

Learning equilibria with personalized incentives in a class of nonmonotone games

Filippo Fabiani, Andrea Simonetto and Paul J. Goulart

Abstract— We consider quadratic, nonmonotone generalized Nash equilibrium problems with symmetric interactions among the agents. Albeit this class of games is known to admit a potential function, its formal expression can be unavailable in several real-world applications. For this reason, we propose a two-layer Nash equilibrium seeking scheme in which a central coordinator exploits noisy feedback from the agents to design personalized incentives for them. By making use of those incentives, the agents compute a solution to an extended game, and then return feedback measures to the coordinator. We show that our algorithm returns an equilibrium if the coordinator is endowed with standard learning policies, and corroborate our results on a numerical instance of a hypomonotone game.

I. INTRODUCTION

Several multi-agent applications are characterized by symmetric interactions, in terms of cost incurred by each entity for a given service, across each pair of agents modelled within a noncooperative game-theoretic framework. Under some fairness condition of the electricity market in smart grids and demand-side management [1], [2], for instance, the cost per energy unit varies in the same way for all the assets, according to the current demand of the market. Likewise, in traffic or congestion games agents usually experience costs that only depend on the number of users occupying shared resources [3], [4]. Recently, symmetric barrier functions are also used to enforce proximity-based constraints among agents in several coordination and formation control applications [5], [6].

In this paper, we consider generalized Nash equilibrium problems (GNEPs), a multi-agent modelling paradigm in which selfish decision-makers compete for shared resources, characterized by symmetric interactions among the agents. The numerous benefit brought by such a symmetric structure to the GNEP are indeed well-known [7]. In particular, the presence of symmetries implies the existence of a *potential function*, which is key for two main reasons: i) it implicitly guarantees the existence of a generalized Nash equilibrium (GNE) for the GNEP at hand, which represents a desirable outcome of the game [8], and ii) in some cases it can be exploited directly in the design of Nash equilibrium seeking algorithms with convergence guarantees. This is a crucial feature in nonconvex/nonmonotone setting [2], [9], [10].

F. Fabiani and P. J. Goulart are with the Department of Engineering Science, University of Oxford, OX1 3PJ, United Kingdom ({filippo.fabiani, paul.goulart}@eng.ox.ac.uk). A. Simonetto is with the UMA, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France (andrea.simonetto@ensta-paris.fr). This work was partially supported through the Government's modern industrial strategy by Innovate UK, part of UK Research and Innovation, under Project LEO (Ref. 104781).

A mathematical expression of the underlying potential function is usually available only when the physics characterizing the GNEP at hand is known, or the agents' cost functions are suitably designed to result in a potential game [11]. In a general scenario, however, retrieving an expression of the potential function is a hard task [7, Ch. 2], or in some cases it may simply be unknown due to, e.g., privacy reasons.

Thus, in the simplified setting of a quadratic, nonmonotone (specifically, hypomonotone [12]) GNEP enjoying the existence of a potential function (§II), which is however assumed to be unavailable, we design a two-layer scheme that allows the agents to compute a GNE. Specifically, in the outer loop we endow a central coordinator with an online learning procedure, e.g., least squares (LS) or Gaussian process (GP), aiming at iteratively exploiting the noisy agents' feedback to learn some of their private characteristic, namely the (pseudo-)gradient mappings obtained from the agents' cost functions. In the spirit of [13]–[15], such a reconstructed information is hence exploited to iteratively design *personalized incentives* for the agents, which on their hand make use of those incentives to compute a GNE in the inner loop, and finally return noisy feedback measures to the central coordinator.

In the proposed setting we leverage the symmetries characterizing the interactions among agents to design parametric personalized incentives (§III) that play a key role to prove the convergence of the algorithm. In particular, i) they serve as regularization terms for the cost functions of the agents, thus enabling them for the practical computation of a variational generalized Nash equilibrium (v-GNE) in the inner loop; and ii) they allow to achieve faster convergence rates through an appropriate tuning of few key parameters. As main results, we prove that the proposed two-layer iterative scheme converges to a GNE by exploiting the consistency bounds typical of available learning procedures for the coordinator, as for example LS or GP (§IV). We remark here that GNE seeking in nonmonotone GNEP is, in general, a hard task. Available solution algorithms are indeed either tailored for Nash equilibrium problems (NEPs) with no coupling constraints, or involve generic variational inequalities (VIs), i.e., possibly not suitable for distributed computation [16]–[19]. We conclude the paper by testing our findings on a numerical instance of the considered hypomonotone GNEP (§V).

Notation: \mathbb{S}^n is the space of $n \times n$ symmetric matrices. For vectors $v_1, \dots, v_N \in \mathbb{R}^n$ and $\mathcal{I} = \{1, \dots, N\}$, we denote $\mathbf{v} := (v_1^\top, \dots, v_N^\top)^\top = \text{col}((v_i)_{i \in \mathcal{I}})$ and $\mathbf{v}_{-i} := \text{col}((v_j)_{j \in \mathcal{I} \setminus \{i\}})$. With a slight abuse of notation, we also use $\mathbf{v} = (v_i, \mathbf{v}_{-i})$. The mapping $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is monotone

on $\mathcal{X} \subseteq \mathbb{R}^n$ if $(F(x) - F(y))^\top(x - y) \geq 0$ for all $x, y \in \mathcal{X}$; strongly monotone if there exists a constant $c > 0$ such that $(F(x) - F(y))^\top(x - y) \geq c\|x - y\|^2$ for all $x, y \in \mathcal{X}$; hypomonotone if there exists a constant $c \geq 0$ such that $(F(x) - F(y))^\top(x - y) \geq -c\|x - y\|^2$ for all $x, y \in \mathcal{X}$.

