Excess Payoff Evolutionary Dynamics With Strategy-Dependent Revision Rates: Convergence to Nash Equilibria for Potential Games

Semih Kara, *Student Member, IEEE*, and Nuno C. Martins, *Senior Member, IEEE*

***Abstract*—Evolutionary dynamics in the context of population games models the dynamic non-cooperative strategic interactions among many nondescript agents. Each agent follows one strategy at a time from a finite set. A game assigns a payoff to each strategy as a function of the so-called population state vector, whose entries are the proportions of the population adopting the available strategies. Each agent repeatedly revises its strategy according to a revision protocol. We focus on a well-known class of protocols that prioritizes strategies with higher excess payoffs relative to a population-weighted average. In contrast to existing work for these protocols, we allow each agent's revision rate to depend explicitly on its current strategy. Motivated by applications and relevance to distributed optimization, we focus on potential games and investigate the population state's convergence to the game's Nash equilibria. Our contributions are twofold: (1) For the considered protocol class, prior work established conditions that ensure convergence under strategy-independent revision rates. We show that these conditions may be violated when the revision rates are strategy-dependent. (2) We prove that a minor, well-motivated modification of the considered protocol class satisfies these conditions for any strategy-dependent revision rates. We also illustrate our results using a distributed task allocation example.**

***Index Terms*—Game theory, stability of nonlinear systems.**

I. Introduction

THE POPULATION games and evolutionary dynamics framework [1], [2] models the dynamic strategic interactions among many agents. At any time, each agent follows a single strategy from a finite strategy set $\{1, \ldots, n\}$. The agents are nondescript and grouped into a population. Consequently, the proportions of the population playing the available strategies suffice to specify the population's strategy profile. The n-dimensional vector composed of these proportions is called the population state. At any time, a payoff mechanism assigns a payoff to each strategy as a function of the population state. Each agent repeatedly revises its strategy in response to the payoffs and the population state, which causes the population state to vary over time. When performing revisions, the agents decide on their subsequent strategies via a stochastic heuristic, characterized by a so-called revision protocol.

So-called revision rates [3], given by positive real numbers, specify how frequently the agents revise their strategies. In this letter, we investigate whether allowing the agents' revision rates to depend on their strategies affects the long-term behavior of the population state. This dependence means that, given a vector of revision rates $\lambda := [\lambda_1 \ \ldots \ \lambda_n]^T$, an agent following strategy i revises it with the rate λ_i, which may differ from λ_j for $i \neq j$. Allowing the revision rates to be strategy-dependent was first proposed and motivated in [3]. This generalization enables modeling a large class of systems that the original framework (in which $\lambda_i = \lambda_j$ for all i, j) cannot. In [3, Example 1] and [3, Remark 1], the authors introduce the hassle vs. price game and a labour market example as instances of such systems. We present another such instance in Section II-C, cast as a distributed task allocation problem. In these systems, strategy-dependent revision rates represent various phenomena, such as component lifetimes, contract turnover rates, and task service times that vary with, respectively, the manufacturer, company, and type of task.

Despite sharing the same framework and motivation, there are two main differences between this letter and [3]: **(1)** We consider excess payoff protocols [4], whereas the focus in [3] is pairwise comparison protocols [5]. Protocols in the excess payoff class prioritize strategies with higher excess payoffs computed as the differences between the strategies' payoffs and a population-weighted (payoff) average. **(2)** Instead of focusing on system-theoretic passivity properties as in [3], here we seek conditions guaranteeing desirable convergence properties for the population state when the payoff mechanism is a potential game [6]. Classical results were general enough

Manuscript received 16 September 2022; revised 11 November 2022; accepted 1 December 2022. Date of publication 16 December 2022; date of current version 29 December 2022. This work was supported by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-19-1-0315, and in part by NSF under Grant CNS-2135561. Recommended by Senior Editor V. Ugrinovskii. *(Corresponding author: Nuno C. Martins.)*

The authors are with the ECE Department and the ISR, University of Maryland, College Park, MD 20742 USA (e-mail: skara@umd.edu; nmartins@umd.edu).

Digital Object Identifier 10.1109/LCSYS.2022.3229962

to handle this case for pairwise comparison protocols, hence potential games were not a focus in [3].

A. Relevance of Studying Population State Convergence

To ascertain the infinite-horizon properties of the population state, we follow a common approach [1], [2], [3], [7], [8], [9], [10] based on deterministic mean-field approximation techniques valid for large populations. Specifically, we focus on a dynamical system, called the mean closed loop, consisting of the feedback interconnection of the so-called evolutionary dynamics model (EDM) and the payoff mechanism. The EDM is a nonlinear system specified by the revision protocol and λ, and its output is referred to as the mean population state. Because of the large number of agents, the mean population state is a good approximation of the population state over any finite time interval [11]. Moreover, the convergence of the mean population state to a set ensures that the population state converges with high probability to the vicinity of this set (see [3, Sec. II.C.4] and references therein).

B. Potential Games and Excess Payoff Protocols

Potential games were introduced in [12] and progressed to be analyzed extensively. They were adapted to the population games and evolutionary dynamics framework in [6], [13] and were further investigated in this context in [10], [14]. Congestion games [1], [15] are celebrated examples of potential games. Notably, potential games have a valuable connection to distributed optimization [14]: a potential game's Nash equilibria coincide with the points that satisfy the Karush-Kuhn-Tucker (KKT) conditions for the problem of maximizing the game's potential over the standard n-simplex.

