosen to be the prior (48) from the previous subsection that biases towards modes that are favorable to the ego-agent. With this choice of prior, modes which describe the controls of non-ego agents and are beneficial to the ego-agent will have high probability.

When computing the posterior in (49), it is assumed that all of the controls in \hat{u}^{-1} correspond to the same mode a. While one could simply use all of the collected observations to compute the posterior, this choice is invalid in the case that non-ego agents switch to a different mode somewhere in the middle of the collected observations. To solve this issue, we can make the less restrictive assumption that only the last kcontrols from non-ego agents correspond to the same mode.

V. GAME THEORETIC MODEL PREDICTIVE CONTROL

In this section, we summarize the MMELQGames algorithm which combines the solution to the information asymmetric game setup described in Section IV-A with the latent mode Bayesian inference in Section IV-B and propose a receding horizon game-theoretic planner that can reason about and infer multiple hypothesis in Algorithm 1, Algorithm 2, Algorithm 3 and summarized in Fig. 4. Algorithm 1 can be summarized as first solving for a set of local Nash equilibrium in Algorithm 2, then using the solutions as inputs to solve a POMDP for the ego agent with uncertainty over the mode in Algorithm 3.

Algorithm 2 starts by first finding a set of modes that each solve the local maximum entropy generalized Nash equilibrium problem (8). Moreover, to obtain the optimal policies for the non-ego agents from previous timesteps, we solve starting from the earliest timestep which we wish to condition our posterior computation (49) on, warm-starting this computation from the previous iteration's solution if available. To solve for each mode, we use the MELQGames algorithm. More specifically, we first compute the nominal state trajectories $\{\bar{x}^a\}_{a=1}^A$ as well as derivatives of the cost functions l^i and dynamics f along (\bar{x}^a, \bar{u}^a) for each mode. Then, we solve for the backward pass by computing the optimal policy (20) and updating the value function (24) for each timestep starting from the terminal time T. Finally, we update the nominal controls



Fig. 4: Schematic diagram of the proposed MPC MMELQGames algorithm presented in (1) and (3).

Algorithm 2 Find Modes

Require: Number of modes A, Number of iterations K, control estimates $\{\bar{u}^a\}_{a=1}^A$, Penalty parameter ρ , Feasibility progress $\tau > 0$, Penalty increase $\gamma > 0$, Thresholds $\epsilon > 0$ 1: Initialize $\lambda_i^i = 0, \mathcal{V}_0^a = \infty, \rho^a = 1$ 2: Compute nominal mean trajectory $\{\mathbf{X}^a\}$ using $\{\mathbf{U}^a\}$. 3: for k = 1 to K do for a = 1 to A in parallel do 4: $\pi^{i,a*}, V_x^{i,a}, V_{xx}^{i,a} \leftarrow \text{Backward pass (20), (25)-(28)}$ 5: 6: end for for a = 1 to A in parallel do 7: $(\bar{\mathbf{X}}^a, \bar{\mathbf{U}}^a) \leftarrow \text{Line Search on } \max_i \|J_{ai}^i\|$ 8: end for 9: Compute \mathcal{V}_k^i using (34) if $\max_i \|J_{u^i}^i\| \le \epsilon$ then \triangleright Unconstrained Converged 10: 11: if $\mathcal{V}_k^a \leq \epsilon$ then ▷ Constraints Converged 12: Save $(\bar{\mathbf{X}}^a, \bar{\mathbf{U}}^a, \pi^{u,a*}, V_x^{i,a}, V_{xx}^{i,a})$ and reinitialize 13: mode a by sampling from π . end if 14: if $\mathcal{V}_k^a > \tau \mathcal{V}_{k-1}^a$ then 15: ▷ Insufficient Decrease $\rho^a \leftarrow \gamma \rho^a$ 16: 17: end if end if 18: 19: end for 20: $\{\bar{\mathbf{U}}^a\}_{a=1}^A \leftarrow \text{Cluster} \{\bar{\mathbf{U}}^a\}$

Algorithm 3 Solve Ego Agent POMDP

- **Require:** Observed controls from past k timesteps $u_{-k:-1}^{\neg 1}$, Value functions $V^{i,a}$, Optimal Policies $\{\pi^{\neg 1,a}\}_{a=1}^{\tilde{A}}$ for non-ego agents
- 1: $p(a) \leftarrow \text{Compute prior with (47) or (48)}$
- 2: $b_0 \leftarrow$ Compute initial belief with (49)
- 3: $\pi^1 \leftarrow$ Solve full (42) or approximate (44) POMDP

using a line search on the maximum of the norm $\max_i ||J_{u^i}^i||$. We exploit the fact that each local Nash equilibrium can be computed independent by solving for A local generalized Nash equilibrium *in parallel*, where A is chosen to be the number of cores available for parallel processing on the CPU. This iteration of backward and forward pass continues for each mode until this norm is smaller than some threshold ϵ , upon which it is saved into a set of converged modes before being reinitialized from the current π . After the maximum number of iterations has been reached, we perform clustering via kmeans to remove potential duplicate modes before returning the reduced set $\{(\bar{x}^a, \bar{u}^a)_{a=1}^{\tilde{A}}\}$ of unique Nash equilibrium.

Then, in Algorithm 3, we use the ego agent's value function, optimal policies for the non-ego agents $\pi^{-i,a}$ and observed controls from non-ego agents to perform Bayesian inference on the mode *a* to compute our belief prior b_0 via (49). Consequently, we use b_0 and $\pi^{-i,a}$ to solve the POMDP (42) and obtain the optimal policy $\pi^{1,*}$. This can be done by either solving the full POMDP (42) using the method from [24] or by solving the approximated one-step POMDP (44).

Finally, we either execute a random sample from $\pi^{1,*}$ or use the mean control before repeating the algorithm again in a receding horizon fashion.

VI. CONNECTIONS TO RELATED WORKS

A. Game-Theoretic Planning

There are two main approaches to solving for game-theoretic equilibria: Iterative Best Response (IBR)-based and direct methods.

