

On the coordination efficiency of strategic multi-agent robotic teams

Marcos M. Vasconcelos¹ and Behrouz Touri²

Abstract—We study the problem of achieving decentralized coordination by a group of strategic decision-makers choosing to engage or not in a task in a stochastic setting. First, we define a class of symmetric utility games that encompass a broad class of coordination games, including the popular framework known as *global games*. To study the extent to which agents engaging in a stochastic coordination game indeed coordinate, we propose a new probabilistic measure of coordination efficiency. Then, we provide a universal information-theoretic upper bound on the coordination efficiency as a function of the amount of noise in the observation channels. Finally, we revisit a large class of global games, and we illustrate that their Nash equilibrium policies may be less coordination efficient than certainty equivalent policies, despite them providing better expected utility. This counter-intuitive result, establishes the existence of a nontrivial trade-off between coordination efficiency and expected utility in coordination games.

I. INTRODUCTION

Coordinated behavior is desirable in many distributed autonomous systems such as robotic, social-economic, and biological networks [1]–[5]. Most of the Engineering literature on coordination assumes that the agents exchange messages over a communication network to asymptotically agree on a common decision variable, such as in opinion dynamics and distributed optimization. However, in the field of Economics, the topic of coordination has been studied from a different point of view, where the agents do not exchange (explicit) messages but instead act strategically. Such an approach is related to coordination games, in which two or more interacting agents are incentivized to take the same action. Deterministic coordination games are characterized by the existence of multiple equilibria, and often lead to the analysis of social dilemmas. One way to address the multiplicity of equilibria uses a framework known as *global games* [6].

A global game is a Bayesian coordination game, where each agent plays an action after observing a noisy signal about the state-of-the-world. The state-of-the-world, which we simply refer as *state* captures features such as the strength of the economy in a bank run model, the political regime in a regime change model, or the difficulty of a task in a task-allocation problem. Under certain assumptions on the utility structure, global games admit a unique Bayesian Nash Equilibrium even in the presence of a vanishingly small noise in the observations, resolving the issue of equilibrium selection in games with multiple equilibria [7].

¹M. M. Vasconcelos is with the Department of Electrical and Computer Engineering, FAMU-FSU College of Engineering, Florida State University, USA. Email: m.vasconcelos@fsu.edu

²Behrouz Touri is with the Department of Electrical Engineering, University of California San Diego, USA. Email: btouri@eng.ucsd.edu




Fig. 1. System architecture for the strategic coordination in a robotic team receiving information about a stochastic state variable available at a remote location.

The recent literature in this class of games focuses on aspects related to existence of Nash equilibria in the presence of different information patterns [8], or the influence of correlation among the agent’s observations [9], [10], and the impact of different local connectivity patterns in terms of externality in the agents’ utility functions [11]. Other recent developments look at non-conventional probabilistic models for the noisy signals [12], and the presence of a multi-dimensional state in a multi-task allocation problem [13].

We consider the federated system architecture outlined in Fig. 1, where the state is available at a remote location (e.g. a cloud server), and broadcast to multiple agents by an edge node or gateway over parallel noisy channels. Upon receiving its noisy signal, an agent makes a binary decision such as to maximize an expected utility function satisfying the strategic complementarity property, leading to coordinated behavior [14]. We study how the coordination in a global game degrades with the level of noise in the communication channels. Moreover, we are interested in characterizing the limits of coordination for a given signal to noise ratio used for communication with the robotic agents.

The main contributions of this paper are:

- We introduce a class of games, namely homogeneous coordination games, that includes a broad class of global games.
- We introduce a novel notion of coordination efficiency used to measure the coordination for the homogeneous coordination games.
- We obtain a fundamental limitation on the coordination efficiency in global games for any policy based on information theoretic tools.

II. SYSTEM MODEL

In this section, we discuss our model for global games, and in particular, we introduce an important subclass of such games, i.e., *homogeneous coordination* games.

A. Utility structure

A global game is an incomplete information game that is played between N players $[N] \stackrel{\text{def}}{=} \{1, \dots, N\}$ and nature. Formally, a global game is a tuple $([N], \mathcal{A}, X, \mathbf{u}, \mathbf{Y})$, where:

- (i) $\mathcal{A} = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_N$ is the joint action set of the N players with \mathcal{A}_i being the action set for player $i \in [N]$,
- (ii) X is a random variable determining the *type* of nature. We refer to X as the state or the underlying fundamental of the game,
- (iii) $\mathbf{u} = (u_1, \dots, u_N) : \mathcal{A} \times \mathbb{R} \rightarrow \mathbb{R}^N$ is the utility of the N players with $u_i : \mathcal{A} \times \mathbb{R} \rightarrow \mathbb{R}$ being the utility of the i -th player that depends on the action of each player and the value of the underlying fundamental X , and
- (iv) $\mathbf{Y} = (Y_1, \dots, Y_N)$ where Y_i is a random variable denoting the player $i \in [N]$ noisy observation (that forms the belief) of the underlying fundamental X .

In this work, for any vector $\mathbf{v} = (v_1, \dots, v_N)$ and any $i \in [N]$, we use the notation $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_N)$ and with abuse of notation, we say $\mathbf{v} = (v_i, \mathbf{v}_{-i})$. For example, for a joint action $\mathbf{a} \in \mathcal{A}$, we write $\mathbf{a} = (a_i, \mathbf{a}_{-i})$ for all $i \in [N]$.