II. MATHEMATICAL PROBLEM SETUP

We study a noncooperative game $\Gamma := (\mathcal{I}, (\mathcal{X}_i)_{i \in \mathcal{I}}, (J_i)_{i \in \mathcal{I}})$ with N agents, indexed by $\mathcal{I} := \{1, \dots, N\}$. Each agent $i \in \mathcal{I}$ controls variable x_i , constrained to a local set $\mathcal{X}_i \subseteq \mathbb{R}^{n_i}$, and aims at solving the following optimization problem:

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in \mathcal{X}_i} & J_i(x_i, \mathbf{x}_{-i}) \\ \text{s.t.} & A_i x_i + \sum_{j \in \mathcal{I} \setminus \{i\}} A_j x_j \leq b, \end{cases} \quad (1)$$

where, for computational purposes, each function $J_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $n := \sum_{i \in \mathcal{I}} n_i$, has a quadratic (possibly aggregative) form

$$J_i(x_i, \mathbf{x}_{-i}) := \frac{1}{2} x_i^\top Q_i x_i + \left(\sum_{j \in \mathcal{I} \setminus \{i\}} C_{i,j} x_j + q_i \right)^\top x_i,$$

for some $Q_i \in \mathbb{S}^{n_i}$, $C_{i,j} \in \mathbb{R}^{n_i \times n_j}$, and $q_i \in \mathbb{R}^{n_i}$. Thus, the collection of optimization problems in (1) amounts to a (quadratic) GNEP, where $A_i \in \mathbb{R}^{m \times n_i}$ constraints each agent to share common resources $b \in \mathbb{R}^m$ with the other agents. Let us define the sets $\mathcal{X} := \prod_{i \in \mathcal{I}} \mathcal{X}_i$, where each \mathcal{X}_i is assumed to be nonempty, compact and convex, $\mathcal{X}_i(\mathbf{x}_{-i}) := \{x_i \in \mathcal{X}_i \mid A_i x_i \leq b - \sum_{j \in \mathcal{I} \setminus \{i\}} A_j x_j\}$, and the feasible set of Γ , $\Omega := \{\mathbf{x} \in \mathcal{X} \mid A \mathbf{x} \leq b\}$, with $A := [A_1 \ \dots \ A_N] \in \mathbb{R}^{m \times n}$. We are interested in the following popular solution to a GNEP.

Definition 1: (Generalized Nash equilibrium [8]) A strategy vector $\mathbf{x}^* \in \Omega$ is a GNE of the game Γ if, for all $i \in \mathcal{I}$,

$$J_i(x_i^*, \mathbf{x}_{-i}^*) \leq \inf_{y_i \in \mathcal{X}_i(\mathbf{x}_{-i}^*)} J_i(y_i, \mathbf{x}_{-i}^*).$$

□

Thus, $\mathbf{x}^* \in \Omega$ is an equilibrium if no agent can decrease their cost by changing unilaterally x_{-i}^* to any other feasible point. In the remainder, we make the following assumption.

Standing Assumption 1: For all $(i, j) \in \mathcal{I}^2$, $C_{i,j} = C_{j,i}$. □

Under the condition of symmetric interactions among agents, we know that the game mapping associated to Γ , $G : \mathbb{R}^n \rightarrow \mathbb{R}^n$, also known as the pseudo-gradient mapping since it is formally defined as $G(\mathbf{x}) := \text{col}((\nabla_{x_i} J_i(x_i, \mathbf{x}_{-i}))_{i \in \mathcal{I}})$, admits a differentiable, yet possibly *unknown*, function $\theta : \mathbb{R}^n \rightarrow \mathbb{R}$ such that $G(\mathbf{x}) = \nabla \theta(\mathbf{x})$, for all $\mathbf{x} \in \Omega$ [20, Th. 1.3.1]. This latter coincides with a potential function for the game Γ [21], for which we define the set of constrained minimizers $\Theta := \text{argmin}_{\mathbf{y} \in \Omega} \theta(\mathbf{y})$, assumed to be nonempty, and set of constrained stationary points Θ^s . Clearly, $\Theta \subseteq \Theta^s$. Note that, according to Definition 1, any $\mathbf{x}^* \in \Theta$ coincides with a GNE for Γ , namely $\Theta \neq \emptyset$ guarantees the existence

of at least a GNE of the underlying game. By considering quadratic costs in (1), $G(\cdot)$ turns into an affine mapping:

$$G(\mathbf{x}) = \begin{bmatrix} Q_1 & \cdots & C_{1,N} \\ \vdots & \ddots & \vdots \\ C_{N,1} & \cdots & Q_N \end{bmatrix} \mathbf{x} + \begin{bmatrix} q_1 \\ \vdots \\ q_N \end{bmatrix} =: Q\mathbf{x} + q.$$

We now introduce the function $\theta : \mathbb{R}^n \rightarrow \mathbb{R}$, which can be characterized as stated immediately below:

$$\theta(\mathbf{x}) := \sum_{i \in \mathcal{I}} \left(\frac{1}{2} x_i^\top Q_i x_i + q_i^\top x_i + \sum_{j \in \mathcal{I}, j < i} (C_{i,j} x_j)^\top x_i \right). \quad (2)$$

Proposition 1: For the function $\theta(\mathbf{x})$ in (2) we have that:

- i) It amounts to a potential function for the GNEP Γ ;
- ii) It is ℓ -weakly convex, with $\ell := |\lambda_{\min}(Q)|$.