IBR-based methods: IBR-based methods rely on either the Jacobi decomposition or Gauss-Seidel decomposition to decompose the problem of finding first-order stationary points into separate optimal control problems for each agent [3, 36]. These methods optimize the controls for one agent while keeping the controls for other agents fixed, repeating this process iteratively until convergence. By holding controls, IBR assumes that $\frac{\partial u_{t+1}^{-i}}{\partial u_t^i} = \frac{\partial u_{t+1}^{-i}}{\partial x_{t+1}} \frac{\partial x_{t+1}}{\partial u_t^i} = 0$. As a result, IBR lacks the term below key to describing inter-agent behavior arising from asymmetric inter-agent coupling terms:

$$\frac{\partial J^{i}}{\partial u_{t}^{i}} = \frac{\partial J^{i}}{\partial \boldsymbol{x}_{t+2}} \cdot \left(\frac{\partial \boldsymbol{x}_{t+2}}{\partial \boldsymbol{u}_{t+1}^{-i}} \frac{\partial \boldsymbol{u}_{t+1}^{-i}}{\partial \boldsymbol{x}_{t}} \frac{\partial \boldsymbol{x}_{t}}{\partial u_{t}^{i}} + \dots \right).$$
(50)

To remedy this, approaches from [31, 32, 33] propose the SE-IBR algorithm which extends IBR by accounting for some of missing information using the sensitivity nformation from the Lagrange multipliers of active constraints. However, their approach relies on a particular structure of the cost function. Additionally, the convergence of IBR-based methods is not well understood [10] and can potentially require many trajectory optimization iterations before convergence.

Direct methods: Direct methods solve for the coupled system of equations resulting from the first-order stationary conditions directly. Examples include methods based on Newton's method [7, 8, 4] and DDP [12, 28]. The proposed

MMELQGames algorithm falls under the category of DDPbased methods as it solves a quadratic approximation of the dynamic game during each timestep. However, unlike previous methods, our method focuses on the constrained MaxEnt version of the game and also incoporates multimodality via a novel formulation of an incomplete information with information asymmetry.

B. Multimodal Planning

There are relatively few works that explore the concept of multimodality within the context of dynamic games. This idea is explored in [16], but requires the multimodality to be known a priori and explicitly encoded into the model. One series of works [13, 14, 27] approach this problem by training a generative model offline conditioned on a discrete latent variable to learn the trajectories of other agents. While this approach does take multimodality into account, it is not done in a game-theoretic context.

The work in this paper is most similar to [24], which also solves a POMDP with a discrete latent variable. However, in their work, the discrete latent variable must be prespecified while the multimodality is discovered in this work. Also, their work only looks at optimality for a single agent while in our work we focus on solving for generalized Nash equilibria, with the POMDP being only one piece of the algorithm.

C. Bayesian Inference and Laplace Approximation

The maximum entropy term transforms the difficulty of each agent's optimization problem into that of sampling from the optimal π^i and computing the normalization term Z^i , a challenge shared by many techniques for Bayesian inference. Some popular techniques for tackling Bayesian inference include Markov chain Monte Carlo (MCMC) and Variational Inference (VI) [1, 26], though these methods are usually much slower and are computationally intractable for planning on real time. We do note however a series of works [22, 35] which successfully apply variational inference via sampling-based optimization. Laplace Approximation [5] is an approximation framework for Bayesian inference by finding a Gaussian approximation to a probability density function MELQGames can be viewed as applying the Laplace Approximation to π^i_t at each timestep t.

VII. EXPERIMENTS

In this section, we compare the proposed MMELQGames algorithm against the IBR and MELQGames algorithms on two examples which illustrate the capabilities of MMELQGames in recovering a rich set of multi-agent interactions. The reader is encouraged to view the videos included in the supplementary material which showcase the behavior of MMELQGames compared to MELQGames and IBR.

A. Multi-agent collision avoidance

We begin with a simple collision avoidance game between agents with unicycle dynamics to highlight the multimodality of our algorithm. Each agent's objective function composes



Fig. 5: Comparison of realizations of the stochastic policy obtained from MELQGames (left) and MMELQGames (right). For MMELQGames, the mode a is sampled using the prior from (47).



Fig. 6: MMELQGames for the three agents case, where the algorithm is able to find both Nash equilibria. The **left** plot shows the planned mean policies for each mode. The square denotes the initial position for each agent, with lines indicating that positions are from the same timestep. The **right** plot shows realizations of the stochastic policy. Again, the prior from (47) is used for sampling the mode a.

of a quadratic task and a soft cost for colliding with other agents. As there is no ego agent in this example, we assume that the ego agent knows the mode *a* from the start and choose the prior for *a* to be of the form (47). The results are shown in Fig. 5. While Mehr et al. [20] claims that the stochastic policy from MELQGames is able to result in multimodal behaviors, our results show that this is not the case. Since the mean control from MELQGames is equivalent to the linear feedback controller from iLQR, the resulting behavior should also result in a "tube" around the mean trajectory, which is indeed what we see on the top row of Fig. 5. On the other hand, the MMELQGames algorithm is able to properly handle the multimodality of this example and successfully find both local Nash equilibria of the game.

B. Game Theoretic Autonomous Racing

We next consider a two-agent racing scenario where the lead vehicle has index i = 1 and the rear vehicle has index i = 2. The objective of each agent is to maximize the difference between its own progress s^i and the progress of the other agent $s^{\neg i}$ (first term) under quadratic control cost (second



MMELQGames

Fig. 7: Comparison of the trajectories resulting from running IBR (top), iLQGames (middle) and MMELQGames (bottom) in MPC for the lead (blue) agent. IBR is used for the rear (green) agent in all cases. We test on the case where the rear agent tries to overtake above (left) and below (right) the lead agent. Colored circles are added every fifth timestep to indicate the position of each agent. Black circles indicating the radius of each agent denote the timestep where the agents are closest to each other.



Fig. 8: Snapshot of iLQGames (left) and MMELQGames (right) at three timesteps after the rear agent starts turning towards the bottom mode. The ego agent's planned trajectory is shown as a solid blue line, with the predicted trajectories of the non-ego agent shown as a dashed line. iLQGames incorrectly predicts that the rear agent will overtake from the top due to the algorithm being stuck in a local minima, while MMELQGames is able to successfully infer that the mode has changed. For MMELQGames, the different colors indicate the different modes.

term), subject to the following

$$\min_{u^{i}} (s^{\neg i} - s^{i}) + \sum_{t=1}^{T-1} \frac{1}{2} u_{t}^{i}{}^{T} R u_{t}^{i},$$
(51)

s.t.
$$h_{\text{track}}^{i}(\mathbf{X}, \mathbf{U}) \le 0,$$
 (52)

$$\mathbb{1}_{s^i < s^{\neg i}} h^i_{\text{collision}}(\mathbf{X}) \le 0, \tag{53}$$

$$h_{\text{velocity}}^{i}(\mathbf{X}) \le 0.$$
(54)

where each vehicle is modeled as a unicycle with linear and angular accelerations as controls.