Due to their suitability for capturing reluctance and moderation in finding best-performing strategies, excess payoff protocols were introduced in [4] as a well-motivated model of individual choice. In subsequent works [2], [8], these protocols are called excess payoff target (EPT) protocols, which is the nomenclature that we will use. Applicability of EPT protocols in engineering settings has been demonstrated in [16] for water distribution systems and in [17] for distributed wireless networks. We present another such application in Section VI, based on the distributed task allocation problem in Section II-C.

C. Nash Stationarity (NS) and Positive Correlation (PC)

When the payoff mechanism is a potential game, the work in [6] shows that if the EDM satisfies the so-called Nash stationarity (NS) and positive correlation (PC) conditions (see Section III for definitions), then the mean population state has desirable convergence properties. For instance, in [6], the author proves that if the payoff mechanism is a potential game and the EDM satisfies (NS) and (PC), then the mean population state converges to a Nash equilibrium of the game from any initial state. Additionally, the author provides criteria on the potential game that guarantee this convergence is to an appropriately-defined social optimum.

In this letter, to ascertain the convergence properties of the mean population state resulting from potential games and EPT protocols, we investigate the (NS) and (PC) properties of the associated EDM so as to use the results in [6]. Assuming that the revision rates are identical, the work in [4] shows that any EPT EDM satisfies (NS) and (PC). However, as we explain in Section I-C and Section IV, existing results [4] on the convergence properties of the mean population state under EPT protocols and potential games are not conclusive when there are strategies i and j for which $\lambda_i \neq \lambda_j$.

Our contributions are twofold: **(1)** We unveil the existence of EPT protocols and λ for which the EDM violates (NS) or (PC). Using these results, we show that for some EPT protocols and λ, the mean population state does not converge to a Nash equilibrium under a broad class of potential games. **(2)** We propose a modification of the EPT class, which we call sign preserving rate-modified EPT protocols. For any protocol in the modified class and any λ, we prove that the EDM satisfies (NS) and (PC). As we discuss in Section V and illustrate in Section VI, the protocols in the modified class are easy to implement.

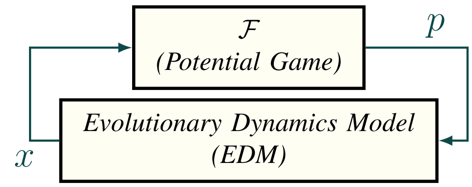

Fig. 1. Interconnection of the EDM and a potential game.

D. Auxiliary Notation

We denote the standard n-simplex by Δ and the n-dimensional column vector of ones by $\mathbb{1}$. We use $\mathbb{R}^n_{\geq 0}$, $\mathbb{R}^n_{>0}$ and $\mathbb{R}^n_{\leq 0}$ to symbolize respectively the n-dimensional non-negative, positive and non-positive orthant. Given $\mathbb{S} \subset \mathbb{R}^n$, we use $int(\mathbb{S})$ and $\delta(\mathbb{S})$ to denote respectively the interior and boundary of $\mathbb{S}$, and define $\mathbb{R}^n_* := \mathbb{R}^n \setminus int(\mathbb{R}^n_{\leq 0})$.

II. Framework Description

The mean closed loop, depicted in Fig. 1, is a deterministic dynamical system composed of the positive feedback interconnection of two sub-systems: the evolutionary dynamics model (EDM) and the payoff mechanism (as we will explain shortly, we focus on payoff mechanisms given by potential games). The input of the EDM, called the deterministic payoff, is denoted by p, and its output, called the mean population state, is denoted by x. In this section, we describe the payoff mechanism and the EDM in more precise terms.

Before we proceed, we revisit Section I-A, which summarizes the relation of the mean closed loop with the finite-agent model described at the beginning of Section I. Recall that, due to the large population assumption, both the finite and infinite horizon behavior of the population state can be approximated via x. We refer to [3] and the references therein for details on the finite-agent model and these approximations.

A. Potential Games

We consider that a Lipschitz continuous map $\mathcal{F} : \Delta \to \mathbb{R}^n$, called a population game (or game in short), determines the deterministic payoff at any time $t \geq 0$ as $p(t) = \mathcal{F}(x(t))$.

Moreover, we focus on $\mathcal{F}$ that is a potential game according to the definition[1] in [6], which we reproduce below.

Definition 1: A game $\mathcal{F}$ is said to be a potential game if there is a differentiable function $f : \mathbb{R}^n \to \mathbb{R}$, called the game's potential, that satisfies $\nabla f(\xi) = \mathcal{F}(\xi)$ for all $\xi \in \Delta$.

The Nash equilibria set is defined in the context of population games as follows.

Definition 2: The Nash equilibria of a game $\mathcal{F}$ is the set

$$\mathbb{NE}(\mathcal{F}) := \big\{\xi \in \Delta \mid \xi_i > 0 \Rightarrow \mathcal{F}_i(\xi) \geq \mathcal{F}_j(\xi) \text{ for all } j\big\}.$$

We note that the Nash equilibria set of any population game is non-empty [1, Th. 2.1.1] and can be interpreted in the "mass-action" sense as explained in [18].

From the definition of the EDM (which we present in the following section), it follows that $x(t) \in \Delta$ for all $t \geq 0$. Since $\mathcal{F}$ is continuous, this implies that p takes values in the bounded set $\mathfrak{P} := \{\mathcal{F}(\xi) \in \mathbb{R}^n \mid \xi \in \Delta\}$.

B. The Evolutionary Dynamics Model

The main object of our analysis in this letter is the EDM, which is specified by λ and the revision protocol. Before proceeding with the description of the EDM, we give further information about the protocol.