□

Proof: i) This part of the proof is akin to [22, Prop. 2] and therefore, due to space limitations, it is here omitted.

ii) We start from the definition of hypomonotonicity for the mapping $\nabla \theta(\cdot)$, which requires the existence of some $\rho > 0$ such that $(\nabla \theta(\mathbf{x}) - \nabla \theta(\mathbf{y}))^\top(\mathbf{x} - \mathbf{y}) \geq -\rho\|\mathbf{x} - \mathbf{y}\|^2$, for all $\mathbf{x}, \mathbf{y} \in \Omega$. Since $\nabla \theta(\mathbf{x}) = G(\mathbf{x}) = Q\mathbf{x} + q$,

$$\begin{aligned} (\nabla \theta(\mathbf{x}) - \nabla \theta(\mathbf{y}))^\top(\mathbf{x} - \mathbf{y}) &= (Q(\mathbf{x} - \mathbf{y}))^\top(\mathbf{x} - \mathbf{y}) \\ &= \|\mathbf{x} - \mathbf{y}\|_Q^2 \geq \lambda_{\min}(Q)\|\mathbf{x} - \mathbf{y}\|^2. \end{aligned}$$

With $\lambda_{\min}(Q)$ that, in principle, may be nonpositive (since each Q_i is only symmetric), we obtain $(\nabla \theta(\mathbf{x}) - \nabla \theta(\mathbf{y}))^\top(\mathbf{x} - \mathbf{y}) \geq -|\lambda_{\min}(Q)|\|\mathbf{x} - \mathbf{y}\|^2$. In view of [23, Lemma 2.1], the $|\lambda_{\min}(Q)|$ -hypomonotonicity of the mapping $\nabla \theta(\cdot)$ is equivalent to the $|\lambda_{\min}(Q)|$ -weak convexity of the function $\theta(\cdot)$, thus concluding the proof. ■

It follows from Proposition 1 that the quadratic GNEP Γ in (1) turns into a hypomonotone GNEP [12] in which the mapping $Q\mathbf{x} + q$ may not be monotone, since $Q + Q^\top$ is not positive semi-definite in general. As a consequence, the function $\theta(\cdot)$ may not be convex, thus posing several limitations in the design of an equilibrium seeking procedure for Γ with most of the available techniques (monotonicity is, indeed, one of the weakest requirements [20]).

On the other hand, we note that Standing Assumption 1 is key to claim that the underlying quadratic GNEP admits a potential function. This latter is key for two main reasons: i) to claim that the game at hand possesses (at least) an equilibrium (in our case, $\Theta \neq \emptyset$), and ii) to design Nash equilibrium seeking algorithms with convergence guarantees, especially in nonconvex/nonmonotone setting, e.g., [2], [9], [10]. In practical applications, however, may not be always reasonable to rely on a formal expression of the potential function [7, Ch. 2] due to privacy reasons for instance. In our mathematical developments, we hence treat the GNEP in (1) as if we were unaware of the fact that it is potential, and therefore the expression for $\theta(\cdot)$ in (2) can not be directly exploited for algorithm design. By relying on standard learning paradigms, in the next section we design a two-layer procedure to steer the agents on some point falling into the

Algorithm 1: Two-layer semi-decentralized method

Initialization: Set $t = 0$, $\mathbf{x}_t^* \in \Omega$, set sequences $(c_t)_{t \in \mathbb{N}} > 2\ell$ ($\xi_t)_{t \in \mathbb{N}} \in [0, 1/c_t)$, choose some $T \in \mathbb{N}$

Iteration ($t \in \mathbb{N}$):

- Integrate recent measures to learn $(\nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*))_{i \in \mathcal{I}}$
 - Compute personalized incentives $(u_{i,t}(\mathbf{x}))_{i \in \mathcal{I}}$
 - Compute a v-GNE of the extended game $\bar{\Gamma}$, $\mathbf{x}_t^* \in \Omega$
 - Obtain noisy agents' feedback $\{(x_{i,t}^*, p_{i,t})\}_{i \in \mathcal{I}}$
-

set Θ , which in view of Definition 1 coincides with a GNE of Γ .

III. THE TWO-LAYER LEARNING PROCEDURE

From the discussion in §II, we identify two critical challenges in designing a Nash equilibrium seeking method to compute a GNE for Γ : 1) the non-monotonicity of $Q\mathbf{x} + q$, and 2) the unknown expression of the potential function θ in (2). We propose a way to circumvent both issues by designing *personalized feedback functionals* $u_i : \mathbb{R}^n \rightarrow \mathbb{R}$ in the spirit of [13], which can be used as “control actions” as described next.

A. The algorithm

We summarize the key steps of the proposed, two-layer approach in Algorithm 1. In particular, black-filled bullets refer to the tasks that have to be performed by a central coordinator, while the empty bullet to the one performed by the agents in \mathcal{I} . After the initialization phase, the coordinator aims at learning online (i.e., while the algorithm is running) the gradient mapping of the unknown function θ by leveraging possibly noisy agents' feedback on the private functions J_i 's in the outer loop. Then, on the basis of the estimated $\hat{J}_{i,t}$, the central entity designs personalized incentive functionals $u_{i,t}$, thus forcing the agents to face with an extended version of the quadratic GNEP Γ in (1), i.e., $\bar{\Gamma} := (\mathcal{I}, (\mathcal{X}_i)_{i \in \mathcal{I}}, (\bar{J}_{i,t})_{i \in \mathcal{I}})$, which stem from replacing each cost function $J_i(x_i, \mathbf{x}_{-i})$ in (1) with $J_i(x_i, \mathbf{x}_{-i}) + u_{i,t}(x_i, \mathbf{x}_{-i}) =: \bar{J}_{i,t}(x_i, \mathbf{x}_{-i})$.

