Quality of Non-Convergent Best Response Processes in Multi-Agent Systems through Sink Equilibria

Rohit Konda   Rahul Chandan   Jason R. Marden

Abstract— Examining the behavior of multi-agent systems is vitally important to many emerging distributed applications - game theory has emerged as a powerful tool set in which to do so. The main approach of game-theoretic techniques is to model agents as players in a game, and predict the emergent behavior through the relevant Nash equilibria. The virtue from this viewpoint is that by assuming that self-interested decision-making processes lead to Nash equilibrium, system behavior can then be captured by Nash equilibrium without studying the decision-making processes explicitly. This approach has seen success in a wide variety of domains, such as sensor coverage, traffic networks, auctions, and network coordination. However, in many other problem settings, Nash equilibria are not necessarily guaranteed to exist or emerge from self-interested processes. Thus the main focus of the paper is on the study of sink equilibria, which are defined as the attractors of these decision-making processes. By classifying system outcomes through a global objective function, we can analyze the resulting approximation guarantees that sink equilibria have for a given game. Our main result is an approximation guarantee on the sink equilibria through defining an introduced metric of misalignment, which captures how uniform agents are in their self-interested decision making. Overall, sink equilibria are naturally occurring in many multi-agent contexts, and we display our results on their quality with respect to two practical problem settings.

I. INTRODUCTION

The analysis of multi-agent systems has received a significant amount of attention recently, primarily due to the emergence of distributed structures in wireless communication, biology, IOT, and many other application domains. This has led to a rich variety of theoretical approaches [1], [2], [3]. In this paper, we specifically consider a *game theoretic* approach, where the emergent properties of the multi-agent system are studied using tools from game theory.

The main idea of this approach is to model agents as self-interested decision makers, where each agent's preference over the collective system outcome is designated through a *utility function*. The agents are then presumed to undergo a *best (or better) response process*, where each agent updates their decision in a self-interested manner to maximize their individual utility. The emergent system outcomes from this process are traditionally expected to be *Nash equilibrium*, which can be expressed as the limit points of the best response process. Thus, many previous works study the emergent system behavior through characterizing properties of the Nash equilibrium; this has been done in many different contexts, such as traffic systems, power networks, etc. [4], [5].

But is it reasonable to expect agents to converge to Nash equilibria from best response processes? In the class of *potential games* [6] and variants [7], best response processes are indeed guaranteed to converge to Nash equilibria [8]. In potential games, agents are fully cooperative, and self-interested decisions made by agents always lead to improvements in a given global objective. However, a variety of natural multi-agent settings fall outside of this class. The system may display competitive interactions; for e.g., business firms may have competing economic interests. In social systems, agents' may be inherently misaligned in their preferences; for e.g., drivers may have different sensitivities to tolls. Even for multi-agent systems that are fully engineered, there are operational concerns such as prediction errors or informational privacy that must be accounted for. All of these scenarios do not exhibit a potential game structure, and thus guarantees of convergence to Nash equilibria can not be established. We describe two examples of this kind in Example 1 and 2. **In these instances, can the emergent system outcomes still be characterized, where either Nash equilibrium do not exist or best response processes does not converge to Nash equilibria?** This is the main focus of this paper.

We utilize *sink equilibria* in [9] as an alternate solution concept [1] to address this, which are specified as the attractors of the best response process. By definition, sink equilibria are well defined and have guarantees of convergence for any given game. Thus, we analyze the behavior of sink equilibria in this paper. Specifically, we assume that system outcomes are evaluated through a given global objective function. The utility function of each agent may not be aligned with this global objective. In these settings, we can characterize performance guarantees of the induced sink equilibria with respect to the global objective - this is the main result of the paper.

To the author's knowledge, a general approach to studying performance guarantees of sink equilibrium has not been done previously. While sink equilibria have not been studied in as much detail as Nash equilibria, we still highlight an

---

[1]A popular alternative to Nash equilibria are *coarse correlated equilibria* which have existence and convergence guarantees [10]. However, these dynamics requires full knowledge of the history of decisions of all of the agents. Additionally, only empirical distribution of play is guaranteed to the converge to coarse correlated equilibria, which may not correspond to the actualized decisions made by the agents.

pertinent selection of past literature on sink equilibria. The seminal work in [9] first established the concept of sink equilibria in a game-theoretic context, as well as provide negative results of sink equilibria in valid utility games. Positive results on sink equilibria behavior were shown in [11], where it was seen that sink equilibria perform much better than mixed Nash equilibria. Analysis of sink equilibria under the name of curb sets was done in [12]. Several complexity results of sink equilibria were introduced in [13]. Sink equilibria were extended in continuous domains in [14]. Recently, design of sink equilibria selection algorithms was done in [15]. While the literature on sink equilibrium is sparse, they naturally emerge when agent utility functions are not aligned perfectly.

