Fixed-Time Seeking and Tracking of Time-Varying Nash Equilibria in Noncooperative Games

Jorge I. Poveda, Miroslav Krstić, Tamer Başar

***Abstract*— We study the solution of *time-varying* Nash equilibrium seeking and tracking problems in non-cooperative games via nonsmooth, model-based and model-free algorithms. Specifically, for potential and non-potential games, we derive tracking bounds for the actions of the players with respect to the Nash Equilibrium Trajectory (NET) of the game using the property of fixed-time input-to-state stability. We show that, in the model-based case, traditional pseudo-gradient flows achieve only exponential tracking with a residual error that is proportional to the time-variation of the NET. In contrast, exact and fixed-time tracking can be achieved by using nonsmooth dynamics with discontinuous vector fields. For continuous but non-Lipschitz dynamics, we show that the residual tracking error can be dramatically decreased whenever the learning gains of the dynamics exceed a particular threshold. In the model-free case, we derive similar semi-global practical input-to-state stability bounds using multi-time scale tools for nonsmooth systems.**

I. Introduction

In many cyber-physical engineering systems a group of autonomous agents interact in a competitive manner. Typical examples include smart grids and energy markets, transportation systems, and autonomous multi-agent robots, to name just a few, see [1] and references therein. Under lack of cooperation, agents usually seek to unilaterally minimize their individual cost functions by controlling their own actions. In this scenario, a non-cooperative game emerges between the agents, and a desirable operating point for the multi-agent system is given by a profile of actions in which each agent lacks any incentive to change their strategy, also called a *Nash equilibrium* (NE) [2]. Algorithms able to guarantee convergence to NE are well-studied in the literature, see [3]–[9] and references therein. However, when the cost functions that describe the game are *time-varying*, the standard (static) concept of NE is no longer appropriate. Instead, desirable action profiles become *time-varying trajectories* that agents seek to *track*. Yet, if the dynamics that govern the time-variation of the NE are unknown to the agents (as in most realistic applications), standard feedback control approaches for tracking problems, such as feedforward control and the internal model principle, become unfeasible. Instead, agents will aim to achieve "approximate" tracking of the Nash Equilibrium Trajectory (NET), accepting a residual error related to the magnitude of the time-variation of the NET. While similar online tracking problems have been studied in the literature on optimization [10], [11], and extremum seeking control [7], [12], [13], in the context of time-varying noncooperative games the NET tracking problem has remained mostly unexplored, and the existing results in the literature have been derived mostly for discrete-time algorithms [14] or dynamics with asymptotic tracking properties [15], [16].

In this paper, we study the tracking problem of NETs in non-cooperative games using model-based and model-free nonsmooth algorithms with high transient performance. Our algorithms and tracking bounds leverage the properties of semi-global practical input-to-state stability (ISS) [17], [18] and semi-global practical ultimate boundedness in non-Lipschitz dynamical systems to provide explicit transient and residual bounds for the trajectories of the algorithms. Our results are inspired by recent model-free algorithms with semi-global fixed-time stability properties studied for traditional (model-free) optimization problems in [19] and [20], and for *time-invariant* games in [21], [22]. However, in contrast to these earlier works, in this paper we focus on the *tracking* problem for *time-varying* games, and we pay special attention to the residual tracking errors of the algorithms and their relations to the ISS gains. In particular, we show that in certain cases the ISS gain can be made equal to zero by relaxing the continuity of the vector field that describes the dynamics, and by choosing the gain of the algorithms sufficiently high with respect to the time-variation of the NET. For continuous but non-Lipschitz vector fields, we derive non-zero ISS gains that provide better tracking bounds compared to the traditional linear ISS gains obtained in standard pseudogradient flows. To our best knowledge, the results of this paper are the first that provide tracking bounds for time-varying games using model-free dynamics with fixed-time stability properties.

The rest of this paper is organized as follows: Section II presents the preliminaries, Section III introduces the problem statement, Section IV studies the tracking properties of nominal model-based algorithms, Section V focuses on model-free algorithms, Section VI presents a numerical example, and finally Section VII ends with some conclusions.

II. PRELIMINARIES

In this paper, we will model our algorithms using the framework of constrained dynamical systems [23], where $x \in \mathbb{R}^n$ is the state of a dynamical system evolving as

$$x \in C, \qquad \dot{x} = F(x), \tag{1}$$

Jorge I. Poveda is with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO, 80309. Email:{jorge.poveda@colorado.edu}.

Miroslav Krstić is with the Department of Mechanical and Aerospace Engineering, University of California, La Jolla, CA.