Our work is motivated by the observation that a vast majority of studies in global games and their applications, the underlying games has the following common features:

- (a) *Symmetric/Permutation invariant*: In many settings, the utility functions of individual agents are invariant under any permutation of other agents' actions. In other words, $u_i(a_i, \mathbf{a}_{-i}) = u_i(a_i, \mathbf{a}_{-i}P)$ for any $(N-1) \times (N-1)$ permutation matrix¹ P .
- (b) *Homogeneous utility functions*: The utility function of the N players are the same in the sense that for any player $i \in [N]$ and any action profile $\mathbf{a} \in \mathcal{A}$, we have $u_i(a_i, \mathbf{a}_{-i}) = u_1(a_i, \mathbf{a}_{-i})$.
- (c) *Homogeneous action sets*: In many global games, we are dealing with a large population, and the action set of all players are identical. For example, in the case of political riots, all players decide to take a risky action or safe action in the face of a political regime. In this case, $\mathcal{A}_i = \{0, 1\}$ where 0 and 1 correspond to the safe and risky actions, respectively.
- (d) *Coordination promoting*: Again, in most settings of interest, the utility structure of the players is such that it promotes coordination. For example, in the case of political uprisings, bank-runs, etc., a well-studied utility function is $u_i(\mathbf{a}, X) = a_i(\sum_{j=1}^N a_j - X)$. Therefore, in the case, where all players have the perfect information about X , i.e., when $Y_i = X$ for all $i \in [N]$, depending on whether $X > \frac{1}{n}$ or $X < \frac{1}{n}$, the only equilibrium of the

¹A matrix is a permutation matrix if all its elements are zero or one and each row and each column has exactly one non-zero element.

game is either $\mathbf{a} = \mathbf{1}$ or $\mathbf{a} = \mathbf{0}$, resulting in coordination among the N players.

Motivated by this, we introduce the notion of *homogeneous coordination* games that formalize a broad class of games satisfying the above properties. For this, let $\Delta^k := \{\mathbf{q} \in \mathbb{R}_+^k \mid \sum_{i=1}^k q_i = 1\}$ be the probability simplex in \mathbb{R}^k . For a finite set $\mathcal{A} = \{0, \dots, M-1\}$ and a vector $\mathbf{v} \in \mathcal{A}^d$, where $d \geq 1$, let us define the empirical mass function $G(\mathbf{v}) = (q_0, \dots, q_{M-1}) \in \Delta^M$ by

$$q_\ell = \frac{1}{\gamma} \sum_j \mathbf{1}(v_j = \ell), \quad \ell \in \mathcal{A}.$$

Basically, q_ℓ is the proportion of the entries of \mathbf{v} that are equal to $\ell \in \mathcal{A}$. Now, we are ready to formalize the class of homogeneous coordination games.

Definition 1 (Homogeneous Coordination Game): A homogeneous coordination game is a game where all the agents have the same action set \mathcal{A} , the same utility function $u : \mathcal{A}^N \times \mathbb{R} \rightarrow \mathbb{R}$, satisfying the following conditions:

- (1) There exists a function $\hat{u} : \mathcal{A} \times \Delta^{N-1} \times \mathbb{R} \rightarrow \mathbb{R}$ where for all $\mathbf{a} \in \mathcal{A}^N$, all $i \in [N]$, and all $x \in \mathbb{R}$, we have

$$\hat{u}(a_i, G(\mathbf{a}_{-i}), x) = u(a_i, \mathbf{a}_{-i}, x). \quad (1)$$

- (2) For all $i \in [N]$, all $x \in \mathbb{R}$, and all $\mathbf{g} \in \Delta^{N-1}$, there exists an optimal action $a^*(x)$ and majority $c^*(x)$ of the players, such that if the majority are playing $a^*(x)$, then player i is better off playing that action. Mathematically, there exists $c^*(x) \in [0, 1]$ such that for any $\mathbf{g} \in \Delta^M$ with $g_{a^*(x)} = \max_{\ell \in [M]} g_\ell \geq c^*(x)$, we have

$$\hat{u}(a_i^*(x), \mathbf{g}, x) \geq \hat{u}(a_i, \mathbf{g}, x), \quad \forall a_i \in \mathcal{A}. \quad (2)$$

Note that Property (1) essentially means that the utility function of each player is symmetric/permutation invariant. In other words, for a finite action set $\mathcal{A} = [M]$, any symmetric/permutation invariant function $f : \mathcal{A}^{N-1} \rightarrow \mathbb{R}$ (as defined in (a)), can be written as a function of the empirical mass function of the M actions, i.e., for such utility functions, it does not matter which player is playing what action, but rather how many or what proportion of the players is playing each action.

For the rest of the paper, with an abuse of notation, instead of $u(a_i, \mathbf{a}_{-i}, x)$, we may view the utility functions of a homogeneous coordination game to be simply a function of the empirical mass and use the notation $u(a_i, \mathbf{g}_{-i}, x)$ instead of $\hat{u}(a_i, G(\mathbf{a}_{-i}), x)$, where $\mathbf{g}_{-i} = G(\mathbf{a}_{-i})$.

B. Information structure and policies

Here, we discuss the assumptions on the fundamental X and individual agents' noisy observation of Y_i . Throughout, we assume that X is a zero-mean Gaussian random variable with variance σ_X^2 , i.e., $X \sim \mathcal{N}(0, \sigma_X^2)$. We assume the commonly studied model (cf. [6], [8], [9]) for the i -th agent noisy observation Y_i to be $Y_i = X + Z_i$. We assume that the noise sequence $\{Z_i\}_{i \in [N]}$, is independent and identically distributed across agents and $Z_i \sim \mathcal{N}(0, \sigma_Z^2)$. Moreover, $\{Z_i\}$ is independent of X .

Note that since X and Y_i are jointly Gaussian and the minimum mean squared error estimate of X given $Y_i = y_i$ is linear and is given by

$$\hat{x}_{\text{mmse}}(y_i) \stackrel{\text{def}}{=} \mathbb{E}[X | Y_i = y_i] = \left(\frac{\sigma_X^2}{\sigma_X^2 + \sigma_Z^2} \right) y_i. \quad (3)$$

In general for games of imperfect information (which includes global games and homogeneous coordination games), the agents take action based on their observation. This leads to the notion of *policy*. For homogeneous coordination games with the action set $\mathcal{A} = \{1, \dots, M\}$, a *policy* is a mapping $\gamma_i : \mathbb{R} \rightarrow \mathcal{A}$ that translates agent i -s observation to action, i.e., agent $i \in [N]$ takes action $a_i = \gamma_i(Y_i)$.