The structure of the paper is as follows. In Section II, we define our game theoretic setup and formally define the best response process and sink equilibrium. Additionally, we introduce two application scenarios where sink equilibria appear. In Section III, we state our main results, where we derive performance guarantees of sink equilibria in our setting with regards the approximation ratio of *price of sinking*. A discussion of sink equilibria induced by better response processes is provided in Section IV. Finally, we conclude in Section V. We include certain proofs in an arxived version of the paper.

II. Preliminaries

Consider a general multi-agent scenario with n agents $\mathcal{I} = \{1, \ldots, n\}$, where each agent is endowed with a finite decision or action set $\mathcal{A}_i$. We denote an action as $a_i \in \mathcal{A}_i$, and a joint action profile as $a \in \mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. The quality of each joint action profile is evaluated with a global objective function $\mathrm{W} : \mathcal{A} \to \mathbb{R}_{\geq 0}$ that characterizes the total system welfare. In other words, the optimal joint decision of the system is described as

$$a^{\mathrm{opt}} \in \arg\max_{a \in \mathcal{A}} \mathrm{W}(a). \tag{1}$$

We assume that solving the decision problem in Eq. (1) cannot be done in a centralized fashion, due to computational, informational, administrative, or communication concerns. Therefore, we assume that each agent selects their decisions in a distributed manner. We assume that each agent i is endowed with a *utility function* $\mathrm{U}_i : \mathcal{A} \to \mathbb{R}$ to classify their preferences over their decision set, resulting in the game tuple $\mathrm{G} \triangleq (\mathcal{I}, \mathcal{A}, \mathrm{W}, \{\mathrm{U}_i\}_{i \in \mathcal{I}})$. Moreover, we assume the utility function U_i depends on the welfare function W, either naturally or by design. In particular, consider a completely cooperative scenario, where the agents exhibit the *common interest* utility design. Here, the utility functions are completely aligned with the global objective W [2], where

$$\mathrm{U}_i(a) \equiv \mathrm{W}(a) \text{ for all } a \in \mathcal{A} \text{ and } i \in \mathcal{I}. \tag{2}$$

---

[2]We can relax this constraint to consider utility functions that are *preference equivalent* to the welfare function, where the ordering of preferences over joint actions is maintained. This is indeed the case for potential and weighted-potential games.

Under the common interest utility, the emergent joint decisions coincide with the set of Nash equilibrium NE of the game. A joint action a^{ne} is considered a Nash equilibrium if the following inequality holds

$$\mathrm{U}_i(a^{\mathrm{ne}}) \geq \mathrm{U}_i(a_i, a^{\mathrm{ne}}_{-i}) \text{ for all } a \in \mathcal{A} \text{ and } i \in \mathcal{I}, \tag{3}$$

where $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ corresponds to the joint action without the decision of agent i. In completely cooperative settings, where they are guaranteed to exist, we can then quantify the emergent behavior through the qualities of the possible resulting Nash equilibrium. This is done through the metric of *price of anarchy* which is defined as

$$\mathrm{PoA(G)} = \frac{\min_{a^{\mathrm{ne}} \in \mathrm{NE}} \mathrm{W}(a^{\mathrm{ne}})}{\max_{a \in \mathcal{A}} \mathrm{W}(a)}, \tag{4}$$

where we take the worst case ratio of the welfare of Nash equilibria over the optimal welfare. Price of anarchy is a well-studied metric, with many results on its characterization in the literature [16], [17], [4]. Thus a standardized approach can be implemented to characterize system behavior in completely cooperative scenarios.

However, due to operational or design constraints, assuming a common interest utility design may not be feasible (see Examples 1 and 2). Throughout the paper, we then allow $\mathrm{U}_i(a) \neq \mathrm{W}(a)$ to be misaligned, and focus our attention to sink equilibrium as our standard solution concept. To define sink equilibrium, we first outline the *best response process*. Under this process, the best decision set for agent i assuming all other agents' decisions are fixed is known as the *best response set*, that is,

$$\mathrm{Br}_i(a) = \arg\max_{\bar{a}_i \in \mathcal{A}_i} \mathrm{U}_i(\bar{a}_i, a_{-i}). \tag{5}$$