To make the problem interesting, the lead vehicle has a lower maximum velocity than the rear vehicle to allow the rear vehicle the opportunity to overtake. To maintain the lead, the lead vehicle must block the rear vehicle from overtaking its position. Furthermore, to more closely model racing in real life, the collision constraint is asymmetric with the burden of respecting the constraint placed solely on the rear vehicle. We compare a MPC version of the IBR, iLQGames and MMELQGames algorithms for the lead vehicle by planning in a receding horizon fashion. Note that MELQGames is equivalent to iLQGames if the mean controls are used instead of sampling from the stochastic policy. In all cases, the rear agent solves for its controls using IBR. The width of the track means that the rear agent can choose to overtake the lead vehicle by going either above or below resulting in multimodality. We test both modes in our experiments, with the results for MMELQGames shown in Fig. 1 and a comparison of the algorithms shown in Fig. 7.

As expected, IBR is unable to take into account the asymmetric coupling effects between the agents. Although the solution it does find is a local GNE, it is one which is disadvantageous for the lead agent and results in the lead agent being overtaken by the rear agent in both modes. iLQGames is able to converge to the advantageous local GNE for the first mode and successfully finish the race ahead of the rear agent. However, the nature of the method means that it is only able to keep track of one of the two modes, resulting in the method failing to block the rear agent when the rear agent overtakes via the bottom. Finally, the multimodal nature of MMELQGames means that it is able to keep track of both local Nash equilibria at the same time. By using the history of the rear agent's control, the lead agent is able to infer whether rear agent is trying to overtake from the top or the bottom and then execute the corresponding control.

We also note that since MELQGames is a local method, it is prone to becoming stuck in local minima. From Fig. 8, even though the rear agent is clearly trying to overtake via the bottom mode, the warmstart for iLQGames means that it still predicts the rear agent will overtake from the top. On the other hand, the multimodality of MMELQGames means that it can keep track of multiple hypothesis at the same time and hence is able to infer the correct mode using the algorithm described in Section IV-B and respond accordingly.

VIII. CONCLUSION

In this paper, we have proposed a constrained MaxEnt dynamic game formulation and presented an algorithm that can solve for multiple modes of the corresponding GNE. We demonstrate its advantages over existing unimodal algorithms in the examples of multi-agent collision avoidance and autonomous racing.

While we have explored the use of constraints on the mean control in our formulation, it may be interesting to look at different ways of including constraints in the MaxEnt dynamic game formulation such as via chance constraints or constraints that hold almost-surely.

Finally, though our algorithm identifies multimodal behavior by exploring the state space, we do not provide any guarantees for how thorough this exploration is or whether there are additional modes which have not been discovered. Quantifying this via uncertainty quantification tools such as Gaussian Processes can provide more structured methods of discovering different modalities in multi-agent interactions.

ACKNOWLEDGMENTS

REFERENCES

- Christophe Andrieu, Nando De Freitas, Arnaud Doucet, and Michael I Jordan. An introduction to mcmc for machine learning. *Machine learning*, 50(1):5–43, 2003.
- [2] Ernesto G Birgin and José Mario Martínez. Practical augmented Lagrangian methods for constrained optimization. SIAM, 2014.
- [3] Andreas Britzelmeier, Axel Dreves, and Matthias Gerdts. Numerical solution of potential games arising in the control of cooperative automatic vehicles. In 2019 Proceedings of the Conference on Control and its Applications, pages 38–45. SIAM, 2019.
- [4] Simon Le Cleac'h, Mac Schwager, and Zachary Manchester. Algames: A fast solver for constrained dynamic games. arXiv preprint arXiv:1910.09713, 2019.
- [5] Erik Daxberger, Agustinus Kristiadi, Alexander Immer, Runa Eschenhagen, Matthias Bauer, and Philipp Hennig. Laplace redux - effortless bayesian deep learning. In A. Beygelzimer, Y. Dauphin, P. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems*, 2021. URL https://openreview.net/ forum?id=gDcaUj4Myhn.
- [6] Andrea De Martino and Daniele De Martino. An introduction to the maximum entropy approach and its application to inference problems in biology. *Heliyon*, 4 (4):e00596, 2018.
- [7] Bolei Di and Andrew Lamperski. Newton's method and differential dynamic programming for unconstrained nonlinear dynamic games. In 2019 IEEE 58th Conference on Decision and Control (CDC), pages 4073–4078. IEEE, 2019.
- [8] Bolei Di and Andrew Lamperski. First-order algorithms for constrained nonlinear dynamic games. *arXiv preprint arXiv:2001.01826*, 2020.
- [9] Benjamin Patrick Evans and Mikhail Prokopenko. A maximum entropy model of bounded rational decisionmaking with prior beliefs and market feedback. *Entropy*, 23(6):669, 2021.
- [10] Francisco Facchinei and Christian Kanzow. Generalized nash equilibrium problems. *Annals of Operations Research*, 175(1):177–211, 2010.
- [11] Alejandra Fonseca-Morales and Onésimo Hernández-Lerma. Potential differential games. *Dynamic Games and Applications*, 8(2):254–279, 2018.

- [12] David Fridovich-Keil, Ellis Ratner, Lasse Peters, Anca D Dragan, and Claire J Tomlin. Efficient iterative linearquadratic approximations for nonlinear multi-player general-sum differential games. In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 1475–1481. IEEE, 2020.
- [13] Boris Ivanovic and Marco Pavone. The trajectron: Probabilistic multi-agent trajectory modeling with dynamic spatiotemporal graphs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2375– 2384, 2019.
- [14] Boris Ivanovic, Edward Schmerling, Karen Leung, and Marco Pavone. Generative modeling of multimodal multi-human behavior. In 2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 3088–3095. IEEE, 2018.
- [15] Jeongho Kim and Insoon Yang. Hamilton-jacobi-bellman equations for maximum entropy optimal control. *arXiv* preprint arXiv:2009.13097, 2020.
- [16] Forrest Laine, David Fridovich-Keil, Chih-Yuan Chiu, and Claire Tomlin. Multi-hypothesis interactions in gametheoretic motion planning. In 2021 IEEE International Conference on Robotics and Automation (ICRA), pages 8016–8023. IEEE, 2021.
- [17] Weiwei Li and Emanuel Todorov. Iterative linear quadratic regulator design for nonlinear biological movement systems. In *ICINCO* (1), pages 222–229. Citeseer, 2004.
- [18] Michael Maschler, Eilon Solan, and Shmuel Zamir. Game theory (translated from the hebrew by ziv hellman and edited by mike borns). *Cambridge University Press, Cambridge, pp. xxvi*, 979:4, 2013.
- [19] David Mayne. A second-order gradient method for determining optimal trajectories of non-linear discretetime systems. *International Journal of Control*, 3(1): 85–95, 1966.
- [20] Negar Mehr, Mingyu Wang, and Mac Schwager. Maximum-entropy multi-agent dynamic games: Forward and inverse solutions. *arXiv preprint arXiv:2110.01027*, 2021.
- [21] Jorge Nocedal and Stephen Wright. *Numerical optimization.* Springer Science & Business Media, 2006.
- [22] Masashi Okada and Tadahiro Taniguchi. Variational inference mpc for bayesian model-based reinforcement learning. In *Conference on Robot Learning*, pages 258– 272. PMLR, 2020.
- [23] Christos H Papadimitriou and John N Tsitsiklis. The complexity of markov decision processes. *Mathematics* of operations research, 12(3):441–450, 1987.
- [24] Dicong Qiu, Yibiao Zhao, and Chris L Baker. Latent belief space motion planning under cost, dynamics, and intent uncertainty. In *Proceedings of Robotics: Science and Systems (RSS)*, 2020.
- [25] Andrzej Ruszczynski. *Nonlinear optimization*. Princeton university press, 2011.
- [26] Tim Salimans, Diederik Kingma, and Max Welling. Markov chain monte carlo and variational inference:

Bridging the gap. In *International Conference on Machine Learning*, pages 1218–1226. PMLR, 2015.

- [27] Edward Schmerling, Karen Leung, Wolf Vollprecht, and Marco Pavone. Multimodal probabilistic model-based planning for human-robot interaction. In 2018 IEEE International Conference on Robotics and Automation (ICRA), pages 3399–3406. IEEE, 2018.
- [28] Wilko Schwarting, Alyssa Pierson, Sertac Karaman, and Daniela Rus. Stochastic dynamic games in belief space. *IEEE Transactions on Robotics*, 2021.
- [29] Herbert A Simon. Theories of bounded rationality. *Decision and organization*, 1(1):161–176, 1972.
- [30] Oswin So, Ziyi Wang, and Evangelos A Theodorou. Maximum entropy differential dynamic programming. *arXiv preprint arXiv:2110.06451*, 2021.
- [31] Riccardo Spica, Eric Cristofalo, Zijian Wang, Eduardo Montijano, and Mac Schwager. A real-time game theoretic planner for autonomous two-player drone racing. *IEEE Transactions on Robotics*, 36(5):1389–1403, 2020.
- [32] Zijian Wang, Riccardo Spica, and Mac Schwager. Game theoretic motion planning for multi-robot racing. In *Distributed Autonomous Robotic Systems*, pages 225–238. Springer, 2019.
- [33] Zijian Wang, Tim Taubner, and Mac Schwager. Multiagent sensitivity enhanced iterative best response: A real-time game theoretic planner for drone racing in 3d environments. *Robotics and Autonomous Systems*, 125: 103410, 2020.
- [34] Ziyi Wang, Oswin So, Jason Gibson, Bogdan Vlahov, Manan Gandhi, Guan-Horng Liu, and Evangelos Theodorou. Variational Inference MPC using Tsallis Divergence. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi: 10.15607/RSS.2021. XVII.073.
- [35] Ziyi Wang, Oswin So, Jason Gibson, Bogdan Vlahov, Manan S Gandhi, Guan-Horng Liu, and Evangelos A Theodorou. Variational inference mpc using tsallis divergence. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi: 10.15607/RSS.2021. XVII.073. URL http://www.roboticsproceedings.org/ rss17/p073.html.
- [36] Grady Williams, Brian Goldfain, Paul Drews, James M Rehg, and Evangelos A Theodorou. Autonomous racing with autorally vehicles and differential games. *arXiv preprint arXiv:1707.04540*, 2017.
- [37] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, Anind K Dey, et al. Maximum entropy inverse reinforcement learning. 2008.

APPENDIX A

EXPRESSION FOR PARTIAL DERIVATIVES OF Q^i

Let Q^i denote the terms inside the expectation in Bellman's equation (11):

$$Q^{i}(\boldsymbol{x}, \boldsymbol{u}) \coloneqq l^{i}(\boldsymbol{x}, \boldsymbol{u}) + V^{i'}(f(\boldsymbol{x}, \boldsymbol{u})).$$
(A.1)

Suppose that the value function at the next timestep $V^{i'}$ is quadratic with form

$$V^{i'}(\boldsymbol{x}) = (\bar{V}^{i})' + (V^{i})'_{x}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}(V^{i})'_{xx}\delta\boldsymbol{x}$$
(A.2)

Using a quadratic approximation of the costs and linear approximation of the dynamics (15) then yields

$$Q^{i}(\boldsymbol{x},\boldsymbol{u}) \approx l^{i}(\bar{\boldsymbol{x}},\bar{\boldsymbol{u}}) + \begin{bmatrix} l_{x}^{i} \\ l_{u}^{i} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} l_{xx}^{i} & l_{xu}^{i} \\ l_{ux}^{i} & l_{uu}^{i} \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}$$
(A.3)

$$+\left(f(\bar{\boldsymbol{x}},\bar{\boldsymbol{u}})+f_{x}\delta\boldsymbol{x}+f_{u}\delta\boldsymbol{u}\right)^{\mathrm{T}}\left((\bar{V}^{i})'+(V^{i})'_{x}\delta\boldsymbol{x}+\frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}(V^{i})'_{xx}\delta\boldsymbol{x}\right)\left(f(\bar{\boldsymbol{x}},\bar{\boldsymbol{u}})+f_{x}\delta\boldsymbol{x}+f_{u}\delta\boldsymbol{u}\right).$$
(A.4)