Tamer Başar is with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, IL.

where $C \subset \mathbb{R}^n$ is a closed set, and $F : \mathbb{R}^n \to \mathbb{R}^n$ is a measurable bounded function. For functions F that are discontinuous, system (1) should be replaced by its Krasovskii regularization [23, Def. 4.13]. A solution x to (1) is an absolutely continuous function $x : \text{dom}(x) \to \mathbb{R}^n$ that satisfies: a) $x(0) \in C$; b) $x(t) \in C$, $\forall\, t \in \text{dom}(x)$; and c) $\dot{x}(t) = F(x(t))$ for almost all $t \in \text{dom}(x)$. A solution is said to be complete if $\text{dom}(x) = [0, \infty)$. System (1) is said to render a compact set $\mathcal{A} \subset \mathbb{R}^n$ *uniformly globally asymptotically stable* (UGAS) if there exists a class $\mathcal{KL}$ function β such that every solution of (1) satisfies $|x(t)|_{\mathcal{A}} \leq \beta(|x(0)|_{\mathcal{A}}, t)$, $\forall\, t \in \text{dom}(x)$. When $\exists\, T^* > 0$ such that $\beta(r, s) = 0$ for all $r > 0$ and $s > T^*$ we say that $\mathcal{A}$ is globally fixed-time stable. We also consider ε-perturbed or parameterized dynamical systems of the form $x \in C$, $\dot{x} = F_\varepsilon(x)$, where F_ε is a function parameterized by a positive number $\varepsilon > 0$. For these systems, we say that the compact set $\mathcal{A} \subset C$ is Semi-Globally Practically Asymptotically Stable (SGPAS) as $\varepsilon \to 0^+$, if there exists a class $\mathcal{KL}$ function β such that $\forall\, \delta > \nu > 0$, $\exists\, \varepsilon^* > 0$ such that $\forall\, \varepsilon \in (0, \varepsilon^*)$ every solution x with $|x(0)|_{\mathcal{A}} \leq \delta$ satisfies $|x(t)|_{\mathcal{A}} \leq \beta(|x(0)|_{\mathcal{A}}, t) + \nu$, $\forall\, t \in \text{dom}(x)$. The notion of SGPAS can be extended to systems that depend on multiple parameters $\varepsilon = [\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_\ell]^\top$. We use $\mathbb{S}^1 \subset \mathbb{R}^2$ to denote the unit circle centered at the origin, and $\mathbb{T}^n = \mathbb{S}^1 \times \ldots \times \mathbb{S}^1$ to denote the n^{th}-Cartesian product of $\mathbb{S}^1$.

III. Problem Statement

We consider time-varying noncooperative games characterized by a finite collection of players $i \in \mathcal{V} := \{1, 2, \ldots, n\}$. Each player has an individual action $u_i \in \mathbb{R}$, and a real-valued, smooth cost function $\phi_i(u, \theta_i)$ that depends on the overall vector of actions $u := [u_1, u_2, \ldots, u_n]^\top$, and on a local time-varying parameter $\theta_i : \mathbb{R}_{\geq 0} \to \mathbb{R}$. For simplicity, we consider scalar individual actions and parameters, but our results can easily be extended to the vectorial case. The goal of each agent i is to individually minimize ϕ_i by controlling u_i. Under this noncooperative setting, a desirable operating point for the system is a *Nash Equilibrium Trajectory* (NET), which is a parameterized profile of actions $u^*(t) := h(\theta(t))$, where $h : \mathbb{R}^n \to \mathbb{R}^n$, such that for each given $\theta(t) = [\theta_1(t), \theta_2(t), \ldots, \theta_n(t)] \in \mathbb{R}^n$, the following holds for each $t \geq 0$:

$$\phi_i(u_i^*, u_{-i}^*, \theta_i) = \inf_{u_i \in \mathbb{R}} \phi_i(u_i, u_{-i}^*, \theta_i), \quad \forall\, i \in \mathcal{V}, \qquad (2)$$

where $u_{-i}^* \in \mathbb{R}^{n-1}$ denotes the vector that excludes from u^* the i^{th} component u_i. When the parameters θ_i are time-invariant, i.e., $\dot{\theta}_i = 0$, the NET u^* defined by (2) reduces to a standard static NE. In that case, the *NE seeking problem* is to design dynamics for the actions u_i such that $\lim_{t\to\infty} u(t) = u^*$. However, when $\dot{\theta}_i \neq 0$, *and no information on* $\dot{\theta}_i$ *is available to the agents*, the NE seeking problem becomes a NET tracking problem, where, ideally, the residual tracking error $\lim\ \sup_{t\to\infty} |u(t) - u^*(t)|$ decreases as $|\dot{\theta}_i| \to 0^+$.

To impose some regularity on the time-variation of the NET, we will consider θ-dynamics modeled by exosystems

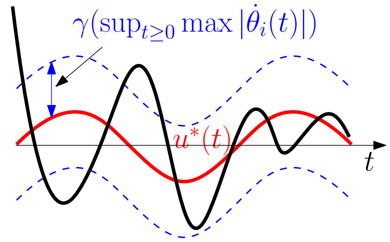

Fig. 1: Approximate tracking of a NET $t \mapsto u^*(t)$, with asymptotic gain $\gamma(\cdot)$.

$$\dot{\theta}_i = \varepsilon_o \Pi_i(\theta_i), \qquad \theta_i \in \Theta_i, \quad \forall\, i \in \mathcal{V}, \qquad (3)$$

where $\varepsilon_o > 0$ is assumed to be the same for all players (otherwise one can work with the maximum of all $\varepsilon_{o,i}$). The following assumption will provide enough regularity to the dynamics (3).

Assumption 1: For all i, the set $\Theta_i \subset \mathbb{R}$ is compact, the mapping $\Pi_i(\cdot)$ is Lipschitz continuous, and the dynamics (3) render the set Θ_i forward invariant. □

To study the tracking properties of our algorithms, we will use the following definition, which is applicable to general NET tracking dynamics with cost functions driven by (3).