Then for every step, a randomly selected agent picks an action from its best response set uniformly. This induces the following Markovian dynamics on the set of joint actions, describing the best response process, as

$$\Pr(\tilde{a}|a) = \begin{cases} \frac{1}{n \cdot |\mathrm{Br}_i(a)|} & \text{if } \tilde{a} \in (\mathrm{Br}_i(a), a_{-i}) \text{ for some } i \in \mathcal{I} \\ 0 & \text{otherwise}, \end{cases} \tag{6}$$

where $\Pr(\tilde{a}|a)$ represents the probability of reaching the joint action $\tilde{a}$ from a in the Markov chain. Note that there is an equal chance for each player i to perform a best response at each time step. We refer to a probability distribution over the joint action set as $\sigma \in \Delta\mathcal{A}$, where $p_\sigma(a)$ denotes the probability of sampling a under σ. We also say that an action $a \in \mathrm{supp}(\sigma) \subseteq \mathcal{A}$ is in the support of σ if the probability $p_\sigma(a) > 0$ is strictly positive. Furthermore, we say that σ is a *stationary distribution* of the Markov chain if the equality $p_\sigma(a) = \mathbb{E}_{\bar{a} \sim \sigma}[\Pr(\tilde{a}|\bar{a}) \cdot p_\sigma(\bar{a})]_a$ for all $a \in \mathcal{A}$ holds. We also specify the *sink strongly connected components* of the Markov chain defined in Eq. (6). A set $S \subseteq \mathcal{A}$ is a sink strongly connected component if there is exists a path of positive probability from a to $\bar{a}$ under (6) for any $a, \bar{a} \in S$ and there are no transitions from S to outside of S. Formally, if $a, \bar{a} \in S$, then there exists a sequence $a_0, a_1, \ldots, a_m$ where $\Pr(a_{j+1}|a_j) > 0$ for all j with $a_0 = a$ and $a_m = \bar{a}$.

Additionally, if $a \in S$ and $\bar{a} \notin S$, then no such sequence exists.

Definition 1. A probability distribution $\sigma \in \Delta\mathcal{A}$ is a sink equilibrium of the game G if σ is a stationary distribution of the Markov chain in Eq. (6) and if $\text{supp}(\sigma) = S$ is a sink strongly connected component.

In other words, sink equilibria are defined as the attractors of the best response process. Given a game G, we characterize the behavior of the sink equilibria in a similar manner to Eq. (4) through the metric of *price of sinking* [9] as

$$\text{PoSE(G)} = \min_{\sigma \in \text{SE}} \mathbb{E}_{a \sim \sigma}[\text{W}(a)]/\text{W}(a^{\text{opt}}), \qquad (7)$$

where SE denotes the total set of sink equilibrium of the game G. We note that under a common interest utility design, the set of Nash equilibrium $\text{NE} \simeq \text{SE}$ is equivalent to the set of sink equilibrium.

We now examine the sink equilibria on two illustrative examples in which they naturally emerge. In these examples, we derive guarantees on the sink equilibria using the tools discussed in Section III. The examples are as follows.

Example 1 (Ecological Monitoring)**.** Ecological monitoring is necessary to understand the well-being of inhabited populations as well as track health of overall ecology. While this can be handled by field ecologists, autonomous agents can supplement or even act as substitutes to gather important ecological data - as was done by the authors in [18]. In this scenario, a high level control objective of the agents is to understand how to orient themselves to monitor the region of interest as best as possible. We can model this as a covering problem [19]. In this way, let $\mathcal{R} = \{r_1, \ldots, r_m\}$ define the region monitored by n agents, where we have finitely partitioned the region into possibly arbitrary segments. For each segment r, the importance of monitoring that segment can be associated with a value $v_r \in \mathbb{R}_{\geq 0}$ which defines the intrinsic quality of data that can be collected in that segment. This can be affected by the number or magnitude of populations in the segment, relevant climactic conditions, etc. These parameters are never known before hand, and thus must be estimated by the agents in the field. Thus each agent has its noisy estimate of the value v_r^i which we assume is drawn from a normal distribution $\mathcal{N}[v_r + c, (d \cdot v_r)^2]$ with bias c and variance $(d \cdot v_r)^2$. Each agent can decide which subset of region to monitor (i.e. $a_i \subset \mathcal{R}$), which depend on its sensor and motor capabilities. Thus, collectively, the goal of the agents is to monitor the most and highest valued portions, as captured by the welfare function below.

$$\text{W}(a) = \sum_{r \in \cup_i a_i} v_r \qquad (8)$$

The objective that each agent witnesses, however is based on their estimate, or that their utility is $\text{U}_i(a) = \sum_{r \in \cup_i a_i} v_r^i$, which is potentially different for each agent. Thus, when running a best response algorithm, the agents converge to a sink equilibrium which may not be a Nash equilibrium. We derive a guarantee on price of sinking in Proposition 1.