The population's revision protocol is a Lipschitz continuous function $\tau : \Delta \times \mathfrak{P} \to \mathbb{R}^{n\times n}_{\geq 0}$ that satisfies for all $i \in \{1, \dots, n\}$ and $(\xi, \pi) \in \Delta \times \mathfrak{P}$ the normalization equality $\sum_{j=1}^n \tau_{ij}(\xi, \pi) = 1$. In the finite-agent model, τ quantifies the agents' strategic behavior by specifying the probabilities with which they switch strategies. We refer to [3] for details.

A protocol class that is of particular interest to this letter is the excess payoff target (EPT) protocols, introduced in [4].

Definition 3: A protocol τ is said to be an excess payoff target (EPT) protocol if we can use $\hat{\pi} := \pi - \mathbb{1}\pi^T\xi$ to write it for all $(\xi, \pi) \in \Delta \times \mathfrak{P}$ and $i, j \in \{1, \dots, n\}$ with $i \neq j$ as:

$$\tau_{ij}(\xi, \pi) = \frac{\varphi_j(\hat{\pi})}{\bar{\tau}}, \;\; \tau_{ii}(\xi, \pi) = 1 - \sum_{\ell=1, \ell\neq i}^{n} \tau_{i\ell}(\xi, \pi). \quad (1)$$

Here, $\varphi : \mathbb{R}^n \to \mathbb{R}^n_{\geq 0}$ is a function that satisfies acuteness, which requires that $\varphi(\eta)^T\eta > 0$ for all $\eta \in int(\mathbb{R}^n_*)$ [4]. Moreover, $\bar{\tau}$ is a constant such that τ_{ii} is non-negative for all $i \in \{1, \dots, n\}$, and $\hat{\pi}$ is called the excess payoff.

Intuitively, acuteness ensures that an agent following an EPT protocol is more likely to switch to strategies with payoffs that are higher than the population-average payoff. An example of an EPT protocol is the Brown-von Neumann-Nash (BNN) protocol, which is obtained by setting $\varphi_j(\hat{\pi}) = \max\{0, \hat{\pi}_j\}$. The BNN protocol was introduced in [19] to analyze symmetric zero-sum games and later used in [20] to show the existence of Nash equilibria in normal-form games.

Having described the revision protocols in more detail, we now characterize the EDM. The EDM is the system with input p, output x, and the state equation

$$\dot{x}(t) = \mathcal{V}\big(x(t), p(t)\big), \quad t \geq 0,$$

where $x(0) \in \Delta$, and $\mathcal{V} : \Delta \times \mathbb{R}^n \to \mathbb{R}^n$ is given for all $(\xi, \pi) \in \Delta \times \mathfrak{P}$ and $i \in \{1, \dots, n\}$ by

$$\mathcal{V}_i(\xi, \pi) := \underbrace{\sum_{j=1}^{n} \lambda_j \tau_{ji}\big(\xi, \pi\big)\xi_j}_{\text{inflow to strategy } i} - \underbrace{\sum_{j=1}^{n} \lambda_i \tau_{ij}\big(\xi, \pi\big)\xi_i}_{\text{outflow from strategy } i}. \quad (2)$$

Observe that $x(0) \in \Delta$ assures that $x(t) \in \Delta$ for all $t \geq 0$.

C. Motivating Example

The following distributed task allocation problem illustrates our framework and motivates allowing the revision rates to be strategy-dependent. Distributed task allocation problems in the population games and evolutionary dynamics framework have been studied in [21]. However, the setting in [21] is different from ours because [21] assumes identical revision rates.

Suppose that there is a large number of agents and 3 types of tasks that the agents perform, which constitute their strategy set. Tasks of type $i \in \{1, 2, 3\}$ have expected service time $\lambda_i^{-1} > 0$, which may differ from $\lambda_j^{-1} > 0$ for $i \neq j$. Each agent undertakes a new task, possibly of a different type, only after completing its current one. Hence, the revision rate associated with tasks of type i is $\lambda_i > 0$ (see [3] for the physical meanings of the revision rates). At any time $t \geq 0$, a central entity assigns a payoff vector $P(t)$ to the types of tasks and disseminates $P(t)$. When an agent completes a task, it decides on the type of its new task in a distributed manner using a protocol τ and announces the types of its completed and new tasks to the central entity. Let us denote the fraction of agents working on tasks of type i at time t by $X_i(t)$. Given $\theta \in \Delta$, the distributed task allocation problem is to find a τ and a payoff mechanism that drives X near θ in the long run. We will revisit this problem in Section VI to demonstrate our results.

III. Problem Formulation

In the upcoming sections, we first study the infinite-horizon properties of x resulting from potential games and EPT protocols. This analysis shows that desirable convergence properties are not ensured by the EPT class when the revision rates are strategy-dependent, thus we propose a modification of this class that does.

As summarized in Section I-C, provided that $\mathcal{F}$ is a potential game, the work in [6] develops a methodology to ascertain the long-run properties of x. Namely, in [6], the author shows that if $\mathcal{F}$ is a potential game and the EDM satisfies the so-called Nash stationarity (NS) and positive correlation (PC) conditions (which we will define shortly), then x has several desirable convergence guarantees. One of these guarantees is that x converges to $\mathbb{NE}(\mathcal{F})$ from any initial state. Denoting the potential of $\mathcal{F}$ by f, the set of points satisfying the KKT conditions for the problem of maximizing $f(\xi)$ over $\xi \in \Delta$ coincides with $\mathbb{NE}(\mathcal{F})$ [6]. Hence, if f is concave, then this convergence is also to the set $\arg\max_{\xi\in\Delta} f(\xi)$. Another guarantee is that, under an additional so-called homogeneity constraint [6] on $\mathcal{F}$, the mean population state x converges to an appropriately defined social optimum from any initial state. For the remaining long-run properties derived in [6] we direct the reader to [6] and proceed by presenting the definitions of (NS) and (PC).