Definition 1: The dynamics $(\dot{u}, \dot{\theta})$ are said to have the (β, γ)-*NET tracking property* if $\beta \in \mathcal{KL}$, $\gamma \in \mathcal{K}$, and every solution $t \mapsto (u(t), \theta(t))$ satisfies the bound

$$\begin{aligned} |u(t) - u^*(t)| \leq& \beta(|u(0) - u^*(0)|, t) \\ &+ \gamma\left(\sup_{0\leq\tau\leq t} \max_i |\dot{\theta}_i(\tau)|\right), \end{aligned} \qquad (4)$$

for all $t \in \text{dom}(u, \theta)$, $u_i(0) \in \mathbb{R}$, $\theta_i(0) \in \Theta_i$, and $i \in \mathcal{V}$. □

Remark 1: Note that inequality (4) essentially describes an input-to-state stability bound with respect to the "input" $\max_i |\dot{\theta}_i|$, see [18] and [24] for related notions. □

The pseudogradient of the time-varying game (2) is defined by the following vector in $\mathbb{R}^n$:

$$G(u, \theta) := [\nabla_{u_1}\phi_1(u, \theta_1), \nabla_{u_2}\phi_2(u, \theta_2), \ldots, \nabla_{u_n}\phi_n(u, \theta_n)]^\top.$$

where $\nabla_{u_i}\phi_i(u, \theta_i)$ stands for the partial derivative of ϕ_i with respect to u_i. In this paper, we will consider time-varying games that satisfy the following assumption.

Assumption 2: The functions ϕ_i, h, and G satisfy:

1) For each $i \in \mathcal{V}$, the functions $\phi_i(\cdot, \cdot)$ and $h(\cdot)$ are continuously differentiable.
2) There exists $\ell_u > 0$ such that
$$|G(u', \theta) - G(u'', \theta)| \leq \ell_u |u' - u''|, \qquad (5)$$
for all $u', u'' \in \mathbb{R}^n$ and all $\theta \in \Theta$.
3) There exists $\kappa > 0$ such that
$$\big(G(u', \theta) - G(u'', \theta)\big)^\top (u' - u'') \geq \kappa |u' - u''|^2, \qquad (6)$$
for all $u', u'' \in \mathbb{R}^n$ and all $\theta \in \Theta$. □

The properties described in Assumption 2 are fairly standard in the literature on fast (time-invariant) NE seeking [5], [9], with the difference that in (5) and (6) the parameters (ℓ_u, κ) are assumed to hold uniformly in θ.

In some cases, we will also consider the following assumption, which describes the so-called *potential games*.

Assumption 3: There exists a function $P : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ such that:

1) $\nabla_u P(u, \theta) = G(u, \theta)$ for all $u \in \mathbb{R}^n$ and all $\theta \in \Theta$.
2) There exists $\ell_\theta > 0$ such that

$$|\nabla_\theta P(u', \theta) - \nabla_\theta P(u'', \theta)| \leq \ell_\theta |u' - u''|, \qquad (7)$$

for all $u', u'' \in \mathbb{R}^n$ and all $\theta \in \Theta$.

The following lemma will be instrumental for our results.

Lemma 1: Suppose that Assumptions 1-2 hold, and let $\Pi = \Pi_1 \times \Pi_2 \times \ldots \times \Pi_n$ and $\Theta = \Theta_1 \times \Theta_2 \times \ldots \times \Theta_n$ be given by (3). Then, there exist constants $m_h, m_\Pi > 0$ such that $|\Pi(\theta)| \leq m_\Pi$ and $|\nabla h(\theta)| \leq m_h$ for all $\theta \in \Theta$. $\square$

IV. FIXED-TIME TRACKING BOUNDS: THE MODEL-BASED CASE

In this Section, we study tracking bounds for NETs given by (2) driven by the dynamics (3). We first assume that agents have access to direct measurements of their gradients $\nabla_{u_i} \phi_i$, and we defer to Section V the model-free case where only measurements of ϕ_i are available to the agents.

A. Potential Games

When a potential function exists, we consider the following *fixed-time pseudogradient* dynamics (FxTPG) implemented by the i^{th} player:

$$\dot{u}_i = -\frac{k_i}{2} \left(\frac{\nabla_{u_i} \phi_i(u, \theta_i)}{|\nabla_{u_i} \phi_i(u, \theta_i)|^\alpha} + \frac{\nabla_{u_i} \phi_i(u, \theta_i)}{|\nabla_{u_i} \phi_i(u, \theta_i)|^{-\alpha}} \right), \qquad (8)$$

where $k_i > 0$ is the individual learning gain of each agent, $\alpha \in [0, 1]$, and the right-hand side of (8) is defined to be zero whenever $\nabla_{u_i} \phi_i(u, \theta_i) = 0$. Note that the normalizing term in (8) is the individual partial derivative of the cost of each agent, which differentiates (8) from fixed-time gradient flows studied in the optimization literature [19], [25]. Thus, when players have access to Oracles that provide measurements or evaluations of their local gradients, equation (8) can be implemented in a decentralized way. In fact, when $\alpha = 0$ system (8) reduces to the standard pseudogradient flow $\dot{u} = -KG(u, \theta)$ [5], [26], where $K = \text{diag}([k_1, k_2, \ldots, k_n])$.

The following result establishes a tracking bound for system (8) with $\alpha = 0$. This bound will serve as a benchmark for the case $\alpha > 0$. All the proofs are omited due to space limitations.