Expanding and collecting terms then results in

$$Q^{i}(\boldsymbol{x}, \boldsymbol{u}) \approx \bar{V}^{i} + \delta Q^{i},$$

$$\bar{V}^{i} \coloneqq l^{i}(\bar{\boldsymbol{x}}, \bar{\boldsymbol{u}}) + (\bar{V}^{i})'(f(\boldsymbol{x}, \boldsymbol{u})),$$
(A.5)

$$\delta Q^{i} \coloneqq \begin{bmatrix} Q_{x}^{i} \\ Q_{u}^{i} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix} + \frac{1}{2} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}^{\mathrm{T}} \begin{bmatrix} Q_{xx}^{i} & Q_{xu}^{i} \\ Q_{ux}^{i} & Q_{uu}^{i} \end{bmatrix} \begin{bmatrix} \delta \boldsymbol{x} \\ \delta \boldsymbol{u} \end{bmatrix}.$$
(A.6)

where the partial derivatives of Q^i are as follows:

$$Q_x^i = l_x^i + f_x^{\rm T} (V^i)_x' \tag{A.7}$$

$$Q_{u}^{i} = l_{u}^{i} + f_{u}^{\mathrm{T}} (V^{i})_{x}^{\prime}$$
(A.8)

$$Q_{u}^{i} = l_{u}^{i} + f_{u}^{T} (V^{i})_{x}^{\prime}$$
(A.8)

$$Q_{xx}^{i} = l_{xx}^{i} + f_{x}^{T} (V^{i})_{xx}^{\prime} f_{x}$$
(A.9)

$$Q_{ux}^{i} = l_{ux}^{i} + f_{u}^{T} (V^{i})_{xx}^{\prime} f_{x}$$
(A.10)

$$Q_{ux}^{i} = l_{u}^{i} + f^{T} (V^{i})_{xx}^{\prime} f_{x}$$
(A.11)

$$Q_{ux}^{i} = l_{ux}^{i} + f_{u}^{1} \left(V^{i} \right)_{xx}^{\prime} f_{x}$$
(A.10)

$$Q_{uu}^{i} = l_{uu}^{i} + f_{u}^{\mathrm{T}} (V^{i})_{xx}' f_{u}$$
(A.11)

APPENDIX B PROOF OF LEMMA 1

We restate the theorem below for convenience.

Lemma 1 (Optimal MELQGames Policy). The optimal policy π^{i*} which solves (19) for each agent i has the form

$$\pi^{i*} = Z^{i^{-1}} \exp\left(-\frac{\alpha}{2} (\delta u^{i} - \delta u^{i*})^{\mathrm{T}} Q^{i}_{u^{i}u^{i}} (\delta u^{i} - \delta u^{i*})\right)$$
(B.1)

where δu^* is the solution to the following system of equations

$$0 = \hat{Q}_{uu}\delta u^* + Q_{ux}\delta x + Q_u = k + K\delta x$$
(B.2)

and the matrix $\hat{Q}_{uu} \in \mathbb{R}^{n_u \times n_u}$ is obtained by vertically stacking the row vectors $Q_{u^i u}^i$ for each $i \in \{1, \ldots, N\}$:

$$\hat{Q}_{uu} \coloneqq \begin{bmatrix} Q_{u^1 u}^1 \\ \vdots \\ Q_{u^N u}^N \end{bmatrix}$$
(B.3)

Proof: The optimal policy for agent i in (19) is the Gibbs distribution

$$\pi^{i*}(\delta u^i) \propto \exp\left(-\frac{1}{\alpha}\left[\left(Q^i_{u^i} + Q^i_{u^i x}\delta \boldsymbol{x} + Q^i_{u^i u^{-i}}\mathbb{E}\left[\delta u^{-i}\right]\right)^{\mathrm{T}}\delta u^i + \frac{1}{2}\delta u^{i^{\mathrm{T}}}Q^i_{u^i u^i}\delta u^i\right]\right).$$
(B.4)

Since the expression inside the exponential function is quadratic in δu^i with the negative definite quadratic term $-\frac{1}{2\alpha}Q^i_{u^iu^i}$, the optimal policy π^{i*} is multivariate Gaussian with covariance matrix $\Sigma = \alpha (Q^i_{u^iu^i})^{-1}$. To find the mean of π^{i*} , we complete the square to obtain

$$\left(Q_{u^{i}}^{i} + Q_{u^{i}x}^{i}\delta x + Q_{u^{i}u^{\neg i}}^{i}\mathbb{E}\left[\delta u^{\neg i}\right]\right)^{\mathrm{T}}\delta u^{i} + \frac{1}{2}\delta u^{i^{\mathrm{T}}}Q_{u^{i}u^{i}}^{i} = \frac{1}{2}(\delta u^{i} - \delta u^{i*})^{\mathrm{T}}Q_{u^{i}u^{i}}^{i}(\delta u^{i} - \delta u^{i*}) - \frac{1}{2}\delta u^{i*^{\mathrm{T}}}Q_{u^{i}u^{i}}^{i}\delta u^{i*}$$
(B.5)

where the mean $\delta u^{i*} = \mathbb{E}[\delta u^i]$ is defined as

$$\delta u^{i*} \coloneqq -(Q^i_{u^i u^i})^{-1}(Q^i_{u^i} + Q^i_{u^i x}\delta \boldsymbol{x} + Q^i_{u^i u^{-i}}\delta \boldsymbol{u}^{-i*})$$
(B.6)

where we have used $\delta u^{\neg i*} := \mathbb{E}[\delta u^{\neg i}]$ to denote the mean control of the other agents. If $\delta u^{\neg i*}$ were known (ex. in the IBR case where $\delta u^{\neg i*}$ is held fixed), then δu^{i*} can be solved from (B.6) alone. However, in this case, $\delta u^{\neg i*}$ is not known. However, by rearranging (B.6) we obtain

$$0 = Q_{u^{i}u^{i}}^{i}\delta u^{i*} + Q_{u^{i}}^{i} + Q_{u^{i}x}^{i}\delta x + Q_{u^{i}u^{-i}}^{i}\delta u^{-i*}$$
(B.7)

$$=Q_{u^{i}}^{i}+Q_{u^{i}x}^{i}\delta x+Q_{u^{i}u}^{i}\delta u^{*},$$
(B.8)

which provides n_{u^i} equations in the unknown variable $\delta u^* \in \mathbb{R}^{n_u}$. Hence, by considering (B.8) for all agents, we obtain the following linear equation for δu^* :

$$\hat{Q}_{uu}\delta u^* + Q_{ux}\delta x + Q_u = 0 \tag{B.9}$$

with \hat{Q}_{uu} defined as the stacked matrix of $Q_{u^i u}$ for all agents:

$$\hat{Q}_{uu} \coloneqq \begin{bmatrix} Q_{u^1 \boldsymbol{u}}^1 \\ \vdots \\ Q_{u^N \boldsymbol{u}}^N \end{bmatrix}$$
(B.10)

Appendix C

PROOF OF LEMMA 2

To simplify the proof, we first begin with the following technical lemma.