[1] In several articles, games satisfying Definition 1 are referred to as full potential games [1], [14]. However, we conform to the terminology in [6] because we frequently invoke the results therein.

Definition 4: The EDM is said to be Nash stationary (NS) if for every $(\xi, \pi) \in \Delta \times \mathfrak{P}$ the following holds:

$$\mathcal{V}(\xi, \pi) = 0 \iff \xi \in \arg\max_{\eta \in \Delta} \pi^T \eta. \tag{NS}$$

Definition 5: The EDM is said to satisfy positive correlation (PC) if for every $(\xi, \pi) \in \Delta \times \mathfrak{P}$ the following holds:

$$\mathcal{V}(\xi, \pi) \neq 0 \Rightarrow \mathcal{V}(\xi, \pi)^T \pi > 0. \tag{PC}$$

Notice that (NS) establishes an important connection between $\mathcal{F}$ and the mean closed loop. Namely, the mean closed loop is at equilibrium if and only if x belongs to $\mathbb{NE}(\mathcal{F})$. As for (PC), it means that $\dot{x}$ and $\mathcal{F}(x)$ make an acute angle as long as x is not stationary.

To perform the analysis stated at the beginning of this section, we employ the methodology in [6] and focus on the following two problems. **(Problem 1)** Does the EDM satisfy (NS) and (PC) for every EPT protocol and $\lambda \in \mathbb{R}^n_{>0}$? When the revision rates are identical, the EDM is known to satisfy (NS) and (PC) for every EPT protocol [4]. However, we show in Section IV that the answer to Problem 1 is negative because we allow $\lambda_i \neq \lambda_j$ for $i \neq j$. Thus, we also pose the ensuing problem. **(Problem 2)** Can we modify the EPT class so that the EDM satisfies (NS) and (PC) for every τ in the modified class and $\lambda \in \mathbb{R}^n_{>0}$? Sign preserving rate-modified EPT (RM-EPT) protocols, which we introduce at the end of Section IV and analyze in Section V, will be our answer to this problem.

IV. EPT Protocols Under Strategy-Dependent Revision Rates and RM-EPT Protocols

In this section, as opposed to the existing results under identical revision rates, we show that the EPT class does not ensure (NS) and (PC) when the revision rates are allowed to be strategy-dependent. Using our results on (NS) we also show that, for some EPT protocols and potential games, the convergence of x to $\mathbb{NE}(\mathcal{F})$ is not robust under perturbations of the revision rates. Subsequently, we modify the EPT class so that the EDM satisfies (NS) and (PC) for all $\lambda \in \mathbb{R}^n_{>0}$.

We begin by showing that even one of the most ubiquitous EPT protocols, the BNN protocol, produces an EDM that violates (PC) for some λ. Recall from Section II-B that the BNN protocol is obtained by setting $\varphi_j(\hat{\pi}) = \max\{0, \hat{\pi}_j\}$ in (1).

Counterexample 1: Consider τ to be the BNN protocol with $\bar{\tau} = 1$ and let $\xi = [1/2 \quad 0 \quad 1/2]^T$, $\pi = [0 \quad 1/2 \quad 3/4]^T$, $\lambda = [1 \quad 1 \quad 16]^T$. Substituting these into (2), we have $\mathcal{V}(\xi, \pi) \neq 0$ and $\mathcal{V}(\xi, \pi)^T \pi = -0.0156 < 0$. Therefore, the EDM specified by τ and λ violates (PC).

Having established that EPT protocols may fail to guarantee (PC) for all $\lambda \in \mathbb{R}^n_{>0}$, we modify the EPT class to eliminate this issue. We construct the modified class, which we name rate-modified EPT (RM-EPT) protocols, by replacing $\hat{\pi}$ in Definition 3 with $\tilde{\pi}$ defined as

$$\tilde{\pi} := \pi - \mathbb{1}\bar{\pi}, \quad \bar{\pi} := \frac{\sum_{i=1}^n \lambda_i \xi_i \pi_i}{\sum_{j=1}^n \lambda_j \xi_j}, \tag{3}$$

where we refer to $\bar{\pi}$ as the rate-weighted average payoff. Thus, given an EPT protocol as in Definition 3, we obtain its rate-modified version by simply replacing $\hat{\pi}$ in (1) with $\tilde{\pi}$. Notice that an EPT protocol coincides with its rate-modified counterpart when the revision rates are identical.

Remark 1: Suppose that, at every revision time, the revising agent declares (possibly to a coordinator) the payoff of the strategy it was following immediately prior to the revision time. For instance, in our distributed task allocation problem in Section II-C, this would correspond to the agents announcing the payoffs they received from the tasks that they recently completed. Since the number of agents is large, the average of the payoffs announced within a short time interval is approximately the rate-weighted average payoff. Consequently, in such a setting, RM-EPT protocols can be implemented effortlessly, without any knowledge of λ.