C. Bayesian Nash Equilibrium

Let u_i be the utility function of a homogeneous coordination game (Definition 1). The agents in this game act in a noncooperative manner, by seeking to maximize their individual *expected* utility with respect to their coordination policies. Let $\gamma \stackrel{\text{def}}{=} (\gamma_1, \dots, \gamma_N)$ be a *policy profile*, i.e., the collection of policies used by all the agents in the system.

Given γ_{-i} , the goal of the i -th agent is to solve the stochastic optimization problem

$$\underset{\gamma_i}{\text{maximize}} \quad \mathcal{J}_i(\gamma_i, \gamma_{-i}) \stackrel{\text{def}}{=} \mathbb{E}[u_i(A_i, G_{-i}, X)],$$

where the expectation is taken over all the exogenous random variables X , and $\{Z_i\}_{i \in [N]}$. This leads to the notion of Bayesian Nash Equilibrium (BNE) strategies.

Definition 2 (Bayesian Nash Equilibrium): A policy profile γ^* is a Bayesian Nash Equilibrium if

$$\mathcal{J}_i(\gamma_i^*, \gamma_{-i}^*) \geq \mathcal{J}_i(\gamma_i, \gamma_{-i}^*), \quad \text{for all } \gamma_i \in \Gamma, \quad i \in [N],$$

where Γ is the space of all admissible coordination policies/strategies.

D. Coordination measure

Given that the state X is not perfectly observed by the agents, full coordination is often unachievable. In a deterministic setting, defining a precise notion of coordination and agreement is a well-posed problem. However, there are multiple ways of defining a metric of coordination efficiency in a stochastic game setting. One essential feature that such a metric should have is to capture that the extent to which agents coordinate around an optimal action degrades with respect to the amount of noise in the observations. The introduction of framework of homogeneous coordination games, allows us to mathematically define a measure of *coordination efficiency*.

Definition 3 (Coordination efficiency): Let γ be a policy profile of N players in a homogeneous coordination game (as defined in Definition 1). We define the average coordination efficiency $\varrho : \gamma \mapsto [0, 1]$ as

$$\varrho(\gamma) \stackrel{\text{def}}{=} \frac{1}{N} \sum_{i=1}^N \mathbb{P}(\gamma_i(Y_i) = a^*(X)),$$

where $a^*(x)$ is the optimal action defined in Definition 1.

III. GLOBAL GAMES REVISITED

An important instance of homogeneous coordination games is a class of binary action global games (i.e., $\mathcal{A} = \{0, 1\}$), where the utility of each agent is given by

$$u_i(a_i, \mathbf{a}_{-i}, x) = a_i \cdot \left(b \left(\sum_{j \neq i} a_j \right) - x \right), \quad (4)$$

where $b : [0, N-1] \rightarrow \mathbb{R}$ is a continuous and increasing function. The function $b(\cdot)$ is called the *benefit function*. One application for this utility is in distributed task allocation in robotic teams [15], [16], where x represents the difficulty of a task. An agent benefits from engaging in the task if the number of other agents engage in the same action is sufficiently large. However, if the variable x is not perfectly observed by the agents it is not clear whether an agent should engage in the task or not.

Our next result establishes that in fact global games with utility structure (4) are homogeneous coordination games.

Lemma 1: A global game with $N \geq 2$ players, binary action set $\mathcal{A} = \{0, 1\}$, and utility function (4) is a homogeneous coordination game for any increasing continuous function $b : [0, N-1] \rightarrow \mathbb{R}$.

Proof: The homogeneity of the action sets and utility functions follow readily from the definition of such games. To show Property (1), for $\mathbf{a} \in \mathcal{A}^N$ and $i \in [N]$, let $p = \frac{\sum_{j \neq i} a_j}{N-1}$. Then, $G(\mathbf{a}_{-i}) = (1-p, p)$. Therefore, letting $\hat{u}(a_i, \mathbf{g}, x) \stackrel{\text{def}}{=} a_i(b((N-1)g_1) - x)$ for all $a_i \in \mathcal{A}$, $\mathbf{g} = (g_0, g_1) \in \Delta^2$, and $x \in \mathbb{R}$, we have

$$u_i(a_i, \mathbf{a}_{-i}, x) = a_i \left(b((N-1)p) - x \right) = \hat{u}(a_i, G(\mathbf{a}_{-i}), x). \quad (5)$$

To show Property (2), fix $x \in \mathbb{R}$ and let $b(\cdot)$ be an increasing benefit function. Then, if $x \leq b(0)$, we have

$$\hat{u}(1, \mathbf{g}, x) = b((N-1)g_1) - x \geq b(0) - x \geq \hat{u}_i(0, \mathbf{g}, x) = 0.$$

Therefore, for any $x \leq b(0)$, (2) holds with $a^*(x) = 1$ and $c^*(x) = 0$. Similarly, it can be shown that for $x \geq b(N-1)$, (2) holds with $a^*(x) = 0$ and $c^*(x) = 0$. For $x \in (b(0), b(N-1))$, we can show that both $a^*(x) = 0$ and $a^*(x) = 1$ are possible coordinating actions. To show $a^*(x) = 1$, let $c^*(x) = \min\{q \in [0, 1] \mid b(q(N-1)) \geq x\}$ (note that the minimum exists due to the continuity of $b(\cdot)$ and compactness of $[0, 1]$). Then for any probability vector $\mathbf{g} = (g_0, g_1) \in \Delta^2$ with $g_1 \geq c^*(x) = q$, we have

$$\begin{aligned} \hat{u}(1, \mathbf{g}, x) &= b((N-1)g_1) - x \\ &\geq b((N-1)q) - x \geq \hat{u}(0, \mathbf{g}, x) = 0. \end{aligned}$$

Therefore, condition (2) holds. Similarly, it can be shown that for $x \in (b(0), b(N-1))$, $a^*(x) = 0$ is a coordinating action with $c^*(x) = \max\{q \in [0, 1] \mid b((1-q)(N-1)) \geq x\}$. ■