Proposition 1: Suppose that Assumptions 1-2 hold, and $\alpha = 0$. Then, the dynamics (3) and (8) have the (β_0, γ_0)-*NET tracking property* with

$$\beta_0(r, s) := \sqrt{\frac{\bar{k}}{\underline{k}}} r e^{-\frac{\underline{k}\kappa(1-\lambda)}{2} s}, \quad \gamma_0(q) := \tilde{k} q, \qquad (9)$$

where $\lambda \in (0, 1)$, $\bar{k} = \max_i k_i$, $\underline{k} = \min_i k_i$, and $\tilde{k} = \frac{2m_h \bar{k}^2}{\lambda \kappa \underline{k}^3}$. $\square$

The exponential ISS bound of Proposition 1 is a tracking-like result, similar to existing results in the literature of optimization [10], [11], [27]. Note that in (9), the ISS gain $\gamma(\cdot)$ is linear, which implies that the time-variation of $\max_i \dot{\theta}_i$ is linearly mapped to the residual tracking error. Since $\beta \in \mathcal{KL}$, it follows that as $t \to \infty$ we obtain

$$\limsup_{t \to \infty} |u(t) - u^*(t)| \leq \frac{2m_h \bar{k}^2}{\lambda \underline{k}^3} \left(\frac{\sup_{t \geq 0} \max_i |\dot{\theta}_i(t)|}{\kappa} \right).$$

Thus, asymptotic tracking is achieved if $\dot{\theta}(t) \to 0$ as $t \to \infty$.

Next, we consider the case $\alpha = 1$ in (8). This choice leads to a discontinuous dynamical system with solutions that are still well-defined in a generalized sense (in the sense of Filippov or Krasovskii). In this case, we focus our attention on potential games.

Proposition 2: Suppose that Assumptions 1-3 hold, $\alpha = 1$, and the learning gains satisfy

$$\min_i \ k_i > \frac{\ell_\theta \sqrt{2n}}{\kappa} \max_i |\dot{\theta}_i|. \qquad (10)$$

Then, the dynamics (3) and (8) have the (β_1, γ_1)-*NET tracking property* with

$$\beta_1(r, s) := c_1 \tan\left(\max\left\{ 0, -c_2 s + \arctan\left(c_3 r^{\frac{2}{\alpha}} \right) \right\} \right)^{\frac{\alpha}{2}}, \qquad (11)$$

and $\gamma_1 := 0$, where $c_j > 0$, $j \in \{1, 2, 3\}$. $\square$

The result of Proposition 2 establishes exact tracking of the NET under the dynamics (8), provided the minimum gain used by the players is larger than a particular threshold proportional to the value of the maximum variation of θ_i among all players. This result generalizes the fixed-time convergence results presented in [22] for static NE seeking problems. However, note that the discontinuity of the vector field can induce chattering along the NET and also in the best response set of each player.

The next proposition considers the case when $\alpha \in (0, 1)$, which eliminates the discontinuity on the right-hand side of system (8).

Proposition 3: Suppose that Assumptions 1-3 hold and $\alpha \in (0, 1)$. Then, the dynamics (3) and (8) have the $(\beta_{01}, \gamma_{01})$-*NET tracking property* with β_{01} equal to β_1 in (11), and γ_{01} given by $\gamma_{01}(q) := \rho^{-1}(q)$, where $\rho^{-1}(\cdot) \in \mathcal{K}_\infty$ is the inverse of the function $\rho : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ given by:

$$\rho(s) := \epsilon_1 s^{1-\alpha} + \epsilon_2 s^{1+\alpha}, \qquad (12)$$

with $\epsilon_1, \epsilon_2 > 0$. $\square$

Remark 2: Note that $\rho^{-1}(q) \approx 0$ when $\alpha \to 1$ and $|q| < 1$. This provides a substantial attenuation of the residual tracking error of the NET. $\square$

B. Non-Potential Games

We now focus on games that do not have potential functions, i.e., Assumption 3 does not hold. In this case, we consider individual NE tracking dynamics normalized by the complete pseudogradient of the game:

$$\dot{u}_i = -\frac{1}{2}k_i\Psi(u,\theta,\alpha)\nabla_{u_i}\phi_i(u,\theta_i), \tag{13}$$

where the scalar-valued mapping Ψ is defined as

$$\Psi(u,\theta,\alpha) = \frac{1}{|G(u,\theta)|^{\alpha}} + |G(u,\theta)|^{\alpha}. \tag{14}$$

Similar dynamics for *time-invariant*, non-potential, strongly monotone games were studied in [22]. In particular, the next result generalizes [22, Prop. 3] for time-varying games, and parallels the tracking bound established in Proposition 2.

Proposition 4: Suppose that Assumptions 1-2 hold, $\alpha = 1$, and the learning gains satisfy

$$\min_i \ k_i > \frac{d_0\sqrt{n}}{2\kappa a_0}\max_i|\dot{\theta}_i|. \tag{15}$$

for a given $c_0, c_1 > 0$. Then, the dynamics (3) and (13) have the $(\tilde{\beta}_1,\tilde{\gamma}_1)$-*NET tracking property* with $\tilde{\beta}_1$ given by (11), and $\tilde{\gamma}_1 = 0$.