Although this modification suffices to ensure (PC), it does not guarantee (NS). Consider the rate-modified version of the EPT protocol in [4, Proposition 2.1], characterized by

$$\varphi_j(\tilde{\pi}) = \sum_{i=1}^n \left(e^{c\tilde{\pi}_i}(k+1)([\tilde{\pi}_j]_+)^k + e^{c\tilde{\pi}_j} c([\tilde{\pi}_i]_+)^{k+1}\right), \tag{4}$$

where $c > 0$, $k > 0$, $(k+1)e^{k+2} + 1 \geq n$ and $[\zeta]_+$ denotes $\max\{0, \zeta\}$ for every $\zeta \in \mathbb{R}$. A key property of this protocol is that it satisfies the positivity condition (Pos) given below:

$$\xi \notin \arg\max_{\eta \in \Delta} \pi^T \eta \Rightarrow \min_{ij} \tau_{ij}(\xi, \pi) > 0, \quad \pi \in \mathbb{R}^n. \tag{Pos}$$

The following counterexample shows that the EDM violates (NS) for the protocol specified by (4) and some λ.

Counterexample 2: Let τ be the RM-EPT protocol characterized by (4) and assume that $\mathcal{F}$ satisfies $int(\Delta) \not\subset \mathbb{NE}(\mathcal{F})$ (for instance, $\mathcal{F}$ given by $\mathcal{F}(\xi) = \xi - (1/n)\mathbb{1}$ satisfies this criterion because $\mathbb{NE}(\mathcal{F}) = \{(1/n)\mathbb{1}\}$). Let us take $\xi \in int(\Delta)$ such that $\xi \notin \mathbb{NE}(\mathcal{F})$ and set $\pi = \mathcal{F}(\xi)$. Note that $\mathcal{V}(\xi, \pi) = 0$ if and only if

$$\tau(\xi, \pi)^T \Xi \lambda = \Xi \lambda, \tag{5}$$

where Ξ denotes the n-by-n diagonal matrix with its (i, i)-th entry given by ξ_i for all $i \in \{1, \ldots, n\}$. From $\xi \notin \mathbb{NE}(\mathcal{F})$ and (Pos), it follows that $\tau_{ij}(\xi, \pi) > 0$ for all $i, j \in \{1, \ldots, n\}$. Therefore, by $\tau(\xi, \pi)$ being a stochastic matrix and the Perron-Frobenius Theorem, $\tau(\xi, \pi)^T$ has an eigenvector ν that corresponds to the eigenvalue 1 and all entries of ν are positive. Because Ξ is invertible, $\Xi^{-1}\nu$ is the unique λ for which $\Xi\lambda = \nu$. In addition, from Ξ^{-1} being a diagonal matrix with positive diagonal entries and ν having positive entries, it follows that each entry of $\Xi^{-1}\nu$ is positive. Thus, choosing $\lambda = \Xi^{-1}\nu$ does not conflict with $\lambda \in \mathbb{R}^n_{>0}$, and with this choice (5) holds. Since $\xi \notin \mathbb{NE}(\mathcal{F})$ implies $\xi \notin \arg\max_{\eta \in \Delta} \pi^T \eta$, we conclude that the EDM generated by τ and λ violates (NS).

Remark 2: Our derivations in Counterexample 2 are valid not only for the protocol specified by (4), but for every τ that satisfies (Pos).

Remark 3: For any τ satisfying (Pos), there are λ, ξ, and a potential game $\mathcal{F}$ (see Counterexample 2) such that $\mathcal{V}(\xi, \mathcal{F}(\xi)) = 0$ and $\xi \notin \mathbb{NE}(\mathcal{F})$. Hence, the resulting trajectory of x with initial state ξ remains at ξ and does not converge to $\mathbb{NE}(\mathcal{F})$. Importantly, the EPT protocol in [4, Proposition 2.1] (not rate-modified version) satisfies (Pos) [4]. Consequently, when the revision rates are strategy dependent, convergence of x to $\mathbb{NE}(\mathcal{F})$ is not guaranteed by the EPT class, even when $\mathcal{F}$ is potential.

In view of Remark 2, to arrive at a protocol class that assures both (PC) and (NS) under every $\lambda \in \mathbb{R}^n_{>0}$, we impose a "sign preservation" requirement on RM-EPT protocols. The resulting subclass of RM-EPT protocols, which we call sign preserving RM-EPT protocols, is explicitly defined below.

Definition 6: A protocol τ is a sign preserving RM-EPT protocol if it can be written for all $(\xi, \pi) \in \Delta \times \mathfrak{P}$ and $i, j \in \{1, \dots, n\}$ with $i \neq j$ as:

$$\tau_{ij}(\xi, \pi) = \frac{\varphi_j(\tilde{\pi})}{\bar{\tau}}, \ \ \tau_{ii}(\xi, \pi) = 1 - \sum_{\ell=1, \ell \neq i}^{n} \tau_{i\ell}(\xi, \pi). \quad (6)$$

Here, $\varphi_j : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is a Lipschitz continuous map that is sign preserving in the j-th component of its argument (meaning $\varphi_j(\tilde{\pi}) > 0$ if and only if $\tilde{\pi}_j > 0$), $\tilde{\pi}$ is the excess payoff relative to the rate-weighted average (see (3)) and $\bar{\tau}$ is a constant such that τ_{ii} is non-negative for all $i \in \{1, \dots, n\}$.

Remark 4: Observe from Definitions 3 and 6 that sign preservation implies acuteness. Therefore, sign preserving RM-EPT protocols form a subclass of RM-EPT protocols.

V. (NS) and (PC) Properties of Sign Preserving RM-EPT Protocols

We proceed by showing that the EDM satisfies (NS) and (PC) for every sign preserving RM-EPT protocol and $\lambda \in \mathbb{R}^n_{>0}$. Later, leveraging this result and the discussion in Section III, we establish a guarantee for the convergence of x to $\mathbb{NE}(\mathcal{F})$.

Theorem 1: If τ belongs to the sign preserving RM-EPT class, then the EDM satisfies (NS) for every $\lambda \in \mathbb{R}^n_{>0}$.