Proposition 1. *Consider the problem defined above. The expected price of sinking is lower bounded by the following expression*

$$\mathbb{E}[\text{PoSE(G)}] \geq \max(\frac{1 - 4n\beta_\Phi}{2}, 0), \qquad (9)$$

where Φ is the normal cumulative distribution function and β_Φ is defined as

$$\beta_\Phi = |\mathcal{R}| \left(d\sqrt{\frac{2}{\pi}} e^{-\frac{c^2}{2d^2}} + c(1 - 2\Phi(-c/d)) \right). \qquad (10)$$

Example 2 (Radio Signalling)**.** Consider the situation in which n agents have to communicate to each other through k communication channels, as seen in [20]. However, when more than one agent selects a channel to communicate under, the signals experience interference. This can captured by the parameter $w_{ij} \geq 0$, which dictates the interference between agent i and agent j. The system as a whole, would like to minimize the total interference experienced between all the agents. In turn, the system welfare can be defined as

$$\text{W}(a) = \sum_i \sum_{j:a_j \neq a_i} w_{ij}, \qquad (11)$$

where $a_i \in \{1, \ldots k\}$ is the channel that agent i decides to transmit their messages on. However, these interference parameters may not be known to the agents and have to be estimated. For simplicity, we can assume that the agents have a margin of error of α or that agent i's estimate is $w_{ij}^i \in [\alpha \cdot w_{ij}, \alpha^{-1} \cdot w_{ij}]$. Again, when agents run a best response algorithm, they are not guaranteed to converge to a Nash equilibrium due to the informational constraints. Then we can characterize the guarantee on the price of sinking in Proposition 2.

Proposition 2. *Consider the problem defined above with two channels ($k = 1$). The price of sinking is lower bounded by the following expression*

$$\text{PoSE(G)} \geq \frac{1}{3\alpha^2 + (1 - \alpha^2)n}. \qquad (12)$$

III. MAIN RESULTS

The main results of this paper are on providing lower bounds for the price of sinking for a given game. To do this, we first recall the notion of *smoothness* [16] as a useful analysis tool for the price of anarchy. In this paper, we use a relaxed version of smoothness to classify a given game. We say that a game is (λ, μ)-*smooth* if, for a fixed $\mu \geq \lambda \geq 0$, we have that

$$\sum_{i \in \mathcal{I}} \left(\text{U}_i(a) - \text{U}_i(a_i^{\text{opt}}, a_{-i}) \right) \leq \mu\text{W}(a) - \lambda\text{W}(a^{\text{opt}}) \quad (13)$$

for all actions $a \in \mathcal{A}$. Given these parameters, the efficiency of Nash equilibria for a given game can easily be determined. This is described in Proposition 3, where a proof is included for completeness.

Proposition 3. *Let G be a (λ, μ)-smooth game. Then the price of anarchy is lower bounded by $\text{PoA(G)} \geq \frac{\lambda}{\mu}$.*

Proof. From applying the definition of Nash equilibrium repeatedly for all $i \in \mathcal{I}$ with respect to the deviation a^{opt}, we have the following inequality

$$\sum_i \left(U_i(a^{\text{ne}}) - U_i(a_i^{\text{opt}}, a_{-i}^{\text{ne}}) \right) \geq 0.$$

Notice that now we can directly substitute the inequality in Eq. (13) to get that

$$\mu W(a^{\text{ne}}) - \lambda W(a^{\text{opt}}) \geq 0.$$

Since a^{ne} is any arbitrary Nash equilibrium, we can rearrange the above equation to get that $\text{PoA}(G) \geq \frac{\lambda}{\mu}$ to show the claim. $\square$

We note that there always exist some λ and μ such that the game is (λ, μ)-smooth, as $\lambda \to 0$ and $\mu \to \infty$ will always satisfy the inequality in Eq. (13). The main analytical benefit of smoothness analysis is that instead of searching across the set of Nash equilibrium directly, we can instead characterize the price of anarchy through a bi-variable optimization problem (over λ and μ). This can be done in various game-theoretic contexts [16]. The optimization problem is written formally below.

Corollary 1. *The price of anarchy for a given game* G *is lower bounded by*

$$\text{PoA}(G) \geq \sup_{\mu \geq \lambda \geq 0} \{ \frac{\lambda}{\mu} : G \text{ is } (\lambda, \mu) - \text{smooth} \}.$$

However, unlike Nash equilibria, when applying the smoothness inequality directly to the analysis of sink equilibria, it is not possible to get non-trivial guarantees. For all valid smoothness parameters, it is possible to construct a corresponding game with trivial performance guarantees on sink equilibria, as stated below.

Proposition 4. *For every $\mu > \lambda \geq 0$, there exists a (λ, μ)-smooth game* G *with a unique sink equilibrium such that the price of sinking is* $\text{PoSE}(G) = 0$.