Lemma 5. Suppose that δu^* satisfies (21). Then,

$$\left(Q_{u^{\neg i}}^{i} + Q_{u^{\neg i}x}^{i} \delta x \right)^{\mathrm{T}} \delta u^{\neg i*} + \frac{1}{2} \delta u^{\neg i*} Q_{u^{\neg i}u^{\neg i}}^{i} \delta u^{\neg i*} - \frac{1}{2} \delta u^{i*} Q_{u^{i}u^{i}}^{i} \delta u^{i*} = \left(Q_{u}^{i} + Q_{ux}^{i} \delta x \right)^{\mathrm{T}} \delta u^{*} + \frac{1}{2} \delta u^{*} Q_{uu}^{i} \delta u^{*}$$

$$(C.1)$$

Proof: Since δu satisfies (21), it must thus satisfy (B.7). Hence,

$$-Q_{u^{i}u^{i}}^{i}\delta u^{i} = Q_{u^{i}}^{i} + Q_{u^{i}x}^{i}\delta \boldsymbol{x} + Q_{u^{i}u^{\neg i}}^{i}\delta \boldsymbol{u}^{\neg i*}$$
(C.2)

Splitting $-\frac{1}{2}$ into $\frac{1}{2} - 1$, using the above identity then simplifying:

$$\left(Q_{u^{-i}}^{i}+Q_{u^{-i}x}^{i}\delta x\right)_{\mathbf{T}}^{\mathrm{T}}\delta u^{-i*}+\frac{1}{2}\delta u^{-i*}Q_{u^{-i}u^{-i}}^{i}\delta u^{-i*}-\frac{1}{2}\delta u^{i*}Q_{u^{i}u^{i}}^{i}\delta u^{i*}$$
(C.3)

$$= \left(Q_{u^{\neg i}}^{i} + Q_{u^{\neg i}x}^{i}\delta x\right)_{\mathsf{T}}^{\mathsf{T}}\delta u^{\neg i*} + \frac{1}{2}\delta u^{\neg i*}Q_{u^{\neg i}u^{\neg i}}^{i}\delta u^{\neg i*} + \frac{1}{2}\delta u^{i*}Q_{u^{i}u^{i}}^{i}\delta u^{i*} - \delta u^{i*}Q_{u^{i}u^{i}}^{i}\delta u^{i}$$
(C.4)

$$= \left(Q_{u^{\neg i}}^{i} + Q_{u^{\neg i}x}^{i}\delta x\right)^{\mathrm{T}}\delta u^{\neg i*} + \frac{1}{2}\delta u^{\neg i*} Q_{u^{\neg i}u^{\neg i}}^{i}\delta u^{\neg i*} + \frac{1}{2}\delta u^{i*} Q_{u^{i}u^{i}}^{i}\delta u^{i*} + \delta u^{i*} Q_{u^{i}}^{i} + Q_{u^{i}x}^{i}\delta x + Q_{u^{i}u^{\neg i}}^{i}\delta u^{\neg i*}\right)$$
(C.5)

$$= \left(Q_{\boldsymbol{u}}^{i} + Q_{\boldsymbol{u}\boldsymbol{x}}^{i}\delta\boldsymbol{x}\right)^{\mathrm{T}}\delta\boldsymbol{u}^{*} + \frac{1}{2}\delta\boldsymbol{u}^{*\mathrm{T}}Q_{\boldsymbol{u}\boldsymbol{u}}^{i}\delta\boldsymbol{u}^{*}$$
(C.6)

We now restate Lemma 2 below for convenience.

Lemma 2 (MELQGames Value Function Update). Suppose the infimum in the Bellman equation (11) is solved with policy π^{i*} with mean $\delta u^* = k + K \delta x$ according to Lemma 1. Then, the value function using $\delta \tilde{u}$ for agent *i* has the form

$$V^{i}(\boldsymbol{x}) = \mathbb{E}_{\delta \boldsymbol{u} \sim \tilde{\pi}} \left[l^{i}(\boldsymbol{x}, \boldsymbol{u}) + V^{i'}(f(\boldsymbol{x}, \boldsymbol{u})) - \alpha H[\tilde{\pi}^{i}] \right]$$
(C.7)

$$= \left(V^{i} + V_{H}^{i}\right) + V_{x}^{i^{\mathrm{T}}}\delta \boldsymbol{x} + \frac{1}{2}\delta \boldsymbol{x}^{\mathrm{T}}V_{xx}^{i}\delta \boldsymbol{x},$$
(C.8)

where the terms V^i , V^i_H , V^i_x and V^i_{xx} have the form

$$V^{i} = \bar{V}^{i} + Q^{i}_{u}\boldsymbol{k} + \frac{1}{2}\boldsymbol{k}^{\mathrm{T}}Q^{i}_{uu}\boldsymbol{k}$$
(C.9)

$$V_H^i = \frac{\alpha}{2} \left(\log |Q_{u^i u^i}^i| - n_u \log(2\pi\alpha) \right)$$
(C.10)

$$+ \alpha \sum_{j=1, j \neq i} \operatorname{tr} \left[(Q_{u^j, u^j}^j)^{-1} Q_{u^j, u^j}^i \right]$$

$$V_x^i = Q_x^i + \mathbf{K}^{\mathrm{T}} Q_{uu}^i \mathbf{k} + \mathbf{K}^{\mathrm{T}} Q_u^i + Q_{xu}^i \mathbf{k}$$

$$V_x^i = Q_x^i + \mathbf{K}^{\mathrm{T}} Q_u^i + \mathbf{K}^{\mathrm{T}} Q_u^i + Q_{xu}^i \mathbf{k}$$
(C.11)
(C.12)

$$V_{xx}^{i} = Q_{xx}^{i} + \mathbf{K}^{\mathrm{T}} Q_{uu}^{i} \mathbf{K} + \mathbf{K}^{\mathrm{T}} Q_{ux}^{i} + Q_{xu}^{i} \mathbf{K}$$
(C.12)

Proof: Let \mathcal{I}^i denote the following integral.

$$\mathcal{I}^{i} \coloneqq \int \exp\left(-\frac{1}{\alpha} \left[Q_{u^{i}}^{i} + Q_{u^{i}x}^{i} \delta \boldsymbol{x} + Q_{u^{i}u^{-i}}^{i} \mathbb{E}\left[\delta \boldsymbol{u}^{-i}\right]\right)^{\mathrm{T}} \delta u^{i} + \frac{1}{2} \delta u^{i^{\mathrm{T}}} Q_{u^{i}u^{i}}^{i} \delta u^{i}\right]\right) \,\mathrm{d}\delta u^{i} \tag{C.13}$$