We prove Theorem 1 in the Appendix.

Theorem 2: If τ belongs to the sign preserving RM-EPT class, then the EDM satisfies (PC) for every $\lambda \in \mathbb{R}^n_{>0}$.

We present a proof of Theorem 2 in the Appendix.

As indicated in Section III, when $\mathcal{F}$ is a potential game, the (NS) and (PC) properties associated with sign preserving RM-EPT protocols lead to the convergence result below.

Corollary 1: If τ is a sign preserving RM-EPT protocol and $\mathcal{F}$ is a potential game, then for every $x_0 \in \Delta$ and $\lambda \in \mathbb{R}^n_{>0}$ the mean population state with initial value x_0 satisfies

$$\lim_{t \to \infty} \inf_{\xi \in \mathbb{NE}(\mathcal{F})} \|x(t) - \xi\| = 0.$$

Corollary 1 is an immediate consequence of [6, Th. 4.5], Theorem 1 and Theorem 2.

Remark 5: Pursuant to the discussion in Section III, there are results in [6] that lead to refinements of Corollary 1. For instance, under the conditions in Corollary 1 and an additional so-called homogeneity constraint on the game, it follows from [6, Th. 5.5] that x converges to a social optimum.

Remark 6: Notice that in order for Corollary 1 to hold, the potential of $\mathcal{F}$ does not need to be concave. Notably, this result holds for any potential game.

VI. Numerical Example

Consider the distributed task allocation problem in Section II-C. Suppose that the agents decide on the types of their new tasks in a distributed manner using the rate-modified BNN (RM-BNN) protocol, obtained by setting $\varphi_j(\tilde{P}(t)) = \max\{0, \tilde{P}_j(t)\}$ in (6). Hence, in addition to $P(t)$, let the central entity also calculate and announce $\tilde{P}(t) = P(t) - \mathbb{1}(\sum_{i=1}^{3} \lambda_i X_i(t) P_i(t) / \sum_{j=1}^{3} \lambda_j X_j(t))$. As stated in Remark 1, the central entity can approximate $\tilde{P}(t)$ accurately and effortlessly by averaging the payoffs of the types of tasks completed within $[t - \epsilon, t]$ for a small $\epsilon > 0$.

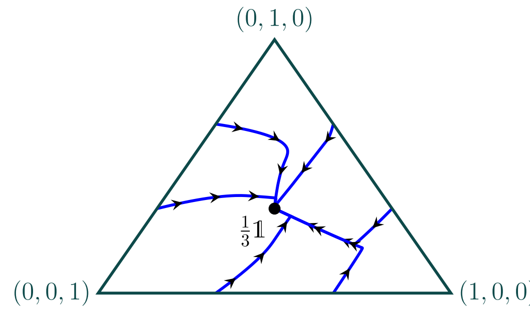

Fig. 2. Trajectories of x for the numerical example.

Now, given any desired allocation $\theta \in \Delta$ of proportions of agents on the types of tasks, let the central entity set $P(t) = \mathcal{F}(X(t))$, where $\mathcal{F}(\xi) = \theta - \xi$ for all $\xi \in \Delta$. Notice that $\mathcal{F}$ is a potential game with $\mathbb{NE}(\mathcal{F}) = \{\theta\}$. Furthermore, the RM-BNN protocol is a sign preserving RM-EPT protocol. Therefore, from Corollary 1, it follows that $\{\theta\}$ is globally attractive under the mean closed loop resulting from $\mathcal{F}$, the RM-BNN protocol and any $\lambda \in \mathbb{R}^n_{>0}$. This is displayed in Fig. 2 for $\theta = (1/3)\mathbb{1}$ and $\lambda = [1 \quad 10 \quad 100]^T$. Together with the mean-field approximation discussion in Section I-A, this global attractivity result ensures that X converges with high probability to the vicinity of θ. As a result, the RM-BNN protocol and the payoff mechanism specified by $P(t) = \theta - X(t)$ is a solution to the distributed task allocation problem in Section II-C. Notably, $\theta \in \mathbb{NE}(\mathcal{F})$ means that, when $X(t) = \theta$, none of the agents can receive a higher payoff by unilaterally changing the types of tasks that they are performing.

VII. Conclusion

In this letter, we consider the mean population state resulting from potential games and EPT protocols and analyze the effects of strategy-dependent revision rates on its infinite-horizon properties. To do so, we investigate the (NS) and (PC) properties of the EDM induced by EPT protocols and employ the results in [6]. Contrary to the conclusions attained under identical revision rates [4], we show the existence of EPT protocols and (strategy-dependent) revision rates that lead to EDM instances that violate (NS) or (PC). Hence, we propose a modification of the EPT class, which we call sign preserving rate-modified EPT (RM-EPT) protocols, and show that the resulting EDM satisfies (NS) and (PC) under any revision rates. We use this result to derive a guarantee for the convergence of the mean population state to Nash equilibria.

Appendix

Before presenting the proofs of Theorems 1 and 2, we note that we are going to use the notation in Section I-D. Furthermore, we define Λ to be the n-by-n diagonal matrix with its (i, i)-th entry given by λ_i. Lastly, given a sign preserving RM-EPT protocol τ, we conform to the notation in Definition 6 and use φ to write it as in (6). We prove Theorem 1 below.