Proof. Let $\mu > \lambda \geq 0$. Consider the following game G with two agents with the action sets $\mathcal{A}_1 = \{e_1, e_2, e_3\}$ and $\mathcal{A}_2 = \{f_1, f_2, f_3\}$. We define the welfare values $W(a)$ for each joint action through Table I below. Similarly, we can

	f_1	f_2	f_3
e_1	1	$(\lambda + \varepsilon)/\mu$	0
e_2	$(\lambda + \varepsilon)/\mu$	0	0
e_3	0	0	0

TABLE I: Welfare $W(a)$ for each joint action $a = (e_i, f_j)$.

define the utility values $U_i(a)$ for each joint action and for each agent in Table II.

	f_1	f_2	f_3
e_1	$(0, 0)$	$(0, \varepsilon)$	$(0, -\varepsilon)$
e_2	$(\varepsilon, 0)$	$(\lambda, -2\lambda)$	$(-2\lambda, \lambda)$
e_3	$(-\varepsilon, 0)$	$(-2\lambda, \lambda)$	$(\lambda, -2\lambda)$

TABLE II: Welfare $(U_1(a), U_2(a))$ for each joint action a.

We can choose $\varepsilon = (\mu - \lambda)/2 > 0$ such that the optimal joint action is $a^{\text{opt}} = (e_1, f_1)$ with an optimal welfare of $W(a^{\text{opt}}) = 1$. Under the best response dynamics, observe that the set $\{(e_2, f_2), (e_3, f_2), (e_2, f_3), (e_3, f_3)\}$ is the unique strongly connected component. It can be verified that each joint action a satisfies the smoothness condition in Eq. (13). Since the welfare of each action in the unique strongly connected component is 0, the price of sinking can be upper bounded by $\text{PoSE}(G) = \mathbb{E}_{a \in \sigma}[W(a)] \leq \sum_{a \in \text{supp}(\sigma)} W(a) = 0$ for the unique sink equilibrium σ. $\square$

This negative result is similar in spirit to the one presented in [9, Lemma 3.2]. However, we emphasize that the inferior guarantees are more indicative of inefficacy of a direct approach rather than the intrinsic behavior of sink equilibria. This sentiment is also reflected in [11], where in certain game settings, it is shown that the quality of sink equilibria is arbitrarily better than the quality of any mixed equilibria. In fact, if we consider games with added structure, we can arrive at nontrivial guarantees on sink equilibria.

Therefore, we consider games in which the deviation from the common interest utility $U_i(a) \neq W(a)$ is bounded. We encapsulate the extent of the deviation through the constant $\beta \in [0, 1]$, where $\beta = 0$ signifies no deviation from the common interest utility and $\beta = 1$ signifies the maximum deviation. We define this formally below [3].

Definition 2. *A game* G *is considered to be β-arithmetically misaligned if*

$$|U_i(a) - W(a)| \leq \beta W(a), \tag{14}$$

or β-geometrically misaligned if

$$1 - \beta \leq \frac{U_i(a)}{W(a)} \leq \frac{1}{1 - \beta}, \tag{15}$$

is satisfied for all actions $a \in \mathcal{A}$ and agents $i \in \mathcal{I}$.

We note that when $\beta = 0$, under the common interest utility, the sink equilibria are equivalent to the Nash equilibria and inherit the price of anarchy guarantees coming from smoothness analysis. Likewise, we will see that the sink equilibria in near-common interest games with β close to 0 inherit similar guarantees dictated by the common interest utility. This observation is also reflected in a different context in [22]. In this vein, let λ_c and μ_c be the parameters that satisfy the smoothness inequality in Eq. (13) for the common interest utility

$$\sum_{i \in \mathcal{I}} \left(W(a) - W(a_i^{\text{opt}}, a_{-i}) \right) \leq \mu_c W(a) - \lambda_c W(a^{\text{opt}}), \tag{16}$$

where we have substituted $U_i(a) \equiv W(a)$. With this, we can state the main result of the paper.

Theorem 1. *Let* G *be a game such that the best response* $\text{Br}_i(a)$ *is always singular valued. Let λ_c and μ_c satisfy*

---

[3]We can generalize the results to instead consider alignment to a potential function. In this way, β characterizes the closeness of the game to a potential game. Near-potential games have been studied in [21].

Eq. (16) *for all* $a \in \mathcal{A}$. *If the game is* β-*arithmetically misaligned, as in Eq.* (14)*, then the price of sinking satisfies*

$$\text{PoSE(G)} \geq \max(\frac{\lambda_c - 4\beta n}{\mu_c}, 0). \tag{17}$$

If the game is β-*geometrically misaligned, as in Eq.* (15)*, then the price of sinking satisfies*

$$\text{PoSE(G)} \geq \frac{\lambda_c}{(1-\beta)^2 \mu_c + (1-(1-\beta)^2)n}. \tag{18}$$

Proof. First, we introduce the following lemma to characterize the sink equilibria in an alternative fashion.