By suitably choosing such incentives, we will show that they serve as regularization terms for $J_i(\cdot, \mathbf{x}_{-i})$, as well as they trade-off convergence properties of Algorithm 1 and robustness to the imperfect knowledge of the potential function θ and associated gradient. Specifically, personalized incentives enable the agents for the practical computation of an equilibrium of the extended game $\bar{\Gamma}$ via standard solution procedures for GNEPs [12], [24], [25]. In the cognate literature, however, algorithmic methods typically return a v-GNE [8], which coincides to any solution to the (extended) GNEP $\bar{\Gamma}$ that is also a solution to the associated VI, i.e., any collective vector of strategies $\mathbf{x}_t^* \in \Omega$ such that, for all $t \in \mathbb{N}$,

$$(\mathbf{y} - \mathbf{x}_t^*)^\top \bar{G}_t(\mathbf{x}_t^*) \geq 0, \text{ for all } \mathbf{y} \in \Omega, \quad (3)$$

where the mapping $\bar{G}_t : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is formally obtained as $\bar{G}_t(\mathbf{x}) := \text{col}((\nabla_{x_i} \bar{J}_{i,t}(x_i, \mathbf{x}_{-i}))_{i \in \mathcal{I}}) = Q\mathbf{x} + q + U_t(\mathbf{x})$, and $U_t(\mathbf{x}) := \text{col}((\nabla_{x_i} u_{i,t}(x_i, \mathbf{x}_{-i}))_{i \in \mathcal{I}})$. This is why the computational step in Algorithm 1 involving the agents requires them to implement an available procedure to compute a v-GNE of the extended game $\bar{\Gamma}$. As a last step, the proposed method requires the coordinator to retrieve feedback measures and equilibrium strategies from the agents, $\{(x_{i,t}^*, p_{i,t})\}_{i \in \mathcal{I}}$, with $p_{i,t} := J_i(\mathbf{x}_t^*) + \varepsilon_{i,t}$ and random variable ε_i .

B. Personalized incentives design

We follow the approach in [14, Alg. 1] by endowing the central coordinator with a *learning procedure* \mathcal{L} such that, at every outer iteration $t \in \mathbb{N}$ of Algorithm 1, it integrates the most recent noisy agents' feedback $\{p_{i,t-1}\}_{i \in \mathcal{I}}$ to calculate an estimate of the gradients $(\nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*))_{i \in \mathcal{I}}$ (since $G(\mathbf{x}) = \nabla\theta(\mathbf{x})$ [20, Th. 1.3.1]). Then, it exploits such estimates to design personalized functionals as follows:

$$u_{i,t}(\mathbf{x}) = \frac{1}{2}c_t \|x_i - x_{i,t}^+\|^2, \text{ for all } i \in \mathcal{I}, \quad (4)$$

where $x_{i,t}^+ := x_{i,t-1}^* + \xi_t \nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*)$, for some $c_t, \xi_t \geq 0$, for all $t \in \mathbb{N}$. Given the parametric form in (4), our goal hence reduces to design suitable conditions for the gain c_t and step-size ξ_t in such a way that the sequence of v-GNE, $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$, monotonically decreases (or non-increases) the unknown potential function $\theta(\cdot)$, and asymptotically converges to one of its constrained minimizers (either local or global).

The structure of the proposed personalized incentives brings us to the following considerations: i) inspired by recent algorithms based on the Heavy Anchor method, e.g., [12], we note that each term $x_{i,t}^+$ requires a positive sign for the gradient step $\xi_t \nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*)$. In the next sections we will show that such a choice enables us to boost the convergence of Algorithm 1, or to reduce the asymptotic error by means of a suitable tuning of ξ_t . In addition, ii) we note that c_t is crucial to allow the agents for the practical computation of a v-GNE through available iterative schemes, as stated next:

Proposition 2: Let $c_t \geq 2\ell$ for all $t \in \mathbb{N}$. Then, with the personalized incentives in (4), the mapping $\bar{G}_t(\cdot)$ is ℓ -strongly monotone, for all $t \in \mathbb{N}$. \square

Proof: By using the definition of incentive functionals in (4), we obtain $\bar{G}_t(\mathbf{x}) = G(\mathbf{x}) + c_t(\mathbf{x} - \mathbf{x}_{t-1}^* - \xi_t \hat{G}_{t-1}(\mathbf{x}_{t-1}^*))$. Thus, for any $\mathbf{x}, \mathbf{y} \in \Omega$, and $t \in \mathbb{N}$, we have that:

$$\begin{aligned} & (\mathbf{x} - \mathbf{y})^\top (\bar{G}_t(\mathbf{x}) - \bar{G}_t(\mathbf{y})) \\ &= (\mathbf{x} - \mathbf{y})^\top (G(\mathbf{x}) + c_t(\mathbf{x} - \mathbf{x}_{t-1}^* - \xi_t \hat{G}_{t-1}(\mathbf{x}_{t-1}^*)) \\ & \quad - G(\mathbf{y}) - c_t(\mathbf{y} - \mathbf{x}_{t-1}^* - \xi_t \hat{G}_{t-1}(\mathbf{x}_{t-1}^*))) \\ &= (\mathbf{x} - \mathbf{y})^\top (G(\mathbf{x}) + c_t\mathbf{x} - G(\mathbf{y}) - c_t\mathbf{y}) \\ &= (\mathbf{x} - \mathbf{y})^\top (\nabla\theta(\mathbf{x}) + c_t\mathbf{x} - \nabla\theta(\mathbf{y}) - c_t\mathbf{y}), \end{aligned}$$

where the last equality follows from Standing Assumption 1 and [20, Th. 1.3.1]. Let us now consider the auxiliary function $\psi_t(\mathbf{x}) := \theta(\mathbf{x}) + \frac{c_t}{2} \|\mathbf{x}\|^2$. In case $c_t \geq 2\ell$ for all