Proof: Assume that τ is a sign preserving RM-EPT protocol. Our aim is to show that the EDM satisfies $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$ if and only if $\mathcal{V}(\xi,\pi) = 0$. But first, we give an auxiliary result. Recall that $\lambda_i > 0$ for all $i \in \{1,\dots,n\}$. Hence, $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$ if and only if $\Lambda\xi/(\lambda^T\xi) \in \arg\max_{\eta\in\Delta}\eta^T\pi$. Moreover, from [4, Proposition 3.4], we know that $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$ if and only if $\hat{\pi} \in \delta(\mathbb{R}^n_*)$. Therefore, we arrive at $\tilde{\pi} \in \delta(\mathbb{R}^n_*)$ if and only if $\tilde{\pi} \in \delta(\mathbb{R}^n_*)$.

We return to the proof of (NS) and let $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$. Then, the auxiliary result above gives $\tilde{\pi} \in \delta(\mathbb{R}^n_*)$. Since φ is sign preserving, $\tilde{\pi} \in \delta(\mathbb{R}^n_*)$ implies that $\mathcal{V}(\xi,\pi) = 0$. Thus, $\mathcal{V}(\xi,\pi) = 0$ whenever $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$.

For the reverse direction, we prove the contrapositive. Assume that $\xi \notin \arg\max_{\eta\in\Delta}\eta^T\pi$ and define $J := \{j \in \{1,\dots,n\} \mid \tilde{\pi}_j > 0\}$. Note that $\xi \in \arg\max_{\eta\in\Delta}\eta^T\pi$ if and only if $(\Lambda\xi/(\lambda^T\xi))^T\pi \geq \pi_j$ for all $j \in \{1,\dots,n\}$. So, it follows from $\xi \notin \arg\max_{\eta\in\Delta}\eta^T\pi$ that $J \neq \emptyset$. We proceed by showing that this implies the existence of a $k \in \{1,\dots,n\}$ for which $\tilde{\pi}_k \leq 0$ and $\xi_k > 0$. Since $(\Lambda\xi/(\lambda^T\xi))^T\tilde{\pi} = 0$, we have $\sum_{j\in J}\lambda_j\xi_j\tilde{\pi}_j + \sum_{j\in J^c}\lambda_j\xi_j\tilde{\pi}_j = 0$, where we denote $J^c := \{1,\dots,n\}\setminus J$. There are two possible cases: either $\xi_j = 0$ for all $j \in J$ or there exists $i \in J$ such that $\xi_i > 0$. The first case implies, because $\sum_{j=1}^n \xi_j = 1$, that there is a $k \in J^c$ such that $\xi_k > 0$. Now consider the second case, i.e., there is an $i \in J$ such that $\xi_i > 0$. Then, due to $i \in J$ and $\sum_{j\in J}\lambda_j\xi_j\tilde{\pi}_j + \sum_{j\in J^c}\lambda_j\xi_j\tilde{\pi}_j = 0$, it must be that $\sum_{j\in J^c}\lambda_j\xi_j\tilde{\pi}_j < 0$. Consequently, there exists $k \in J^c$ such that $\xi_k > 0$. Hence, in both cases, there exists a $k \in \{1,\dots,n\}$ for which $\tilde{\pi}_k \leq 0$ and $\xi_k > 0$. Finally, we leverage this to show that $\mathcal{V}(\xi,\pi) \neq 0$. Since φ is sign preserving and $\tilde{\pi}_k \leq 0$ we have that $\varphi_k(\tilde{\pi}) = 0$. Therefore,

$$\begin{aligned}\mathcal{V}_k(\xi,\pi) &= \sum_{j=1,j\neq k}^n \lambda_j \tfrac{1}{\bar{\tau}}\varphi_k(\tilde{\pi})\xi_j - \sum_{j=1,j\neq k}^n \lambda_k \tfrac{1}{\bar{\tau}}\varphi_j(\tilde{\pi})\xi_k \\ &= -\lambda_k\xi_k\tfrac{1}{\bar{\tau}}\sum_{j=1,j\neq k}^n \varphi_j(\tilde{\pi}). \end{aligned} \tag{7}$$

Observe that $k \notin J$, together with $J \neq \emptyset$, implies $\sum_{j=1,j\neq k}^n \varphi_j(\tilde{\pi}) > 0$. Combining this with $\xi_k > 0$, $\bar{\tau} > 0$ and $\lambda_i > 0$ for all $i \in \{1,\dots,n\}$, we conclude that (7) is negative. As a result, (ξ,π) is not a rest point of $\mathcal{V}$. ■

Now, we present a proof of Theorem 2.

Proof: Let τ be a sign preserving RM-EPT protocol. Our aim is to show that the EDM satisfies $\mathcal{V}(\xi,\pi)^T\pi > 0$ whenever $\mathcal{V}(\xi,\pi) \neq 0$. Observe that

$$\mathcal{V}(\xi,\pi)^T\pi = \left(\tfrac{\Lambda\xi}{\lambda^T\xi}\right)^T\pi\sum_{i=1}^n \mathcal{V}_i(\xi,\pi) + \sum_{i=1}^n \tilde{\pi}_i\mathcal{V}_i(\xi,\pi).$$

From $\sum_{i=1}^n \mathcal{V}_i(\xi,\pi) = 0$, the above equality simplifies to

$$\mathcal{V}(\xi,\pi)^T\pi = \sum_{i=1}^n \tilde{\pi}_i\tfrac{1}{\bar{\tau}}\sum_{j=1,j\neq i}^n \big(\lambda_j\varphi_i(\tilde{\pi})\xi_j - \lambda_i\varphi_j(\tilde{\pi})\xi_i\big).$$