Lemma 1. *Let* G *be a game such that* $\text{Br}_i(a)$ *is always singular valued and let* $\sigma \in \text{SE}$ *be any sink equilibrium of the game. For any function* $g : \mathcal{A} \to \mathbb{R}$, *the following equality must hold*

$$\mathbb{E}_{a \sim \sigma}[\sum_{i \in \mathcal{I}} g(a) - g(\text{Br}_i(a), a_{-i})] = 0. \tag{19}$$

Proof. Let σ be a sink equilibrium of the game G. Since the sink equilibrium is a stationary distribution under the dynamics outlined in Eq. (6), we have that $p_\sigma(a) = \sum_{\bar{a} \in \mathcal{A}} \text{Pr}(a|\bar{a}) p_\sigma(\bar{a})$. Under this statement, we have the series of equalities below

$$n \sum_a p_\sigma(a) g(a) = n \sum_a \sum_{\bar{a}} \text{Pr}(a|\bar{a}) p_\sigma(\bar{a}) g(a)$$

$$\mathbb{E}_{a \sim \sigma}[\sum_{i \in \mathcal{I}} g(a)] = \mathbb{E}_{\bar{a} \sim \sigma}[n \sum_a \text{Pr}(a|\bar{a}) g(a)]$$

$$= \mathbb{E}_{a \sim \sigma}[\sum_{i \in \mathcal{I}} g(\text{Br}_i(a), a_{-i})],$$

where we change the the naming convention from $\bar{a}$ to a in the last line. Rearranging the terms and using linearity of expectation gives us the claim. □

Proof of Arithmetic. For ease of notation, let $a_{\text{br}}^i = (\text{Br}_i(a), a_{-i})$. We can apply Lemma 1 with respect to the welfare function W to get

$$\mathbb{E}_{a \sim \sigma}[\sum_{i \in \mathcal{I}} W(a) - W(a_{\text{br}}^i)] = 0. \tag{20}$$

Since we assume the game is β-arithmetically misaligned, we have that $U_i(a) \geq (1-\beta)W(a) \geq W(a) - \beta W(a^{\text{opt}})$. Likewise, we can also bound $U_i(a_{\text{br}}^i) \leq W(a_{\text{br}}^i) + \beta W(a^{\text{opt}})$. We can substitute these two inequalities in Eq. (20) to get

$$W(a) - W(a_{\text{br}}^i) \leq U_i(a) - U_i(a_{\text{br}}^i) + 2\beta W(a^{\text{opt}}) \tag{21}$$

We can apply this inequality to Eq. (20) for

$$\mathbb{E}_{a \sim \sigma}[2\beta n W(a^{\text{opt}}) + \sum_{i \in \mathcal{I}} U_i(a) - U_i(a_{\text{br}}^i)] \geq 0.$$

Further, observe that since $U_i(a_{\text{br}}^i) \geq U_i(a_i^{\text{opt}}, a_{-i})$ from the definition of a best response, we can replace $U_i(a_{\text{br}}^i)$ $U_i(a_i^{\text{opt}}, a_{-i})$. We can utilize the β-misalignment and substitute for the utility functions using Eq. (21) to arrive at

$$\mathbb{E}_{a \sim \sigma}[4\beta n W(a^{\text{opt}}) + \sum_{i \in \mathcal{I}} W(a) - W(a_i^{\text{opt}})] \geq 0.$$

Applying the definition of λ_c and μ_c as in Eq. (16) results in the final inequality.

$$\mathbb{E}_{a \sim \sigma}[4\beta n W(a^{\text{opt}}) + \mu_c W(a) - \lambda_c W(a^{\text{opt}})] \geq 0.$$

Notice that the above inequality holds for any arbitrary sink equilibrium σ. Thus rearranging terms and using linearity of expectation gives us the price of sinking guarantee in Eq. (17). □

Proof of Geometric. We can apply Lemma 1 with respect to the welfare function W to get Eq. (20). For ease of notation, let $\bar{\beta} = 1 - \beta$. We can successively apply the geometric misalignment property in Eq. (15), as well as using the fact that $U_i(a_{\text{br}}^i) \geq U_i(a_i^{\text{opt}}, a_{-i})$, to arrive at the following set of inequalities.

$$W(a_{\text{br}}^i) \geq \bar{\beta} U_i(a_{\text{br}}^i) \geq \bar{\beta} U_i(a_i^{\text{opt}}, a_{-i}) \geq \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})$$

Substituting these inequalities back into Eq. (20) gives

$$\mathbb{E}_{a \sim \sigma}[\sum_i W(a) - \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})] \geq 0.$$

Now we can substitute the definitions of λ_c and μ_c in Eq. (16) to a portion of the terms and simplify to get

$$\mathbb{E}_{a \sim \sigma}[n(1 - \bar{\beta}^2)W(a) + \bar{\beta}^2 (\mu_c W(a) - \lambda_c W(a^{\text{opt}}))] \geq 0.$$

Notice that the above inequality holds for any arbitrary sink equilibrium σ. Thus rearranging terms and using linearity of expectation gives us the price of sinking guarantee in Eq. (18). □