Then, completing the square as in (B.5) and simplifying, we obtain

$$\mathcal{I}^{i} = \exp\left(-\frac{1}{\alpha}\left[-\frac{1}{2}\delta u^{i*}{}^{\mathrm{T}}Q^{i}_{u^{i}u^{i}}\delta u^{i*}\right]\right) \int \exp\left(-\frac{1}{\alpha}\left[\frac{1}{2}(\delta u^{i}-\delta u^{i*}){}^{\mathrm{T}}Q^{i}_{u^{i}u^{i}}(\delta u^{i}-\delta u^{i*})\right]\right) \,\mathrm{d}\delta u^{i} \tag{C.14}$$

$$= \exp\left(-\frac{1}{\alpha} \left[-\frac{1}{2} \delta u^{i*}{}^{\mathrm{T}} Q^{i}_{u^{i}u^{i}} \delta u^{i*}\right]\right) \left(2\pi\right)^{\frac{n_{u}}{2}} \left(\left|\alpha(Q^{i}_{u^{i}u^{i}})^{-1})\right|\right)^{\frac{1}{2}}$$
(C.15)

$$= \exp\left(-\frac{1}{\alpha} \left[-\frac{1}{2} \delta u^{i*}{}^{\mathrm{T}} Q^{i}_{u^{i}u^{i}} \delta u^{i*}\right]\right) \frac{\left(2\pi\alpha\right)^{-2}}{\left|Q^{i}_{u^{i}u^{i}}\right|^{\frac{1}{2}}}$$
(C.16)

$$= \exp\left(-\frac{1}{\alpha}\left[-\frac{1}{2}\delta u^{i*}{}^{\mathrm{T}}Q^{i}_{u^{i}u^{i}}\delta u^{i*} + \frac{\alpha}{2}\left(\ln|Q^{i}_{u^{i}u^{i}}| - n_{u^{i}}\ln(2\pi\alpha)\right)\right]\right)$$
(C.17)

Then, the partition function Z^i (13) takes the form

$$Z^{i} = \int \exp\left(-\frac{1}{\alpha}\mathbb{E}_{\boldsymbol{u}^{\neg i}}\left[V^{i'}(f(\boldsymbol{x},\boldsymbol{u})) + l^{i}(\boldsymbol{x},\boldsymbol{u})\right]\right) du^{i}$$
(C.18)

$$= \exp\left(-\frac{1}{\alpha}\left[\bar{V}^{i} + Q_{\boldsymbol{x}}^{i}{}^{\mathrm{T}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}Q_{\boldsymbol{x}\boldsymbol{x}}^{i}\delta\boldsymbol{x} + \left(Q_{\boldsymbol{u}^{\neg i}}^{i} + Q_{\boldsymbol{u}^{\neg i}\boldsymbol{x}}^{i}\delta\boldsymbol{x}\right)^{\mathrm{T}}\delta\boldsymbol{u}^{\neg i\ast} + \frac{1}{2}\delta\boldsymbol{u}^{\neg i\ast}{}^{\mathrm{T}}Q_{\boldsymbol{u}^{\neg i}\boldsymbol{u}^{\neg i}}^{i}\delta\boldsymbol{u}^{\neg i\ast} + \alpha\sum_{j=1, j\neq i}^{N} \operatorname{tr}\left[\left(Q_{\boldsymbol{u}^{j}, \boldsymbol{u}^{j}}^{j}\right)^{-1}Q_{\boldsymbol{u}^{j}, \boldsymbol{u}^{j}}^{i}\right]\right]\right)$$
(C.19)

$$\int \exp\left(-\frac{1}{\alpha}\left[\left(Q_{\boldsymbol{u}^{i}}^{i} + Q_{\boldsymbol{u}^{i}\boldsymbol{x}}^{i}\delta\boldsymbol{x} + Q_{\boldsymbol{u}^{i}\boldsymbol{u}^{\neg i}}^{i}\mathbb{E}\left[\delta\boldsymbol{u}^{\neg i}\right]\right)^{\mathrm{T}}\delta\boldsymbol{u}^{i} + \frac{1}{2}\delta\boldsymbol{u}^{\mathrm{T}}Q_{\boldsymbol{u}^{i}\boldsymbol{u}^{i}}^{i}\delta\boldsymbol{u}^{i}\right]\right) d\delta\boldsymbol{u}^{i}$$

$$= \exp\left(-\frac{1}{\alpha}\left[\bar{V}^{i} + Q_{\boldsymbol{x}}^{i}{}^{\mathrm{T}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}Q_{\boldsymbol{x}\boldsymbol{x}}^{i}\delta\boldsymbol{x} + \left(Q_{\boldsymbol{u}^{\neg i}}^{i} + Q_{\boldsymbol{u}^{\neg i}\boldsymbol{x}}^{i}\delta\boldsymbol{x}\right)^{\mathrm{T}}\delta\boldsymbol{u}^{\neg i\ast} + \frac{1}{2}\delta\boldsymbol{u}^{\neg i\ast}{}^{\mathsf{T}}Q_{\boldsymbol{u}^{\neg i}\boldsymbol{u}^{\neg i}}^{i}\delta\boldsymbol{u}^{\neg i\ast} + \alpha\sum_{j=1, j\neq i}^{N}\operatorname{tr}\left[\left(Q_{\boldsymbol{u}^{j}, \boldsymbol{u}^{j}\right)^{-1}Q_{\boldsymbol{u}^{j}, \boldsymbol{u}^{j}}^{i}\right]\right]\right)\mathcal{I}$$

$$= \exp\left(-\frac{1}{\alpha}V^{i}(\boldsymbol{x})\right)$$
(C.21)

where the trace term comes from taking the expectation of the quadratic term in $\delta u^{\neg i}$ and

$$V^{i}(\boldsymbol{x}) = (\bar{V}^{i})' + Q_{\boldsymbol{x}}^{i}{}^{\mathrm{T}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}Q_{\boldsymbol{x}\boldsymbol{x}}^{i}\delta\boldsymbol{x} + \left(Q_{\boldsymbol{u}^{\neg i}}^{i} + Q_{\boldsymbol{u}^{\neg i}\boldsymbol{x}}^{i}\delta\boldsymbol{x}\right)^{\mathrm{T}}\delta\boldsymbol{u}^{\neg i*} + \frac{1}{2}\delta\boldsymbol{u}^{\neg i*}{}^{\mathrm{T}}Q_{\boldsymbol{u}^{\neg i}\boldsymbol{u}^{\neg i}}^{i}\delta\boldsymbol{u}^{\neg i*} - \frac{1}{2}\delta\boldsymbol{u}^{i*}{}^{\mathrm{T}}Q_{\boldsymbol{u}^{i}\boldsymbol{u}^{i}}^{i}\delta\boldsymbol{u}^{i*} + \frac{\alpha}{2}\left(\ln|Q_{\boldsymbol{u}^{i}\boldsymbol{u}^{i}}^{i}| - n_{\boldsymbol{u}^{i}}\ln(2\pi\alpha)\right) + \alpha\sum_{j=1, j\neq i}^{N}\operatorname{tr}\left[(Q_{\boldsymbol{u}^{j},\boldsymbol{u}^{j}}^{j})^{-1}Q_{\boldsymbol{u}^{j},\boldsymbol{u}^{j}}^{i}\right]$$
(C.22)