Remark 3: The tracking result of Proposition 4 holds for smooth strongly monotone games that are not necessarily potential games, which is why the result is not covered by existing stability results. While the implementation of the dynamics requires full information of the game via the pseudogradient G, this information can be estimated in a distributed way by each agent of the system using multi-time scale consensus-based techniques. This approach will be studied in the next section. □

We finish this section with a result that holds for the case $\alpha \in (0,1)$.

Proposition 5: Suppose that Assumptions 1-2 hold and $\alpha \in (0,1)$. Then, the dynamics (3) and (13) have the $(\tilde{\beta}_{01},\tilde{\gamma}_{01})$-NET tracking property with $\tilde{\beta}_{01}$ of the form (11), and $\tilde{\gamma}_{01}$ given by $\tilde{\gamma}_{01}(q) := \tilde{\rho}^{-1}(q)$, where $\tilde{\rho}^{-1}(\cdot) \in \mathcal{K}_{\infty}$ is the inverse of the function $\tilde{\rho} : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, which is of the form (12). □

V. PRACTICAL FIXED-TIME TRACKING BOUNDS: THE MODEL-FREE CASE

In this section, we leverage Propositions 1-5 to design model-free implementations of the tracking dynamics studied in Section IV. In particular, we now consider the scenario where each player has access only to measurements of their own cost function ϕ_i, and, in some cases, to the states of other neighboring players $\mathcal{N}_i$ characterized by a communication graph that is assumed to be undirected and connected. The model-free dynamics aim to *emulate* the behavior of the algorithms (8) and (13).

A. Potential Games

To emulate the behavior of system (8), each player implements the following multi-time scale model-free dynamics:

$$\dot{\hat{u}}_i = -k_i\left(\frac{\xi_i}{|\xi_i|^{\alpha}} + \frac{\xi_i}{|\xi_i|^{-\alpha}}\right), \quad \hat{u}_i \in \mathbb{R}, \tag{16a}$$

$$\dot{\xi}_i = \frac{1}{\varepsilon_f}\left(-\xi_i + \phi_i(u,\theta_i)M_i(\mu_i)\right), \quad \xi_i \in \mathbb{R}^n, \tag{16b}$$

$$\dot{\mu}_i = \frac{1}{\varepsilon_p}\mathcal{R}_{\kappa,i}\mu_i, \quad \mu_i \in \mathbb{S}^1, \tag{16c}$$

where the right-hand side of (16a) is defined to be zero when $\xi_i = 0$. In (16b), $M_i(\mu_i) := 2\varepsilon_a^{-1}\mu_{i,1}$, and in (16c) the matrix $\mathcal{R}_{\kappa,i}$ is given by $\mathcal{R}_i = 2\pi[0,\tilde{\kappa}_i;-\tilde{\kappa}_i,0] \in \mathbb{R}^{2\times 2}$, $i \in \{1,2,\ldots,n\}$, with $\tilde{\kappa}_i > 0$. The parameter α still satisfies $\alpha \in [0,1]$, but now the actions u_i of the players are updated as follows:

$$u_i = \hat{u}_i + \varepsilon_a\mu_{i,1}, \tag{17}$$

The algorithm (16)-(17) is based on ideas of extremum seeking control [28], [7]. In particular, equation (16c) describes a dynamic oscillator evolving on the unit circle $\mathbb{S}^1$, generating periodic dither signals $\mu_{i,1}(t) = \mu_{i,1}(0)\cos\left(\frac{2\pi t}{\varepsilon_p}\tilde{\kappa}_i\right) + \mu_{i,2}(0)\sin\left(\frac{2\pi t}{\varepsilon_p}\tilde{\kappa}_i\right)$, with initial conditions satisfying $\mu_{i,1}(0)^2 + \mu_{i,2}(0)^2 = 1$. We make the following assumptions on the parameters of (16c).

Assumption 4: For all $i \in \mathcal{V}$ and $j \neq i$ we have that $\tilde{\kappa}_i \neq \tilde{\kappa}_j$, $\tilde{\kappa}_i \neq 2\tilde{\kappa}_j$, $\tilde{\kappa}_i \neq 3\tilde{\kappa}_j$. Moreover, $\tilde{\kappa}_i > 0$ is a rational number. □

In addition to the learning gains k_i and the frequencies $\tilde{\kappa}_i$, the dynamics (16) have three main tunable parameters: $(\varepsilon_a,\varepsilon_f,\varepsilon_p)$. To simplify the notation we assume that these parameters are the same for all players, but it is straightforward to extend our results to dynamics with heterogenous parameters. The parameter ε_a corresponds to the amplitude of the sinusoidal signal added in (17) to $\hat{u}_i$. This dither signal allows players to perform a local exploration of their cost function in a neighborhood of their current action. The parameter ε_p characterizes the time scale of the frequencies of the dither signals. As $\varepsilon_p \to 0^+$, the faster oscillatory behavior induced by the dither signals will permit the application of averaging theory to analyze the model-free dynamics. The parameter ε_f characterizes the gain of the low pass filter. When ε_f is small, this filter will permit a transparent stability analysis by enabling a clean computation of the average dynamics of (16a)-(16b) along the solutions of (16c) by removing from the right-hand side of (16a) any dependence on μ. □