Thus we have shown the bounds for both geometric and arithmetic misalignment cases. □

We see that when the utility functions are close to the common interest utility design with $\beta \sim 0$, the price of sinking guarantees match the guarantees for the common interest utility. We note that while our approach allows us to get nontrivial guarantees on the sink equilibria, we still suffer from the degradation of the guarantee as the number of agents $n \to \infty$ increases arbitrarily. However, we assume worst case deviations (see $W(a_{\text{br}}^i) \geq \bar{\beta}^2 W(a_i^{\text{opt}}, a_{-i})$) for all actions in the game, which is not true for most natural games and produces a conservative bound. Thus the focus of future work will be to address this concern to get tighter guarantees. We can also get alternative guarantees if we consider sink induced by better responses. This is discussed in the next section.

IV. SINK EQ. FROM BETTER RESPONSES

In this section we consider sink equilibria that are induced by a better (rather than best) response process. In contrast to the best response set in Eq. (5), we consider the *better response set* defined as

$$\text{br}_i(\bar{a}) = \{a_i \in \mathcal{A}_i : U_i(a_i, \bar{a}_{-i}) \geq U_i(\bar{a})\} \tag{22}$$

for a given agent i. The better response process is then defined by a random walk, similar to Eq. (6) as

$$\Pr(\tilde{a}|a) = \begin{cases} \frac{1}{n \cdot |\mathrm{br}_i(a)|} & \text{if } \tilde{a} \in (\mathrm{br}_i(a), a_{-i}) \text{ for some } i \in \mathcal{I} \\ 0 & \text{otherwise.} \end{cases}$$
(23)

The sink equilibria are similarly defined for these dynamics. If we consider sink equilibria that are induced by better responses, it is possible to get positive guarantees on the behavior of sink equilibrium. More specifically, we show that there always exists a joint action in the support of the sink equilibria that has similar guarantees to the Nash equilibria.

Proposition 5. *Let* G *be* (λ, μ)-*smooth. Every sink equilibrium in* G *induced by better responses contains a joint action* $\tilde{a} \in \mathrm{supp}(\sigma)$ *in its support such that* $\mathrm{W}(\tilde{a}) \geq \frac{\lambda}{\mu}\mathrm{W}(a^{\mathrm{opt}})$.

Proof of Proposition 5. We show that $\mathrm{W}(a) \geq \frac{\lambda}{\mu}\mathrm{W}(a^{\mathrm{opt}})$ for some $a \in \mathrm{supp}(\sigma)$ in the sink induced by better responses. We first claim that for any sink σ, there exists a joint action $a \in \mathrm{supp}(\sigma)$ such that for all $i \in \mathcal{I}$,

$$\mathrm{U}_i(a) - \mathrm{U}_i(a_i^{\mathrm{opt}}, a_{-i}) \geq 0. \tag{24}$$

Consider an arbitrary action $a \in \mathrm{supp}(\sigma)$ in which the condition does not hold true. Let i be the smallest number such that a does not satisfy Eq. (24) for agent i. Then the action $\hat{a} = (a_i^{\mathrm{opt}}, a_{-i})$ is a better response to a and therefore $\hat{a} \in \mathrm{supp}(\sigma)$ is in the support of σ as well. Note that $\hat{a}$ also satisfies Eq. (24) for agent i. By induction, we can then derive an action $a^* \in \mathrm{supp}(\sigma)$ such that Eq. (24) is satisfied for all i. By the smoothness inequality in Eq. (13), we have that

$$\mu\mathrm{W}(a^*) - \lambda\mathrm{W}(a^{\mathrm{opt}}) \geq \sum_i \mathrm{U}_i(a^*) - \mathrm{U}_i(a_i^{\mathrm{opt}}, a_{-i}^*) \geq 0.$$

Therefore for some $a^* \in \mathrm{supp}(\sigma)$, the efficiency is lower bounded by $\mathrm{W}(a^*) \geq \frac{\lambda}{\mu}\mathrm{W}(a^{\mathrm{opt}})$. $\qquad\square$

V. Conclusion

The main focus of this paper is on studying the limiting behavior of self-interested agents when it is not guaranteed that they converge to Nash equilibria. As such, we consider sink equilibria as our main solution concept, which is naturally defined as a limiting distribution of the best response process. To quantify performance of the sink equilibrium of a game, we examine the price of sinking as an appropriate approximation metric. For Nash equilibrium, smoothness has been classically used to derive guarantees for price of anarchy, the respective approximation metric. However, we show that smoothness is not enough to guarantee anything for the price of sinking. But by using a novel method of misalignment parametrization, this work provides nontrivial guarantees on the sink equilibria. We implement our guarantees into two natural application settings and provide bounds on the resulting sink equilibria. Future work consists of considering additional settings where sink equilibria may appear and providing tighter approximation guarantees.