$t \in \mathbb{N}$, $\psi_i(\cdot)$ turns out to be ℓ -strongly convex in view of the ℓ -weak convexity of θ (Proposition 1), which directly leads to

$$\begin{aligned} (\mathbf{x} - \mathbf{y})^\top (\bar{G}_t(\mathbf{x}) - \bar{G}_t(\mathbf{y})) &= (\mathbf{x} - \mathbf{y})^\top (\nabla\theta(\mathbf{x}) + c_t\mathbf{x} - \nabla\theta(\mathbf{y}) - c_t\mathbf{y}) \\ &= (\mathbf{x} - \mathbf{y})^\top (\nabla\psi_t(\mathbf{x}) - \nabla\psi_t(\mathbf{y})) \geq \ell\|\mathbf{x} - \mathbf{y}\|^2, \end{aligned}$$

i.e., the definition of strongly monotone mapping for \bar{G}_t . ■

It follows that, if c_t is large enough for all $t \in \mathbb{N}$, then at every outer iteration the agents compute the unique [20, Th. 2.3.3] v-GNE associated to the extended GNEP $\bar{\Gamma}$.

IV. MAIN RESULTS

We now establish convergence results for Algorithm 1 by distinguishing between two cases: online perfect reconstruction, if the learning procedure \mathcal{L} allows the central coordinator to leverage $\nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*) = \nabla_{x_i} J_i(\mathbf{x}_{t-1}^*)$, for all $i \in \mathcal{I}$ and $t \in \mathbb{N}$, and imperfect reconstruction. To this end, we introduce the fixed point residual $\Delta_t^* := \mathbf{x}_t^* - \mathbf{x}_{t-1}^*$ as a key quantity to “measure” the distance between the strategy profile at the current iteration and the points in Θ^s .

Lemma 1: Let $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$ be the sequence of v-GNE generated by Algorithm 1 with $c_t \xi_t < 1$. Moreover, assume perfect reconstruction of the mapping $Q\mathbf{x} + q$, and that there exists some $\mathbf{x}_t^* \in \Omega$ such that $\|\Delta_t^*\| = 0$. Then, some $\mathbf{x}_t^* \in \Omega$, i.e., \mathbf{x}_t^* is a stationary point for the function $\theta(\mathbf{x})$. □

Proof: Any stationary point of the function $\theta(\mathbf{x})$ satisfies the following first-order optimality condition:

$$w(Q\mathbf{x} + q) + \mathcal{N}_\Omega(\mathbf{x}) \ni \mathbf{0},$$

where $w > 0$ scales the cost associated to $\theta(\cdot)$, while \mathcal{N}_Ω denotes the normal cone operator of the feasible set Ω . At every $t \in \mathbb{N}$, the v-GNE \mathbf{x}_t^* obtained with perfect reconstruction satisfies $(Q\mathbf{x}_t^* + q) + c_t(\mathbf{x}_t^* - \mathbf{x}_{t-1}^* - \xi_t(Q\mathbf{x}_{t-1}^* + q)) + \mathcal{N}_\Omega(\mathbf{x}_t^*) \ni \mathbf{0}$. If $\|\Delta_t^*\| = 0$ then we have $\mathbf{x}_t^* = \mathbf{x}_{t-1}^*$, namely \mathbf{x}_t^* is the solution to the inclusion

$$(1 - c_t \xi_t)(Q\mathbf{x}_t^* + q) + \mathcal{N}_\Omega(\mathbf{x}_t^*) \ni \mathbf{0},$$

thus satisfying the first-order optimality conditions for $\theta(\mathbf{x})$, and therefore it amounts to one of its stationary points. ■

In case of inexact reconstruction of the gradient mappings, as a metric for the convergence of the sequence $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$ to the stationary point set we also adopt the average value of $\|\Delta_t^*\|$ over a certain iterations horizon of length $T \in \mathbb{N}$, formally defined as $\frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^*\|$, $\mathcal{T} := \{1, \dots, T\}$. For the imperfect reconstruction case, our result will hence be of the form $(1/T) \sum_{t \in \mathcal{T}} \|\Delta_t^*\| = O(1)$ if the reconstruction error is persistent, while $(1/T) \sum_{t \in \mathcal{T}} \|\Delta_t^*\| = 0$ otherwise.