Remark 4: The rationale behind the dynamics (16) is the following: as $\varepsilon_p \to 0^+$ the dynamic oscillator generates a sinusoidal signal μ_i with high frequency. Since this signal is added to the argument of ϕ_i via (17), and since ϕ_i is multiplied again by μ_i via the mapping M_i in (16b), the resulting signal ϕ_iM_i generates an approximation of the derivative $\frac{\partial\phi_i}{\partial\hat{u}_i}$, on compact sets. In turn, as $\varepsilon_f \to 0^+$, the state of the low pass filter (16b) converges exponentially fast

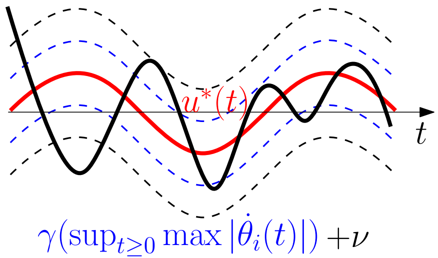

Fig. 2: Semi-global practical tracking of a NET $t \mapsto u^*(t)$.

to the derivative $\frac{\partial \phi_i}{\partial \hat{u}_i}$. It then follows that, as $(\varepsilon_f, \varepsilon_a) \to 0^+$, the dynamics (16a) behave similar to (8) on compact sets and compact time-domains, with a residual error induced by the fast stable dynamics (16b)-(16c). These arguments can be formalized using singular perturbation theory and averaging theory for perturbed nonsmooth systems to establish a "semi-global practical" fixed-time stability result for the main state $\hat{u}$. □

The following definition will be used to characterize the stability and convergence properties of the model-free dynamics.

Definition 2: The dynamics $(\dot{\hat{u}}, \dot{\theta}, \dot{\xi}, \dot{\mu})$ are said to have the *semi-global practical* (β, γ)*-NET tracking property* if $\exists\ \beta \in \mathcal{KL}$, $\gamma \in \mathcal{K}_\infty$ such that the following holds: $\forall$ $k_i, \varepsilon_o > 0$ and $\forall\ \Delta > \nu > 0$, $\exists\ \varepsilon_f^* > 0$, such that $\forall$ $\varepsilon_f \in (0, \varepsilon_f^*)$, $\exists\ \varepsilon_a^* > 0$, such that $\forall\ \varepsilon_a \in (0, \varepsilon_a^*)$, $\exists\ \varepsilon_p^* > 0$ such that $\forall\ \varepsilon_p \in (0, \varepsilon_p^*)$, all solutions $t \mapsto (\hat{u}, \theta, \xi, \mu)$ with initial conditions satisfying:

$$|\hat{u}(0) - u^*(0)| \leq \Delta,\ |\xi(0)| \leq \Delta,\ \mu(0) \in \mathbb{T}^n,\ \theta(0) \in \Theta,$$

also satisfy the following bound for all $t \geq 0$:

$$|\hat{u}(t) - u^*(t)| \leq \beta(|\hat{u}(0) - u^*(0)|, t) + \gamma(\varepsilon_o) + \nu, \quad (18)$$

and $\limsup_{t\to\infty} |\xi(t)| \in \mathcal{O}(\gamma(\varepsilon_o) + \nu + \varepsilon_a)$. □

The bound (18) describes a semi-global practical bound with two residual terms: 1) the term $\gamma(\varepsilon_0)$, which gauges the residual tracking error induced by the variation of $\dot{\theta}$, which is of order $\mathcal{O}(\varepsilon_0)$; 2) the term $\nu > 0$, which is the residual error induced by the multi-time scale model-free dynamics. The bound (18) is related to the notions of semi-global practical ultimate boundedness and semi-global practical input-to-state stability [29].

The following Theorem is the first main result of this paper.

Theorem 1: Suppose that Assumptions 1-4 hold. Then, the dynamics (3) and (16) have the *semi-global practical* (β, γ)*-NET tracking property*. Moreover, the following holds:

1) If $\alpha = 0$, then $(\beta,\ \gamma)$ are as given by (9).
2) If $\alpha = 1$ and (15) holds, then β is as given in (11) and γ is equal to zero.
3) If $\alpha \in (0, 1)$, then β is as given by (11), and γ is the inverse of (12).

B. Non Potential Games

For non-potential games, we seek to emulate the behavior of the NE tracking dynamics (13), in a distributed and model-free way. Thus, we consider the following algorithm:

$$\dot{\hat{u}}_i = -k_i \left(\frac{\xi_{ii}}{|\xi_i|^{\frac{\alpha}{2}}} + \frac{\xi_{ii}}{|\xi|^{\frac{-\alpha}{2}}} \right), \quad (19a)$$

$$\dot{\xi}_{ij} = \frac{1}{\varepsilon_f} \sum_{k \in \mathcal{N}_i} \Big(\xi_{kj} - \xi_{ij} \Big) + b_{ij} \left(\frac{2}{\varepsilon_a} \phi_i(u, \theta_i) M_i(\mu_i) - \xi_{ij} \right) \quad (19b)$$

$$\dot{\mu}_i = \frac{1}{\varepsilon_p} \mathcal{R}_{\kappa,i} \mu_i, \quad \mu_i \in \mathbb{S}^1, \quad (19c)$$

In these dynamics, each player is endowed with three types of auxiliary states $(\hat{u}_i, \xi_i, \mu_i)$, where $\xi_i := [\xi_{i1}, \xi_{i2}, \xi_{i3}, \ldots, \xi_{iN}]^\top \in \mathbb{R}^n$ is now an individual estimate of the pseudogradient G. In the consensus mechanism (19b) the constants b_{ij} satisfy $b_{ij} = 1$ if $i = j$, and $b_{ij} = 0$ for all $i \neq j$. Also, as in (16a), the right-hand side of (19a) is defined to be zero whenever $\xi = 0$. The individual action of the players is updated as in (17), and the parameter α satisfies $\alpha \in (0, 1]$.