In this work, we study the so-called *best-response* policy for the above games. In our case, for a joint policy γ , agent

i -s best-response policy is

$$\text{BR}(y_i, \gamma_{-i}) = \begin{cases} 1 & \text{if } \mathbb{E}[b(\sum_{i \neq j} \gamma(Y_j)) \mid Y_i = y_i] \\ & \geq \mathbb{E}[X \mid Y_i = y_i] \\ 0 & \text{otherwise} \end{cases}.$$

A. Threshold policies and their best-response

In many games of imperfect information, including global games, we are interested in the class of threshold policies. For global games with binary actions, these are policies where an agent compares its observed signal $Y_i = y_i$ to a threshold τ_i and decides whether to take the risky action ($a_i = 1$) or not ($a_i = 0$), i.e.,

$$\gamma_i(y_i) = \begin{cases} 1, & y_i \leq \tau_i \\ 0, & \text{otherwise.} \end{cases}$$

Using the next result, we will show that the best response to homogeneous threshold policies is a threshold policy.

Lemma 2: If the function $b(\cdot)$ is nonnegative and strictly increasing, and all other agents $j \neq i$ utilize a threshold policy γ with the same threshold τ , then

$$\mathbb{E}\left[b\left(\sum_{i \neq j} \gamma(Y_j)\right) \mid Y_i = y_i\right]$$

is a strictly decreasing function of y_i .

Proof: Let B_{-i} denote a function of random variables $\{Y_j\}_{j \neq i}$ given by $B_{-i} \stackrel{\text{def}}{=} b\left(\sum_{j \neq i} \gamma(Y_j)\right)$, with the Cumulative Distribution Function (CDF)

$$F_{B_{-i}|Y_i=y_i}(\xi) \stackrel{\text{def}}{=} \mathbb{P}(B_{-i} \leq \xi \mid Y_i = y_i).$$

Since $b(\cdot) \geq 0$, B_{-i} is a nonnegative random variable. Therefore (cf. [17, Chapter 1.5, Property E.6]),

$$\mathbb{E}[B_{-i} \mid Y_i = y_i] = \int_0^\infty (1 - F_{B_{-i}|Y_i=y_i}(\xi)) d\xi.$$

Because the function b is strictly increasing, it admits a unique inverse function $b^{-1} : \mathbb{R}_+ \rightarrow \mathbb{R}_+$. Also, since the observations $\{Y_j\}_{j \in [N]}$ are conditionally (on X) independent, therefore,

$$F_{B_{-i}|Y_i=y_i}(\xi) = \int_{\mathbb{R}} \mathbb{P}\left(\sum_{i \neq j} \gamma(Y_j) \leq b^{-1}(\xi) \mid X = x\right) \times f_{X|Y_i=y_i}(x) dx.$$

Let $A_j = \gamma(Y_j)$ for $j \neq i$. Conditioned on $X = x$, the collection of Bernoulli random variables $\{A_j\}_{j \neq i}$ is mutually independent with

$$\mathbb{P}(A_j = 1 \mid X = x) = \Phi\left(\frac{\tau_j - x}{\sigma_Z}\right),$$

where Φ is the CDF of a standard Gaussian random variable².

²The CDF of a standard Gaussian random variable is given by

$$\Phi(x) \stackrel{\text{def}}{=} \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\xi^2}{2}\right) d\xi.$$

Under the assumption of an homogeneous threshold strategy profile where $\tau_j = \tau$ for all $j \neq i$, $\{A_j\}_{j \neq i}$ is identically distributed, which implies that

$$\sum_{j \neq i} A_j \mid X = x \sim \mathcal{B}\left(N - 1, \Phi\left(\frac{\tau - x}{\sigma_Z}\right)\right),$$

where $\mathcal{B}(k, p)$ is a binomial distribution with parameters (k, p) . Therefore,

$$F_{B_{-i}|Y_i=y_i}(\xi) = \mathbb{E}_V \left[\sum_{\ell=0}^{\lfloor b^{-1}(\xi) \rfloor} \binom{N-1}{\ell} \Phi\left(\frac{\tau - V - \alpha y_i}{\sigma_Z}\right)^\ell \times \left(1 - \Phi\left(\frac{\tau - V - \alpha y_i}{\sigma_Z}\right)\right)^{N-1-\ell} \right],$$

where the expectation is with respect to a random variable V with

$$V \sim \mathcal{N}\left(0, \frac{\sigma_X^2 \sigma_Z^2}{\sigma_X^2 + \sigma_Z^2}\right) \text{ and } \alpha \stackrel{\text{def}}{=} \frac{\sigma_X^2}{\sigma_X^2 + \sigma_Z^2}. \quad (6)$$

Let $p(\tau, v, y_i) \stackrel{\text{def}}{=} \Phi\left(\frac{\tau - v - \alpha y_i}{\sigma_Z}\right)$. Note that $p(\cdot, \cdot, y_i)$ is strictly decreasing in y_i , and since the CDF of a binomial random variable computed at a point is a strictly decreasing function in the probability parameter p , we have

$$\frac{\partial}{\partial y_i} (1 - F_{B_{-i}|Y_i=y_i}(\xi)) < 0.$$

Therefore,

$$\frac{\partial}{\partial y_i} \int_0^\infty (1 - F_{B_{-i}|Y_i=y_i}(\xi)) d\xi < 0. \quad \blacksquare$$

Theorem 1: If the benefit function $b : \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is nonnegative and strictly increasing, the best-response map to a homogeneous threshold strategy profile is a threshold strategy.