Remark 1: Since our algorithm produces monotonically non-increasing values for θ through the iterations, it is worth mentioning that the application of traditional perturbation techniques, such as in [26], can ensure that the stationary points to which we converge are, in practice, local minima in Θ^s , and hence GNE of the original game Γ . □

A. Perfect reconstruction of the pseudo-gradients

In the desirable, yet unrealistic, case in which the procedure \mathcal{L} enables for the perfect reconstruction, i.e., $\nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*) = \nabla_{x_i} J_i(\mathbf{x}_{t-1}^*)$, for all $i \in \mathcal{I}$ and $t \in \mathbb{N}$, we show that a careful choice of the gain c_t and the step-size ξ_t characterizing the personalized incentives in (4) allows Algorithm 1 to produce a convergent sequence of v-GNE:

Proposition 3: For all $t \in \mathbb{N}$, let $c_t \geq 2\ell$ and $\xi_t \in [0, 1/c_t)$. Then, with the incentives in (4) the sequence of v-GNE $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$ generated by Algorithm 1 has limit point in Θ^s . □

Proof: We first show that, if c_t (resp., ξ_t) is large (small) enough, the personalized functionals in (4) allow to point a descent direction for the unknown potential function θ , i.e., $\Delta_t^{*\top} \nabla\theta(\mathbf{x}_{t-1}^*) < 0$. Since \mathbf{x}_t^* amounts to a v-GNE at every outer iteration $t \in \mathbb{N}$, we have by (3) that $(\mathbf{y} - \mathbf{x}_t^*)^\top \bar{G}_t(\mathbf{x}_t^*) \geq 0$ for all $\mathbf{y} \in \Omega$. Then, since $\mathbf{x}_{t-1}^* \in \Omega$ as it is a v-GNE at $t-1$, we also have $\Delta_t^{*\top} \bar{G}_t(\mathbf{x}_t^*) \leq 0$. By adding and subtracting the term $\Delta_t^{*\top} G(\mathbf{x}_{t-1}^*)$, we hence obtain

$$\Delta_t^{*\top} G(\mathbf{x}_{t-1}^*) \leq \Delta_t^{*\top} (G(\mathbf{x}_{t-1}^*) - (G(\mathbf{x}_t^*) + U_t(\mathbf{x}_t^*))).$$

In addition, by combining the definition of the incentives in (4) with the perfect reconstruction of the pseudo-gradients, we also have that $U_t(\mathbf{x}_t^*) = c_t(\Delta_t^* - \xi_t G(\mathbf{x}_{t-1}^*))$. Then, in view of the symmetry postulated in Standing Assumption 1, it follows that

$$\begin{aligned} (1 - c_t \xi_t) \Delta_t^{*\top} \nabla\theta(\mathbf{x}_{t-1}^*) &\leq \Delta_t^{*\top} (\nabla\theta(\mathbf{x}_{t-1}^*) + c_t \mathbf{x}_{t-1}^* - \nabla\theta(\mathbf{x}_t^*) - c_t \mathbf{x}_t^*). \end{aligned}$$

We now define $\alpha_t := 1 - c_t \xi_t$ for all $t \in \mathbb{N}$, and introduce the auxiliary function $\psi_t(\mathbf{x}) := \theta(\mathbf{x}) + \frac{c_t}{2} \|\mathbf{x}\|^2$, which is known to be ℓ -strongly convex in case $c_t \geq 2\ell$ (Proposition 1). For this reason, we have $\Delta_t^{*\top} (-\nabla\psi_t(\mathbf{x}_{t-1}^*) + \nabla\psi_t(\mathbf{x}_t^*)) \geq \ell \|\Delta_t^*\|^2$, and hence

$$\alpha_t \Delta_t^{*\top} \nabla\theta(\mathbf{x}_{t-1}^*) \leq \Delta_t^{*\top} (\nabla\psi_t(\mathbf{x}_{t-1}^*) - \nabla\psi_t(\mathbf{x}_t^*)) \leq -\ell \|\Delta_t^*\|^2.$$

By imposing $\alpha_t > 0$, yielding to $\xi_t < 1/c_t \leq 1/2\ell$, for all $t \in \mathbb{N}$, we finally obtain $\Delta_t^{*\top} \nabla\theta(\mathbf{x}_{t-1}^*) \leq -(\ell/\alpha_t) \|\Delta_t^*\|^2 < 0$. This result can be then combined with the descent lemma [27, Prop. A.24] to claim that the sequence $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$ satisfies

$$\begin{aligned} \theta(\mathbf{x}_t^*) &\leq \theta(\mathbf{x}_{t-1}^*) + (\Delta_t^*)^\top \nabla\theta(\mathbf{x}_{t-1}^*) + \frac{\ell}{2} \|\Delta_t^*\|^2 \\ &\leq \theta(\mathbf{x}_{t-1}^*) - \ell \frac{2 - \alpha_t}{2\alpha_t} \|\Delta_t^*\|^2. \end{aligned} \quad (5)$$

Thus, by imposing $(2 - \alpha_t)/2\alpha_t > 0$, which implies that $0 \leq \xi_t < 1/c_t \leq 1/2\ell$, it must happen that the sequence $(\theta(\mathbf{x}_t^*))_{t \in \mathbb{N}}$ tends to a finite value, as $\theta(\mathbf{x}_t^*) \rightarrow -\infty$ can not happen in view of the boundedness of Ω . Since θ is continuous, the convergence of $(\theta(\mathbf{x}_t^*))_{t \in \mathbb{N}}$ implies that $\lim_{t \rightarrow \infty} \|\Delta_t^*\| = 0$, and hence the bounded sequence of feasible points $(\mathbf{x}_t^*)_{t \in \mathbb{N}} \in \Omega$ (since any \mathbf{x}_t^* is a v-GNE of $\bar{\Gamma}$) has a limit point in Θ^s in view of Lemma 1. ■

From (5) it is clear that making the product $c_t \xi_t$ close to one allows us to actually boost the convergence of Algorithm 1 to some point in Θ^s . This indeed supports the choice for a positive sign in the gradient step of (4). In view of the presence of noise in the agents’ feedback $\{p_{i,t-1}\}_{i \in \mathcal{I}}$,

however, at least at the beginning of the iterative scheme in Algorithm 1 it seems unlikely that the learning procedure \mathcal{L} guarantees a perfect reconstruction of $(\nabla_{x_i} \hat{J}_{i,t-1}(\mathbf{x}_{t-1}^*))_{i \in \mathcal{I}}$. For this reason, we investigate next the effect of the inexact reconstruction on the convergence properties of Algorithm 1.