Remark 5: The NE tracking dynamics (19) follow a similar rationale as the dynamics (16). Here, the dynamics (19b) allows players to estimate the overall pseudogradient G in a distributed way and also on a faster time scale compared to (19a), parameterized by ε_f. □

The following Theorem is the second main result of this paper.

Theorem 2: Suppose that Assumptions 1, 2 and 4 hold. Then, the dynamics (3) and (19) have the *semi-global practical* (β, γ)*-NET tracking property*. Moreover, the following holds:

1) If $\alpha = 1$ and (15) holds, then (β, γ) are as given in Proposition 4.
2) If $\alpha \in (0, 1)$, then (β, γ) are as given by Proposition 5.

VI. NUMERICAL EXAMPLE

To illustrate our theoretical results, we consider a *time-varying duopoly game*, similar to the one studied in [5, Sec. II] for time-invariant games. In a duopoly, two companies that produce the same good have dominant control over a market, and compete for profit by controlling their individual prices u_i. The payoffs of the companies are given by $J_i = s_i(u_i - m_i)$, where s_i is the number of sales of the i^{th} company, and m_i is the marginal cost. The sales s_i are modeled as $s_1 = S - s_2$ $s_2 = \frac{1}{p}(u_1 - u_2)$, where $p > 0$ is the preference of the consumer for company 1, and S is the total consumer demand. Given that in problem (2) every agent minimizes their cost, we define $\phi_i = -J_i$.

In contrast to [5, Sec. II], we consider time-varying duopoly games characterized by dynamic demands of the form $S(\theta(t)) = 100 + \theta(t)$, where $\theta(t) = 40\sin(t)$. This sinusoidal parameter can be easily generated by a linear exosystem that satisfies Assumption 1. Note that, in this

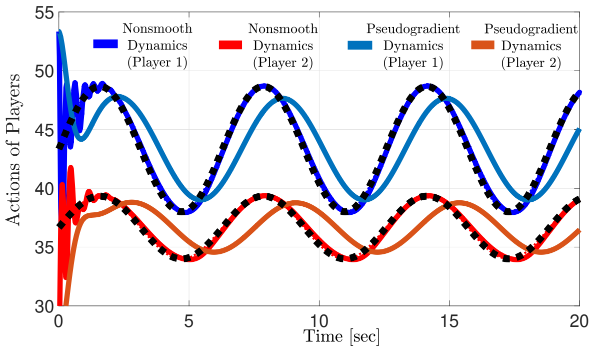

Fig. 3: Evolution in time of the actions of the players via the model-free NET tracking dynamics, with $\alpha = 0$ (light color) and with $\alpha = 0.95$ (dark color). The dashed black curves indicate the NET of the time-varying duopoly game.

case it suffices to consider one common parameter θ in both cost functions. The pseudogradient of the game is given by $G(u,\theta) = [-2u_1 + u_2 + (m_1 + S(\theta)p), -2u_2 + u_1 + m_2]^\top$, which satisfies Assumption 2 with $\ell_u = 3$ and $\kappa = 1$. This game also satisfies Assumption 3 since $S(\theta)$ appears linearly in the payoff function. The resulting NET is then given by

$$u_1^*(t) = \frac{1}{3}\Big(2m_1 + m_2 + 2pS(\theta(t))\Big)$$
$$u_2^*(t) = \frac{1}{3}\Big(m_1 + 2m_2 + pS(\theta(t))\Big).$$

To track this NET, we implement the decoupled model-free dynamics (16) with $\alpha = 0$ (corresponding to [5]) and $\alpha = 0.95$. The resulting trajectories are shown in Figure 3. The black dashed line indicates the NET $t \mapsto u^*(t)$. It can be observed that the nonsmooth dynamics ($\alpha = 0.95$) achieve much better tracking performance compared to the model-free pseudogradient flow studied in [5] and corresponding to $\alpha = 0$. In all simulations players used the same learning gains $k_1 = k_2 = 0.2$ and also the same frequencies $\tilde{\kappa}_i$ and parameters $(\varepsilon_f, \varepsilon_p, \varepsilon_a)$

VII. Conclusions

We studied the tracking problem of Nash equilibrium trajectories in time-varying non-cooperative games. We characterized different tracking bounds for smooth and nonsmooth algorithms in potential and non-potential games using semi-global practical input-to-state stability tools. In the nonsmooth case, we established semi-global practical fixed-time ISS, and showed that the ISS gain can dramatically attenuate the residual tracking error of the algorithms. Future research directions will focus on incorporating constraints into the action space of the agents, and designing model-free NET tracking dynamics using notions of homogeneity.

References

[1] J. S. Shamma, "Game theory, learning, and control systems," *National Science Review*, vol. 7, no. 7, pp. 1118–1119, 2020.

[2] J. Nash, "Non-cooperative games," *Annals of Mathematics*, vol. 54, no. 2, pp. 286–295, 1951.

[3] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. San Diego, CA: Academic Press, 1995.