Proof: Let

$$f(y_i) \stackrel{\text{def}}{=} \mathbb{E}[b(\sum_{j \neq i} \gamma(Y_j)) \mid Y_i = y_i] - \mathbb{E}[X \mid Y_i = y_i].$$

Lemma 2 implies that $\mathbb{E}[b(\sum_{j \neq i} \gamma(Y_j)) \mid Y_i = y_i]$ is monotonically decreasing in y_i while

$$\mathbb{E}[X \mid Y_i = y_i] = \left(\frac{\sigma_X^2}{\sigma_X^2 + \sigma_Z^2}\right) y_i$$

is a strictly increasing function of y_i . Therefore, $f(y_i)$ is strictly decreasing. Also, since b is an increasing function, $\mathbb{E}[b(\sum_{j \neq i} \gamma(Y_j)) \mid Y_i = y_i] \in [b(0), b(N-1)]$ and hence, $\lim_{y_i \rightarrow -\infty} f(y_i) = \infty$ and $\lim_{y_i \rightarrow \infty} f(y_i) = -\infty$. Therefore, there exists a single crossing point $\bar{\tau}$ such that $f(y_i) > 0$ for $y_i < \bar{\tau}$ and $f(y_i) < 0$ for $y_i > \bar{\tau}$. \blacksquare

B. Linear benefit functions

Theorem 1 guarantees that the best response to homogeneous thresholds is a threshold policy for a broad class of Global Games. Once a new threshold is found, other agents imitate by using the same best response threshold. We recursively use this scheme, which may converge to a BNE

policy. However, the convergence of such a scheme for an arbitrary benefit function $b(\cdot)$ might not be easy establish, in general. In addition, finding the optimal strategy $a^*(x)$ may not be feasible for general benefit functions. However, such characterization is possible for the class of *linear* benefit functions.

Consider the following linear benefit function indexed by N , $b_N^{\text{lin}} : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$b_N^{\text{lin}}(\xi) \stackrel{\text{def}}{=} \lambda \cdot \left(\frac{\xi}{N} \right), \quad (7)$$

where $\lambda > 0$. Define the *belief* function $\pi_{ij} : \mathbb{R}^2 \rightarrow \mathbb{R}$ as

$$\pi_{ij}(\tau_j, y_i) \stackrel{\text{def}}{=} \mathbb{P}(Y_j \leq \tau | Y_i = y_i).$$

After a few algebraic manipulations, we can write

$$\pi_{ij}(\tau_j, y_i) = \mathbb{E} \left[\Phi \left(\frac{\tau_j - V - \alpha y_i}{\sigma_Z} \right) \right], \quad (8)$$

where V and α are defined in Eq. (6).

Corollary 1 (Corollary to Theorem 1): Assuming a linear benefit function (given by (7)), if each agent $j \neq i$ uses a threshold policy with threshold τ_j , then the best-response to any threshold strategy profile is given by the unique solution \bar{y}_i of $\frac{\lambda}{N} \sum_{i \neq j} \pi_{ij}(\bar{y}_i) = \alpha \bar{y}_i$.

Based on Corollary 1, we can define a BR map in the space of threshold policies, which takes a vector of N thresholds and maps into N thresholds. Let $\mathcal{F} : \mathbb{R}^N \rightarrow \mathbb{R}^N$, where

$$\mathcal{F}_i(\tau_i, \tau_{-i}) \stackrel{\text{def}}{=} \arg \min_{\xi \in \mathbb{R}} \left(\frac{\lambda}{N} \sum_{i \neq j} \pi_{ij}(\tau_j, \xi) - \alpha \xi \right)^2,$$

for all $i \in [N]$.

Remark 1: The existence of a Bayesian Nash-equilibrium in threshold policies is easy to show, but whether it is unique depends on establishing a contraction property of \mathcal{F} , which is a topic for future work.

C. Homogeneous agents using a threshold τ

The problem is simpler when we focus only on homogeneous strategy profiles. In that case, the BR to a threshold strategy profile where every agent uses threshold τ is the unique solution to the following equation

$$\text{BR}(\tau) = \left\{ \xi^* : \lambda \frac{(N-1)}{N} \pi(\xi^*; \tau) = \alpha \xi^* \right\}, \quad (9)$$

here

$$\pi(\xi; \tau) \stackrel{\text{def}}{=} \mathbb{E} \left[\Phi \left(\frac{\tau - \tilde{\sigma}W - \alpha \xi}{\sigma_Z} \right) \right], \quad (10)$$

with $W \sim \mathcal{N}(0, 1)$ and $\tilde{\sigma}^2 = \alpha \sigma_Z^2$.

Example 1: Figure 2 shows the BR function of Eq. (9) and corresponding Nash equilibrium (NE) thresholds for different values of noise variance σ_Z^2 . Two observations from this numerical experiment is that as the noise variance increases, so do the NE thresholds $\tau^*(\sigma_Z^2)$. Less obvious is the limit of the NE threshold as the variance of the noise goes to zero, that is, with perfect observations. In the noiseless case, Fig. 2




Fig. 2. Best response function to a homogeneous strategy profile with threshold τ . Here $N = 10$, $\lambda = 1$, and $\sigma_X^2 = 1$.

shows that $\tau^* = 1/2$, which implies that, in this example, $a^*(x) = \mathbf{1}(x \leq 0.5)$.

First, let us discuss the following properties of the Gaussian CDF whose proofs are omitted due to space limitations.

Lemma 3: Let $V \sim \mathcal{N}(0, 1)$. Then, for any $c \in \mathbb{R}$ and any $\epsilon \in \mathbb{R}$, we have

$$|\mathbb{E}_V[\Phi(cV + \epsilon)] - \frac{1}{2}| \leq |\epsilon|. \quad (11)$$

In particular, $\Phi(cV) = \frac{1}{2}$ for all $c \in \mathbb{R}$.

Using this result, we can show the following important estimate of the fixed point of (10).