B. Inexact estimate of the pseudo-gradients

In this section we leverage a further assumption characterizing the reconstruction of the pseudo-gradient mappings. Inspired by [14], [28], however, we make the following assumption on $\hat{G}_t(\cdot)$ directly, rather than on each gradient.

Assumption 1: For all $t \in \mathbb{N}$ and $\mathbf{x} \in \mathcal{X}$, $\hat{G}_t(\mathbf{x}) := Q\mathbf{x} + q + \epsilon_t$, and, for any $\delta \in (0, 1]$, there exists $\bar{t} < \infty$ such that $\mathbb{P}\{\|\epsilon_t\| \leq e(\bar{t}) \mid \forall t \geq \bar{t}\} \geq 1 - \delta$, for some nonincreasing function $e: \mathbb{N} \rightarrow \mathbb{R}_{\geq 0}$ such that $e(t) < \infty$, for all $t \in \mathbb{N}$. \square

With Assumption 1 we essentially require that, for all $\mathbf{x} \in \mathcal{X}$, the reconstruction error on $Q\mathbf{x} + q$ made by the learning procedure \mathcal{L} is bounded with (possibly high) probability $1 - \delta$ by some function of the available agents' feedback. Then, by introducing quantities $\kappa_t := (1 - \alpha_t)/2\alpha_t$ and $\beta_t := \ell(2 - \alpha_t)/2\alpha_t$, we can prove the following result.

Lemma 2: Let $c_t \geq 2\ell$ and $\xi_t \in [0, 1/c_t]$ for all $t \in \mathbb{N}$, and let Assumption 1 be valid for some $\delta \in (0, 1]$. Then, the sequence of v-GNE $(\mathbf{x}_t^*)_{t \in \mathbb{N}}$ generated by Algorithm 1 with personalized incentives in (4) satisfies, with probability $1 - \delta$,

$$\Delta_t^{*\top} \nabla \theta(\mathbf{x}_{t-1}^*) \leq \frac{\alpha_t \kappa_t^2}{\ell} e^2(t-1) - \left(\sqrt{\frac{\ell}{\alpha_t}} \|\Delta_t^*\| - \kappa_t \sqrt{\frac{\alpha_t}{\ell}} e(t-1) \right)^2. \quad (6)$$

\square

Proof: We start by making use of the same steps performed at the beginning of the proof of Proposition 3, which in view of Assumption 1 leads to $\alpha_t \Delta_t^{*\top} \nabla \theta(\mathbf{x}_{t-1}^*) \leq \Delta_t^{*\top} (\nabla \theta(\mathbf{x}_{t-1}^*) + c_t \mathbf{x}_{t-1}^* - \nabla \theta(\mathbf{x}_t^*) - c_t \mathbf{x}_t^*) + c_t \xi_t \Delta_t^{*\top} \epsilon_{t-1}$, for all $t \geq \bar{t}$. Therefore, if $c_t \geq 2\ell > 0$ for all $t \in \mathbb{N}$, we can upper bound $\Delta_t^{*\top} (\nabla \theta(\mathbf{x}_{t-1}^*) + c_t \mathbf{x}_{t-1}^* - \nabla \theta(\mathbf{x}_t^*) - c_t \mathbf{x}_t^*)$ with $-\ell \|\Delta_t^*\|^2$ due to the ℓ -strong convexity of the auxiliary function $\psi_t(\cdot)$. Since $\xi_t \geq 0$, $\Delta_t^{*\top} \epsilon_{t-1}$ attains its maximum positive module when Δ_t^* and ϵ_{t-1} are aligned, and hence we obtain the following chain of inequalities

$$\begin{aligned} \alpha_t \Delta_t^{*\top} \nabla \theta(\mathbf{x}_{t-1}^*) &\leq -\ell \|\Delta_t^*\|^2 + c_t \xi_t \|\Delta_t^*\| \|\epsilon_{t-1}\| \\ &\leq -\ell \|\Delta_t^*\|^2 + c_t \xi_t \|\Delta_t^*\| e(t-1). \end{aligned}$$

Note that the last inequality holds with probability $1 - \delta$, as it follows directly from Assumption 1, for any $t \geq \bar{t}$. Thus, for $\alpha_t > 0$, i.e., $\xi_t < 1/c_t$, and $c_t \xi_t = 1 - \alpha_t$, we obtain

$$\Delta_t^{*\top} \nabla \theta(\mathbf{x}_{t-1}^*) \leq -\frac{\ell}{\alpha_t} \|\Delta_t^*\|^2 + \frac{1 - \alpha_t}{\alpha_t} \|\Delta_t^*\| e(t-1). \quad (7)$$

Finally, completing the square in the RHS of (7) by adding and subtracting $((1 - \alpha_t)^2 / 4\alpha_t \ell) e^2(t-1)$ leads to

$$\begin{aligned} \Delta_t^{*\top} \nabla \theta(\mathbf{x}_{t-1}^*) &\leq -\left(\sqrt{\frac{\ell}{\alpha_t}} \|\Delta_t^*\| - (1 - \frac{\alpha_t}{2\alpha_t}) \sqrt{\frac{\alpha_t}{\ell}} e(t-1) \right)^2 \\ &\quad + \frac{(1 - \alpha_t)^2}{4\alpha_t \ell} e^2(t-1). \end{aligned}$$

The proof ends by substituting κ_t in the inequality above. \blacksquare

From the relation in (6), we note that the vector Δ_t^* does not necessarily point a descent direction for the function θ . The term $e^2(t-1)$, in fact, prevents the LHS in (6) from being strictly negative, though it can be made small through κ_t by means of a suitable choice of the step-size ξ_t . By introducing the quantities $\bar{\beta} := \sum_{t \in \mathcal{T}} \beta_t$ and $\underline{\beta} := \min_{t \in \mathcal{T}} \beta_t$ over some horizon \mathcal{T} of length $T \geq 0$, the following result characterizes the sequence of v-GNE generated by Algorithm 1.