[4] T. Başar, "Asynchronous algorithms in non-cooperative games," *Journal of Economic Dynamics and Control*, vol. 12, pp. 167–172, 1988.

[5] P. Frihauf, M. Krstic, and T. Başar, "Nash equilibrium seeking in noncooperative games," *IEEE Transactions on Automatic Control*, vol. 57, no. 5, pp. 1192–1207, May 2012.

[6] R. J. Kutadinata, W. H. Moase, and C. Manzie, "Dither re-use in Nash equilibrium seeking," *IEEE Trans. Autom. Contr.*, vol. 60, pp. 1433–1438, 2015.

[7] J. I. Poveda and A. R. Teel, "A framework for a class of hybrid extremum seeking controllers with dynamic inclusions," *Automatica*, vol. 76, pp. 113–126, 2017.

[8] M. Stanković, K. H. Johansson, and D. M. Stipanović, "Distributed seeking of Nash equilibria with applications to mobile sensor networks," *IEEE Trans. Autom. Contr.*, vol. 57, no. 4, pp. 904–919, 2012.

[9] C. D. Persis and S. Grammatico, "Distributed averaging integral Nash equilibrium seeking on networks," *Automatica*, vol. 110, pp. 1–7, 2019.

[10] A. Cherukuri, E. Mallada, S. Low, and J. Cortés, "The role of convexity in saddle-point dynamics: Lyapunov function and robustness," *IEEE Trans. Autom. Control.*, vol. 63, no. 8, pp. 2449–2464, 2018.

[11] G. Bianchin, J. Cortes, J. I. Poveda, and E. DallAnese, "Time-varying optimization of LTI systems via projected primal-dual gradient flows," *IEEE Trans. Control of Network Systems, appeared online.*, 2021.

[12] A. Scheinker and M. Krstic, "Extremum seeking-based tracking for unknown systems with unknown control directions," *51st IEEE Conference on Decision and Control*, pp. 6065–6070, 2012.

[13] V. Grushkovskaya, H-B., Durr, C. Ebenbauer, and A. Zuyev, "Extremum seeking for time-varying functions using lie bracket approximations," *IFAC-Papers Online*, vol. 50, pp. 5522–5528, 2017.

[14] B. Duvocelle, P. Mertikopoulos, M. Staudigl, and D. Vermeulen, "Multi-agent online learning in time-varying games," *arXiv:1809.03066*, 2021.

[15] B. Huang, C. Yang, Z. Meng, F. Chen, and W. Ren, "Distributed nonlinear placement for multicluster systems: A time-varying Nash equilibrium-seeking approach," *IEEE Transactions on Cybernetics, DOI: 0.1109/TCYB.2021.3085583*, pp. 1–10, 2021.

[16] M. Ye and G. Hu, "Distributed seeking of time-varying Nash equilibrium for non-cooperative games," *IEEE Trans. Autom. Contr.*, vol. 60, pp. 3000–3005, 2015.

[17] E. D. Sontag, "Smooth stabilization implies coprime factorization," *IEEE Transactions on Automatic Control*, vol. 34, pp. 435–443, 1989.

[18] D. Nešíc and A. R. Teel, "Input-to-state stability for nonlinear time-varying systems via averaging," *Mathematics of Controls, Signals and Systems*, vol. 14, pp. 257–280, 2001.

[19] J. I. Poveda and M. Krstić, "Non-smooth extremum seeking control with user-prescribed fixed-time convergence," *IEEE Transactions on Automatic Control*, vol. 66, no. 12, pp. 6156 – 6163, 2021.

[20] ——, "Fixed-time seeking and tracking of time-varying extrema," *In Proc. 60th IEEE Conference on Decision and Control, Austin, TX, USA.*, pp. 108–113, 2021.

[21] J. I. Poveda, M. Krstić, and T. Başar, "Fixed-time Nash equilibrium seeking in non-cooperative games," *In Proc. of 59th IEEE Conference on Decision and Control*, pp. 3514–3519, 2020.

[22] ——, "Fixed-time Nash equilibrium seeking in time-varying networks," *Submitted*.

[23] R. Goebel, R. Sanfelice, and A. R. Teel, *Hybrid Dynamical System*. Princeton, NJ: Princeton University Press, 2012.

[24] D. Angeli, E. Sontag, and Y. Wang, "Input-to-state stability with respect to inputs and their derivatives," *International Journal of Robust and Nonlinear Control*, vol. 13, pp. 1035–1056, 2003.

[25] K. Garg and D. Panagou, "Fixed-time stable gradient-flow schemes: Applications to continuous-time optimization," *IEEE Transactions on Automatic and Control*, vol. 66, no. 5, pp. 2002–2015, 2021.

[26] J. B. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, pp. 520–534, 1965.

[27] G. Bianchin, J. I. Poveda, and E. DallAnese, "Online optimization of switched LTI systems using continuous-time and hybrid accelerated gradient flows," *arXiv preprint arXiv:2008.03903*, 2020.

[28] K. Ariyur and M. Krstic, *Real-Time Optimization by Extremum-Seeking Control*. Hoboken, NJ: Wiley, 2003.

[29] W. Wang, D. Nesic, and A. R. Teel, "Input-to-state stability for a class of hybrid dynamical systems via averaging," *Mathematics of Control, Signals, and Systems*, vol. 23, pp. 223–256, 2012.