References

[1] M. Kubisch, H. Karl, A. Wolisz, L. C. Zhong, and J. Rabaey, "Distributed algorithms for transmission power control in wireless sensor networks," in *2003 IEEE Wireless Communications and Networking, 2003. WCNC 2003.*, vol. 1, pp. 558–563, IEEE, 2003.

[2] O. Babaoglu, G. Canright, A. Deutsch, G. A. D. Caro, F. Ducatelle, L. M. Gambardella, N. Ganguly, M. Jelasity, R. Montemanni, A. Montresor, *et al.*, "Design patterns from biology for distributed computing," *ACM Transactions on Autonomous and Adaptive Systems (TAAS)*, vol. 1, no. 1, pp. 26–66, 2006.

[3] M. Stolpe, "The internet of things: Opportunities and challenges for distributed data analysis," *Acm Sigkdd Explorations Newsletter*, vol. 18, no. 1, pp. 15–34, 2016.

[4] T. Roughgarden, *Selfish routing and the price of anarchy*. MIT press, 2005.

[5] I. Atzeni, L. G. Ordóñez, G. Scutari, D. P. Palomar, and J. R. Fonollosa, "Noncooperative and cooperative optimization of distributed energy generation and storage in the demand-side of the smart grid," *IEEE transactions on signal processing*, vol. 61, no. 10, pp. 2454–2472, 2013.

[6] D. Monderer and L. S. Shapley, "Potential games," *Games and economic behavior*, vol. 14, no. 1, pp. 124–143, 1996.

[7] J. R. Marden, G. Arslan, and J. S. Shamma, "Regret based dynamics: convergence in weakly acyclic games," in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, (New York, NY, USA), pp. 1–8, Association for Computing Machinery, 2007.

[8] B. Swenson, R. Murray, and S. Kar, "On best-response dynamics in potential games," *SIAM Journal on Control and Optimization*, vol. 56, no. 4, pp. 2734–2767, 2018.

[9] M. Goemans, V. Mirrokni, and A. Vetta, "Sink equilibria and convergence," in *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, (USA), pp. 142–151, IEEE, IEEE, 2005.

[10] S. Hart and A. Mas-Colell, "A simple adaptive procedure leading to correlated equilibrium," *Econometrica*, vol. 68, no. 5, pp. 1127–1150, 2000.

[11] R. D. Kleinberg, K. Ligett, G. Piliouras, and É. Tardos, "Beyond the nash equilibrium barrier.," in *ICS*, (USA), pp. 125–140, ICS, 2011.

[12] K. Basu and J. W. Weibull, "Strategy subsets closed under rational behavior," *Economics Letters*, vol. 36, no. 2, pp. 141–146, 1991.

[13] A. Fabrikant and C. H. Papadimitriou, "The complexity of game dynamics: Bgp oscillations, sink equilibria, and beyond.," in *SODA*, vol. 8, (USA), pp. 844–853, Citeseer, Society for Industrial and Applied Mathematics, 2008.

[14] C. Papadimitriou and G. Piliouras, "Game dynamics as the meaning of a game," *ACM SIGecom Exchanges*, vol. 16, no. 2, pp. 53–63, 2019.

[15] R. Yan, X. Duan, Z. Shi, Y. Zhong, J. R. Marden, and F. Bullo, "Policy evaluation and seeking for multi-agent reinforcement learning via best response," *IEEE Transactions on Automatic Control*, vol. 67, no. 4, pp. 1898–1913, 2021.

[16] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *Journal of the ACM (JACM)*, vol. 62, no. 5, pp. 1–42, 2015.

[17] D. Paccagnan and J. R. Marden, "Utility design for distributed resource allocation—part ii: Applications to submodular, covering, and supermodular problems," *IEEE Transactions on Automatic Control*, vol. 67, no. 2, pp. 618–632, 2021.

[18] G. Notomista, Y. Emam, and M. Egerstedt, "The slothbot: A novel design for a wire-traversing robot," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 1993–1998, 2019.

[19] V. Chvatal, "A greedy heuristic for the set-covering problem," *Mathematics of operations research*, vol. 4, no. 3, pp. 233–235, 1979.

[20] L. Gourves and J. Monnot, "On strong equilibria in the max cut game," in *International Workshop on Internet and Network Economics*, pp. 608–615, Springer, 2009.

[21] O. Candogan, I. Menache, A. Ozdaglar, and P. A. Parrilo, "Flows and decompositions of games: Harmonic and potential games," *Mathematics of Operations Research*, vol. 36, no. 3, pp. 474–503, 2011.

[22] O. Candogan, A. Ozdaglar, and P. A. Parrilo, "Dynamics in near-potential games," *Games and Economic Behavior*, vol. 82, pp. 66–90, 2013.