Lemma 4: For $\sigma_X^2, \sigma_Z^2 > 0$, let $\tau_N^* = \tau_N^*(\sigma_X^2, \sigma_Z^2)$ be the unique solution to the fixed point equation

$$\lambda \frac{N-1}{N} \pi(\tau_N^*; \tau_N^*) = \alpha \tau_N^*. \quad (12)$$

Then,

$$\frac{1}{2} \leq \left(\frac{\sigma_X^2}{\sigma_X^2 + \sigma_Z^2} \right) \frac{\tau_N^*}{\lambda(N-1)/N} \leq \left(\frac{1}{2} + \frac{\sigma_Z}{\sigma_X} \lambda(N-1)/N \right). \quad (13)$$

Proof: First, note that $0 \leq \Phi(\cdot) \leq 1$ and (10), implies that $0 \leq \pi(\tau_N^*; \tau_N^*) \leq 1$. Therefore, the solution to (12), satisfies,

$$0 \leq \tau_N^* \leq \frac{\lambda(N-1)/N}{\alpha}. \quad (14)$$

Using $\pi(\tau_N^*; \tau_N^*) = \mathbb{E} \left[\Phi \left(\frac{(1-\alpha)\tau_N^* - \tilde{\sigma}W}{\sigma_Z} \right) \right]$, we get

$$\begin{aligned} \frac{1}{2} &\stackrel{(a)}{=} \mathbb{E} \left[\Phi \left(\frac{-\tilde{\sigma}W}{\sigma_Z} \right) \right] \\ &\stackrel{(b)}{\leq} \mathbb{E} \left[\Phi \left(\frac{(1-\alpha)\tau_N^* - \tilde{\sigma}W}{\sigma_Z} \right) \right] \\ &\stackrel{(c)}{\leq} \mathbb{E} \left[\Phi \left(\frac{(1-\alpha)\frac{\lambda(N-1)/N}{\alpha} - \tilde{\sigma}W}{\sigma_Z} \right) \right] \end{aligned}$$

$$\stackrel{(d)}{\leq} \frac{1}{2} + \frac{(1-\alpha)}{\sigma_Z} \lambda(N-1)/N.$$

Here, (a) follows from (11) for $\epsilon = 0$, (b) and (c) follow from Eq. (14), the fact $0 < \alpha < 1$, the monotonicity of expectation, and the monotonicity of $\Phi(\cdot)$, and (d) follows from (11). Using the above inequality and the fact that τ_N^* satisfies (12), we have

$$\frac{1}{2} \leq \alpha \frac{\tau_N^*}{\lambda(N-1)/N} \leq \frac{1}{2} + \frac{(1-\alpha)}{\sigma_Z} \lambda(N-1)/N.$$

Finally, noting $\alpha = \frac{\sigma_X^2}{\sigma_X^2 + \sigma_Z^2}$ concludes the proof. ■

The following result follows immediately from Lemma 4.

Theorem 2 (Diffuse Gaussian priors): Consider a Global Game with the linear benefit function of Eq. (7). Then, for all $\sigma_Z^2 > 0$,

$$\lim_{\sigma_X^2 \rightarrow \infty} \tau_N^* = \frac{\lambda}{2} \cdot \left(1 - \frac{1}{N}\right).$$

Remark 2: In the Economics literature on global games [6]–[8], it is customary to assume a diffuse prior distribution on X . However, from the Engineering perspective, the assumption of a diffuse Gaussian distribution on X with $\sigma_X^2 \rightarrow \infty$ leads to effectively having parallel Gaussian communication channels of infinite capacity³. Therefore, from an information theoretic perspective, the effect of the channel noise becomes negligible, leading to perfect estimates of the input X given Y_i in the mean-squared error sense [18].

Remark 3: A special case of our result is when $N \rightarrow \infty$ and $\lambda = 1$, we get $\tau_\infty^*(1) \rightarrow 1/2$, which is established in Morris and Shin [6].

The significance of Theorem 2 is that it leads to an Bayesian Nash equilibrium threshold policy corresponding to the case where the signals are observed through perfect channels. Remarkably, for a linear benefit function such policy has a closed form. We refer to this as the *oracle* policy, and is defined as:

$$\gamma_{\text{oracle}}^*(x) \stackrel{\text{def}}{=} \mathbf{1}(x \leq \tau_N^*(\sigma_Z^2)). \quad (15)$$

D. Certainty equivalent policies

The knowledge of the optimal policy for $\sigma_X^2 \rightarrow \infty$ motivates the definition of a *certainty equivalent* (CE) policy [19] when $\sigma_X^2 < \infty$ in which the agents first form a estimate of the fundamental using the MMSE estimator of Eq. (3), followed by using the oracle policy in Eq. (15):

$$\gamma_{\text{ce}}(y) \stackrel{\text{def}}{=} \mathbf{1}(\hat{x}_{\text{mmse}}(y) \leq \tau_N^*(\sigma_Z^2)).$$

Thus,

$$\gamma_{\text{ce}}(y) = \mathbf{1}\left(y \leq \left(1 + \frac{\sigma_Z^2}{\sigma_X^2}\right) \cdot \frac{\lambda}{2} \cdot \left(1 - \frac{1}{N}\right)\right).$$




Fig. 3. Diagram showing how to compute the coordination error event between a generic agent and an omniscient agent with access to perfect information playing a stochastic coordination game.

IV. A FUNDAMENTAL LIMIT ON COORDINATION

We now obtain a universal upper bound on the efficiency of *any* policy, regardless of their structure. Our result is based on *Fano's inequality* [20]. Fano's inequality provides a bound on the probability of estimation error of the estimate of a discrete random variable on the basis of side information.

Theorem 3 (Upper bound on coordination efficiency):

For a global game with a linear benefit functions, the coordination efficiency of any homogeneous strategy profile satisfies the following bound

$$\varrho \leq 1 - h^{-1}\left(H(A^*(X) | Y_i)\right),$$

where $H(\cdot | \cdot)$ is the conditional entropy function⁴, and $h^{-1}(\cdot)$ is the inverse of the binary entropy function over the interval $[0, 1/2]$.