Furthermore, notice that

$$\begin{aligned}&\sum_{i=1}^n \tilde{\pi}_i\Big(\sum_{j=1,j\neq i}^n \lambda_j\tfrac{1}{\bar{\tau}}\varphi_i(\tilde{\pi})\xi_j - \sum_{j=1,j\neq i}^n \lambda_i\tfrac{1}{\bar{\tau}}\varphi_j(\tilde{\pi})\xi_i\Big) \\ &= \tfrac{1}{\bar{\tau}}\sum_{i=1}^n \tilde{\pi}_i\varphi_i(\tilde{\pi})\sum_{j=1,j\neq i}^n \lambda_j\xi_j - \tfrac{1}{\bar{\tau}}\sum_{i=1}^n \tilde{\pi}_i\lambda_i\xi_i\sum_{j=1,j\neq i}^n \varphi_j(\tilde{\pi}) \\ &= \tfrac{1}{\bar{\tau}}\sum_{i=1}^n \tilde{\pi}_i\varphi_i(\tilde{\pi})\sum_{j=1,j\neq i}^n \lambda_j\xi_j, \end{aligned} \tag{8}$$

where the last equality follows from $\sum_{i=1}^n \tilde{\pi}_i\lambda_i\xi_i = 0$. Now, assume that $\mathcal{V}(\xi,\pi) \neq 0$. Then, since sign preserving RM-EPT protocols satisfy (NS), there exists a strategy k for which $\tilde{\pi}_k > 0$. Moreover, it must be that $\xi_k < 1$, because otherwise $\tilde{\pi}_k$ would have been 0. These, combined with the sign preservation of φ, imply that

$$\sum_{i=1}^n \tilde{\pi}_i\varphi_i(\tilde{\pi})\sum_{j=1,j\neq i}^n \lambda_j\xi_j > \tilde{\pi}_k\varphi_k(\tilde{\pi})\sum_{j=1,j\neq k}^n \lambda_j\xi_j > 0.$$

Since $\bar{\tau}$ is also positive, we conclude that (8) is positive. ■

References

[1] W. H. Sandholm, *Population Games and Evolutionary Dynamics*. Cambridge, MA, USA: MIT Press, 2010.
[2] S. Park, N. C. Martins, and J. S. Shamma, “Payoff dynamics model and evolutionary dynamics model: Feedback and convergence to equilibria,” 2019, *arXiv:1903.02018*.
[3] S. Kara and N. C. Martins, “Pairwise comparison evolutionary dynamics with strategy-dependent revision rates: Stability and delta-passivity (expanded version),” 2021, *arXiv:2107.02835*.
[4] W. H. Sandholm, “Excess payoff dynamics and other well-behaved evolutionary dynamics,” *J. Econ. Theory*, vol. 124, no. 2, pp. 149–170, 2005.
[5] W. H. Sandholm, “Pairwise comparison dynamics and evolutionary foundations for Nash equilibrium,” *Games*, vol. 1, no. 1, pp. 3–17, 2010.
[6] W. H. Sandholm, “Potential games with continuous player sets,” *J. Econ. Theory*, vol. 97, pp. 81–108, Mar. 2001.
[7] J. W. Weibull, *Evolutionary Game Theory*. Cambridge, MA, USA: MIT Press, 1995.
[8] M. J. Fox and J. S. Shamma, “Population games, stable games, and passivity,” *Games*, vol. 4, pp. 561–583, Oct. 2013.
[9] M. Arcak and N. C. Martins, “Dissipativity tools for convergence to Nash equilibria in population games,” *IEEE Trans. Control Netw. Syst.*, vol. 8, no. 1, pp. 39–50, Mar. 2021.
[10] S. Kara, N. C. Martins, and M. Arcak, “Population games with Erlang clocks: Convergence to Nash equilibria for pairwise comparison dynamics,” 2022, *arXiv:2204.00593*.
[11] T. G. Kurtz, “Solutions of ordinary differential equations as limits of pure jump Markov processes,” *J. Appl. Probab.*, vol. 7, no. 1, pp. 49–58, Apr. 1970.
[12] D. Monderer and L. S. Shapley, “Potential games,” *Games Econ. Behav.*, vol. 14, pp. 124–143, May 1996.
[13] J. Hofbauer, “From Nash and Brown to Maynard Smith: Equilibria, dynamics and ESS,” *Selection*, vol. 1, no. 1, pp. 81–88, 2000.
[14] J. Barreiro-Gomez, N. Quijano, and C. Ocampo-Martinez, “Constrained distributed optimization: A population dynamics approach,” *Automatica*, vol. 69, pp. 101–116, Jul. 2016.
[15] T. Roughgarden, *Algorithmic Game Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007, ch. 18.
[16] A. Pashaie, L. Pavel, and C. J. Damaren, “Population dynamics approach for resource allocation problems,” in *Proc. Amer. Control Conf. (ACC)*, 2015, pp. 5231–5237.
[17] H. Tembine, E. Altman, R. ElAzouzi, and Y. Hayel, “Correlated evolutionarily stable strategies in random medium access control,” in *Proc. Int. Conf. Game Theory Netw.*, 2009, pp. 212–221.
[18] J. W. Weibull, *The Mass-Action Interpretation of Nash Equilibrium* (Working Paper Series). Stockholm, Sweden: Industriens Utredningsinstitut, 1995.
[19] G. W. Brown and J. von Neumann, “Solutions of games by differential equations,” RAND Corp., Santa Monica, CA, USA, Rep. P-142, Apr. 1950.
[20] J. Nash, “Non-cooperative games,” *Ann. Math.*, vol. 54, no. 2, pp. 286–295, 1951.
[21] S. Park, Y. D. Zhong, and N. E. Leonard, “Multi-robot task allocation games in dynamically changing environments,” in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2021, pp. 8678–8684.