Applying Lemma 5, using the definition of δu^* and collecting terms then yields

$$V^{i}(\boldsymbol{x}) = (\bar{V}^{i})' + Q_{\boldsymbol{x}}^{i}{}^{\mathrm{T}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}Q_{\boldsymbol{x}\boldsymbol{x}}^{i}\delta\boldsymbol{x} + \left(Q_{\boldsymbol{u}}^{i} + Q_{\boldsymbol{u}\boldsymbol{x}}^{i}\delta\boldsymbol{x}\right)^{\mathrm{T}}\delta\boldsymbol{u}^{*} + \frac{1}{2}\delta\boldsymbol{u}^{*}{}^{\mathrm{T}}Q_{\boldsymbol{u}\boldsymbol{u}}^{i}\delta\boldsymbol{u}^{*} + \frac{\alpha}{2}\left(\ln|Q_{\boldsymbol{u}^{i}\boldsymbol{u}^{i}}^{i}| - n_{\boldsymbol{u}^{i}}\ln(2\pi\alpha)\right) + \alpha\sum_{j=1, j\neq i}^{N}\operatorname{tr}\left[(Q_{\boldsymbol{u}^{j},\boldsymbol{u}^{j}}^{j})^{-1}Q_{\boldsymbol{u}^{j},\boldsymbol{u}^{j}}^{i}\right]$$

$$(C.23)$$

$$= \left(V^{i} + V_{H}^{i}\right) + V_{x}^{i^{\mathrm{T}}}\delta\boldsymbol{x} + \frac{1}{2}\delta\boldsymbol{x}^{\mathrm{T}}V_{xx}^{i}\delta\boldsymbol{x}$$
(C.24)

where

$$V^{i} = \bar{V}^{i} + Q^{i}_{u}\boldsymbol{k} + \frac{1}{2}\boldsymbol{k}^{\mathrm{T}}Q^{i}_{uu}\boldsymbol{k}$$
(C.25)

$$V_H^i = \frac{\alpha}{2} \left(\log |Q_{u^i u^i}^i| - n_u \log(2\pi\alpha) \right)$$
(C.26)

$$+ \alpha \sum_{j=1, j \neq i}^{N} \operatorname{tr} \left[(Q_{u^{j}, u^{j}}^{j})^{-1} Q_{u^{j}, u^{j}}^{i} \right]$$
(C.20)

$$V_x^i = Q_x^i + \mathbf{K}^{\mathrm{T}} Q_{uu}^i \boldsymbol{k} + \mathbf{K}^{\mathrm{T}} Q_u^i + Q_{xu}^i \boldsymbol{k}$$
(C.27)

$$V_{xx}^{i} = Q_{xx}^{i} + \mathbf{K}^{\mathrm{T}} Q_{uu}^{i} \mathbf{K} + \mathbf{K}^{\mathrm{T}} Q_{ux}^{i} + Q_{xu}^{i} \mathbf{K}$$
(C.28)

APPENDIX D Proof of Lemma 3

We restate Lemma 3 below for convenience.

Lemma 3. Suppose that the augmented Lagrangian MELQGames converges to a tuple $(\pi^*, \lambda_j^i, \rho)$ for some $\rho > 0$ which satisfies dual feasibility of λ , LICQ and SOSC. Then, π^* is a GNE for (5).

Proof: Since the solution satisfies LICQ, SOSC and $\lambda > 0$, each π^{i*} is a local minima for its corresponding constrained optimal control problem (5). Thus, there does not exist any feasible control $\pi \in \Pi(\pi^{-i*})$ within a neighborhood of π^{i*} that can achieve a lower cost. Hence, π^* satisfies (9) and is a GNE.

Appendix E

PROOF OF LEMMA 4

Before proving Lemma 4, we first introduce the following theorem from Maschler et al. [18] which presents a sufficient condition for an event A to be common belief among the players at a particular state of the world ω :

Theorem 1. Let N denote the number of agents in the game, $\omega \in Y$ be a state of the world and $A \subseteq Y$ be an event satisfying the following two conditions:

$$p_i(A|\omega) = 1, \quad \forall i \in \{1, \dots, N\}$$
(E.1)

$$p_i(A|\omega') = 1, \quad \forall i \in \{1, \dots, N\}, \ \forall \omega' \in A$$
(E.2)

Then, A is common belief among the agents at ω .

Using Theorem 1, we can now prove (4), which we restate below for convenience.

Lemma 4. Let $E = \{\omega_{a,s}\}_{a=1}^{A}$ denote the event that the ego agent knows the correct mode with probability 1. Then, for all $i \neq 1, \omega \in B_i E$, i.e. all non-ego agents believe that the ego agent knows the mode perfectly. Furthermore, it is common belief that $B_i E_a$ for all $i \neq 1$. In particular, the ego agent (correctly) believes that all non-ego agents (incorrectly) believes that it knows the correct mode.

Proof: First, note that at $\omega \in E$, all agents, including the ego agent, know the true mode. Using the definition of the belief operator B_i (38), we have that for $i \neq 1$:

$$B_i E = \{\omega \in Y | p_i(E|\omega) = 1\} = Y.$$
(E.3)

In other words, non-ego agent *i* believes that *E* will attains no matter what the true world state ω_* is. Since $p_i(Y|\omega) = 1$ for any $\omega \in Y$, $B_i E = Y$ is thus common belief among the agents at all world states. Hence, by the definition of common belief, we have that

$$B_1 B_i E = B_1 Y = Y \tag{E.4}$$

i.e., the ego agent believes that the non-ego agents believe it knows the true mode. On the other hand, we have that

$$B_1 E = \{ \omega \in Y | p_1(E|\omega) = 1 \} = E,$$
(E.5)

$$B_1 E^{\complement} = \{ \omega \in Y | p_1(E^{\complement} | \omega) = 1 \} = E^{\complement}.$$
(E.6)

where E^{\complement} denotes the complement of E in Y. In other words, when $\omega \in E^{\complement}$, the ego agent (correctly) believes that it does not know the true mode while simultaneously knowing that the non-ego agents (incorrectly) believe that it knows the true mode.