Proof: Let $A^*(X) \in \{0, 1\}$ such that

$$A_i^*(X) = \mathbf{1}(X \leq \tau_N^*(\sigma_Z^2))$$

and let $\hat{A}_i(Y) \in \{0, 1\}$ denote any estimate of $A^*(X)$, on the basis of Y_i . Then, notice that the following Markov relation is satisfied

$$A^*(X) \leftrightarrow X \leftrightarrow Y \leftrightarrow \hat{A}_i(Y_i).$$

Considering the block diagram in Fig. 3, define the error random variable

$$E_i \stackrel{\text{def}}{=} \mathbf{1}(\hat{A}_i(Y) \neq A^*(X)),$$

and notice that the probability of making an error when estimating $A^*(X)$ is at most $1/2$. Fano's inequality [20] is a bound on the conditional entropy of the optimal decision computed using the oracle policy given the signal Y_i available to the i -th agent

$$H(A^*(X) | Y_i) \leq h(E_i) + \mathbb{P}(E_i = 1) \log_2(|\mathcal{A}| - 1),$$

³The Shannon capacity of a Gaussian channel is given by

$$C = \frac{1}{2} \log_2 \left(1 + \frac{\sigma_X^2}{\sigma_Z^2}\right) \text{ bits/channel use.}$$

⁴The entropy of a random variable $X \sim f(x)$ is defined as

$$H(X) \stackrel{\text{def}}{=} -\mathbb{E}[\log_2(f(X))].$$

where $|\mathcal{A}|$ is the cardinality of the decision variable A_i . Since out decision variables are binary, we have

$$H(A^*(X) | Y_i) \leq h(E_i). \quad (16)$$

Assuming that we can compute the LHS of Eq. (16), we obtain a bound on $\mathbb{P}(E_i = 1)$, by finding the inverse of the *binary entropy function*⁵ within the interval $[0, 0.5]$. Finally, notice that for a homogeneous strategy profile the coordination efficiency is

$$\varrho = 1 - \mathbb{P}(E_i = 1) \leq 1 - h^{-1}\left(H(A^*(X) | Y_i)\right).$$

A. Computing the bound on coordination efficiency

Using properties of the entropy function, we obtain:

$$H(A^*(X) | Y_i) = H(A^*(X)) - H(Y_i) + H(Y_i | A^*(X)).$$

We proceed to compute each of these three terms: the first is the entropy of the optimal decision variable as computed by the oracle:

$$h(A^*(X)) = h\left(\mathbb{P}\left(X \leq \frac{\lambda}{2}\left(1 - \frac{1}{N}\right)\right)\right),$$

where $h(\cdot)$ denotes the binary entropy function.

The second term is the differential entropy of the signal Y_i , which is a Gaussian random variable with variance $\sigma_X^2 + \sigma_Z^2$. Therefore,

$$H(Y_i) = \frac{1}{2} \log_2(2\pi e(\sigma_X^2 + \sigma_Z^2)).$$

The third term is more challenging must be computed numerically.

$$H(Y_i | A^*(X) = 1) = H\left(Y_i | X \leq \frac{\lambda}{2}\left(1 - \frac{1}{N}\right)\right).$$

To evaluate this entropy, we must use the conditional probability density function

$$f_{Y_i | X \leq \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i) = \frac{\int_{-\infty}^{\frac{\lambda}{2}(1 - \frac{1}{N})} f_Z(y_i - x) f_X(x) dx}{\int_{-\infty}^{\frac{\lambda}{2}(1 - \frac{1}{N})} f_X(x) dx}.$$

Similarly,

$$H(Y_i | A^*(X) = 0) = H\left(Y_i | X > \frac{\lambda}{2}\left(1 - \frac{1}{N}\right)\right).$$

To evaluate this entropy, we must use the conditional probability density function

$$f_{Y_i | X > \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i) = \frac{\int_{\frac{\lambda}{2}(1 - \frac{1}{N})}^{\infty} f_Z(y_i - x) f_X(x) dx}{\int_{\frac{\lambda}{2}(1 - \frac{1}{N})}^{\infty} f_X(x) dx}.$$

Finally, we can compute:

$$H(Y_i | A^*(X)) = H(Y_i | A^*(X) = 0)\mathbb{P}(A^*(X) = 0) + H(Y_i | A^*(X) = 1)\mathbb{P}(A^*(X) = 1),$$

⁵The binary entropy function is defined as

$$h(p) \stackrel{\text{def}}{=} -p \log_2 p - (1-p) \log_2(1-p).$$

where

$$H(Y_i | A^*(X) = 0) = - \int_{\mathbb{R}} f_{Y_i | X > \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i) \times \log_2\left(f_{Y_i | X > \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i)\right) dy_i$$

and

$$H(Y_i | A^*(X) = 1) = - \int_{\mathbb{R}} f_{Y_i | X \leq \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i) \times \log_2\left(f_{Y_i | X \leq \frac{\lambda}{2}(1 - \frac{1}{N})}(y_i)\right) dy_i.$$

Lastly, the computation of the inverse of the binary entropy function can be efficiently performed numerically.

V. NUMERICAL RESULTS

The characterization we have provided thus far assumes that the agents choose their actions according to a policy that ideally tracks the behavior of an omniscient agent that has access to perfect information about the state. Since the agents receive noisy signals about the state, they are not able to perfectly coordinate with the omniscient agent using a threshold policy indexed by τ^* .

Assuming that the agents use a generic homogeneous threshold policy indexed by τ , the probability of miscoordination is given $X = x$ is given by:

$$\mathbb{P}(E_i = 1 | X = x) = \left(1 - \Phi\left(\frac{\tau - x}{\sigma_Z}\right)\right) \mathbf{1}(x \leq \tau^*) + \Phi\left(\frac{\tau - x}{\sigma_Z}\right) \mathbf{1}(x > \tau^*). \quad (17)$$

Therefore,

$$\varrho(\gamma) = 1 - \int_{\mathbb{R}} \mathbb{P}(E_i = 1 | X = x) f(x) dx. \quad (18)$$

Assume a global game with linear benefit function, and a number of agents $N \rightarrow \infty$. The optimal threshold used by the omniscient agent is $\tau^* = \lambda/2$. From the agent standpoint, we consider two strategies: 1. computing the NE threshold for the global game, using the prior information σ_X^2 and σ_Z^2 , and the parameter λ ; 2. estimate the state variable X using a MMSE estimator and using the certainty equivalent policy.

Figure 4 (left) shows the thresholds corresponding to these two types of coordination policies for a system $\sigma_X^2 = 1$ and $\lambda = 1$ as a function of the noise variance σ_Z^2 . We can clearly see how different these two policies are. Moreover, there is also a larger computational cost of solving for the NE in the first strategy, whereas the CE strategy can be obtained in closed form in this case. Figure 4 (center) shows the expected utility of these two strategies. There is a substantial gap between the utilities of an agent using the NE versus CE. This is also clear, because CE in stochastic control and optimization is a suboptimal strategy, in general. More surprisingly is the fact that CE yields a better coordination efficiency, as shown in Fig. 4 (right).

Finally, Fig. 4 (right) also shows the information theoretic upper bound on coordination efficiency for any homogeneous policy profile (not just threshold policies). The significance of this figure is that it establishes that certain coordination




Fig. 4. Numerical results for a global game with linear benefit function with $\lambda = 1$ for a non-diffuse prior with variance $\sigma_X^2 = 1$, and a countably infinite number of agents. The NE and CE thresholds as a function of the noise variance (left); The expected utility as a function of the noise variance (center); The coordination efficiency and its universal lower bound (right).

efficiencies cannot be achieved by any policy for a given level of noise in the communication channel between the gateway and the robotic agents, in a practical application. Therefore, when properly planning for a distributed implementation of a collective task performed by strategic self-interested agents, the system designer needs to communicate at a certain signal to noise ratio, which is not determined by the bit error rate at the receiver, but instead by the level of collective coordination it is interested in achieving.

VI. CONCLUSIONS AND FUTURE WORK

We defined the class of homogeneous coordination games which encompass the popular class of global games. Then, we proposed a Bayesian metric of coordination based on the probabilities that the agents will take the “right” action by aligning their decisions with the ones from an omniscient agent with access to the perfect state of the system. We show that this metric of coordination efficiency can be bounded using information theoretic inequalities, establishing regimes in which certain levels of coordination are impossible to achieve. To the best of our knowledge, this is the first time such methods are used in the context of global games.

Future work on this topic will include design of new learning algorithms (for threshold policies) in the presence of local data at the agents, the presence of partially connected influence graphs on the agent’s benefit functions, and the characterization of better upper bounds on coordination efficiency that would take into account the structure of the policy (e.g. threshold).

REFERENCES

- [1] L. Arditti, G. Como, F. Fagnani, and M. Vanelli, “Equilibria and learning dynamics in mixed network coordination/anti-coordination games,” in *2021 60th IEEE Conference on Decision and Control (CDC)*, 2021, pp. 4982–4987.
- [2] P. Ramazi and M. H. Roohi, “Characterizing oscillations in heterogeneous populations of coordinators and anticordinators,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 4615–4620.
- [3] K. Paarporn, B. Canty, P. N. Brown, M. Alizadeh, and J. R. Marden, “The impact of complex and informed adversarial behavior in graphical coordination games,” *IEEE Transactions on Control of Network Systems*, vol. 8, no. 1, pp. 200–211, 2021.
- [4] K. Paarporn, M. Alizadeh, and J. R. Marden, “A risk-security tradeoff in graphical coordination games,” *IEEE Transactions on Automatic Control*, vol. 66, no. 5, pp. 1973–1985, 2021.
- [5] S. Das and C. Ekin, “Approximate submodularity of maximizing anticoordination in network games,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 3151–3157.
- [6] S. Morris and H. S. Shin, *Global Games: Theory and Applications*, ser. Econometric Society Monographs. Cambridge University Press, 2003, vol. 1, pp. 56–114.
- [7] H. Carlsson and E. Van Damme, “Global games and equilibrium selection,” *Econometrica: Journal of the Econometric Society*, pp. 989–1018, 1993.
- [8] M. A. Dahleh, A. Tahbaz-Salehi, J. N. Tsitsiklis, and S. I. Zoumpoulis, “Coordination with local information,” *Operations Research*, vol. 64, no. 3, pp. 622–637, 2016.
- [9] B. Touri and J. Shamma, “Global games with noisy sharing of information,” in *53rd IEEE Conference on Decision and Control*. IEEE, 2014, pp. 4473–4478.
- [10] H. Mahdaviifar, A. Beirami, B. Touri, and J. S. Shamma, “Global games with noisy information sharing,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 4, no. 3, pp. 497–509, 2017.
- [11] C. M. Leister, Y. Zenou, and J. Zhou, “Social connectedness and local contagion,” *The Review of Economic Studies*, vol. 89, no. 1, pp. 372–410, 2022.
- [12] M. M. Vasconcelos, “Bio-inspired multi-agent coordination games with Poisson observations,” *IFAC-PapersOnLine*, vol. 55, no. 13, pp. 180–185, 2022.
- [13] Y. Wei and M. M. Vasconcelos, “Strategic multi-task coordination over regular networks of robots with limited computation and communication capabilities,” *arXiv preprint arXiv:2212.10968*, 2022.
- [14] E. J. Hoffmann and T. Sabarwal, “Global games with strategic complements and substitutes,” *Games and Economic Behavior*, vol. 118, pp. 72–93, 2019.
- [15] A. Kanakia, B. Touri, and N. Correll, “Modeling multi-robot task allocation with limited information as global game,” *Swarm Intelligence*, vol. 10, no. 2, pp. 147–160, 2016.
- [16] S. Berman, A. Halasz, M. A. Hsieh, and V. Kumar, “Optimized stochastic policies for task allocation in swarms of robots,” *IEEE Transactions on Robotics*, vol. 25, no. 4, pp. 927–937, 2009.
- [17] B. Hajek, *Random Processes for Engineers*. Cambridge University Press, 2015.
- [18] D. Guo, S. Shamai, and S. Verdú, “Mutual information and minimum mean-square error in Gaussian channels,” *IEEE Transactions on Information Theory*, vol. 51, no. 4, pp. 1261–1282, 2005.
- [19] D. Bertsekas, *Dynamic programming and optimal control: Volume I*. Athena scientific, 2012, vol. 1.
- [20] T. M. Cover and J. A. Thomas, *Elements of Information Theory*, 2nd ed., ser. Wiley Series in Telecommunications and Signal Processing. Wiley, 2006.