Theorem 1: Let $c_t \geq 2\ell$ and $\xi_t \in [0, 1/c_t]$, for all $t \in \mathbb{N}$, and let Assumption 1 be valid for some $\delta \in (0, 1]$. Let some $T \in \mathbb{N}$ be fixed, $\mathcal{T} := \{\bar{t} + 1, \dots, T + \bar{t}\}$ and, for any global minimizer $\mathbf{x}^* \in \Theta$, $\Delta_{\bar{t}} := \theta(\mathbf{x}_{\bar{t}}^*) - \theta(\mathbf{x}^*)$. The sequence of v-GNE $(\mathbf{x}_t^*)_{t \in \mathcal{T}}$ generated by Algorithm 1 with the personalized incentives in (4) satisfies, with probability $1 - \delta$,

$$\begin{aligned} \frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^*\| &\leq \frac{1}{T \underline{\beta}} \sqrt{\sum_{t \in \mathcal{T}} \left(\beta_t \Delta_{\bar{t}} + \frac{\bar{\beta} \kappa_t^2}{\beta_t} e^2(t) \right)} \\ &\quad + \frac{1}{T \underline{\beta}} \sum_{t \in \mathcal{T}} \kappa_t e(t). \end{aligned} \quad (8)$$

\square

Proof: From Assumption 1, by combining the descent lemma [27, Prop. A.24] and the relation in (7) we obtain

$$\theta(\mathbf{x}_t^*) \leq \theta(\mathbf{x}_{t-1}^*) - \ell \frac{2 - \alpha_t}{2\alpha_t} \|\Delta_t^*\|^2 + \frac{1 - \alpha_t}{\alpha_t} \|\Delta_t^*\| e(t-1)$$

with probability $1 - \delta$, for all $t \geq \bar{t}$. Then, by exploiting the definition of β_t , which is strictly greater than zero if $\xi_t \in [0, 1/c_t]$, and by focusing on $\ell(2 - \alpha_t)/2\alpha_t \|\Delta_t^*\|^2$ and $(1 - \alpha_t)/\alpha_t \|\Delta_t^*\| e(t-1)$, we can complete the square by adding and subtracting $((1 - \alpha_t)^2 / 4\alpha_t^2 \beta_t) e^2(t-1)$. With standard manipulations one obtains the following expression

$$\begin{aligned} \theta(\mathbf{x}_t^*) &\leq \theta(\mathbf{x}_{t-1}^*) - \beta_t \left(\|\Delta_t^*\| - \frac{1 - \alpha_t}{2\alpha_t \beta_t} e(t-1) \right)^2 \\ &\quad + \frac{(1 - \alpha_t)^2}{4\alpha_t^2 \beta_t} e^2(t-1). \end{aligned} \quad (9)$$

We now fix $T \in \mathbb{N}$ and $\kappa_t = (1 - \alpha_t)/2\alpha_t$, and for any $\mathbf{x}^* \in \Theta$, by summing up the inequality above over $t \in \mathcal{T} := \{\bar{t} + 1, \dots, T + \bar{t}\}$, with probability $1 - \delta$ we obtain $\theta(\mathbf{x}^*) \leq \theta(\mathbf{x}_T^*) \leq \theta(\mathbf{x}_{\bar{t}}^*) - \sum_{t \in \mathcal{T}} \beta_t (\|\Delta_t^*\| - \frac{\kappa_t}{\beta_t} e(t))^2 + \sum_{t \in \mathcal{T}} \frac{\kappa_t^2}{\beta_t} e^2(t)$. Then, by moving $\theta(\mathbf{x}^*)$ to the RHS, the term with β_t to the LHS, and by introducing $\Delta_{\bar{t}} := \theta(\mathbf{x}_{\bar{t}}^*) - \theta(\mathbf{x}^*)$, a quantity that is always nonnegative, we obtain:

$$\sum_{t \in \mathcal{T}} \beta_t \left(\|\Delta_t^*\| - \frac{\kappa_t}{\beta_t} e(t) \right)^2 \leq \Delta_{\bar{t}} + \sum_{t \in \mathcal{T}} \frac{\kappa_t^2}{\beta_t} e^2(t). \quad (10)$$

Since $\beta_t > 0$ for all $t \in \mathcal{T}$, the expression in (10) is amenable to apply the Jensen's inequality [29, Th. 3.4] on the convex function $(\cdot)^2$ to lower bound the summation in the LHS. Before doing that, we normalize both sides by multiplying and dividing by $\bar{\beta} := \sum_{t \in \mathcal{T}} \beta_t$, and we define $\hat{\beta}_t := \beta_t / \bar{\beta}$, thus obtaining $\bar{\beta} (\sum_{t \in \mathcal{T}} \hat{\beta}_t (\|\Delta_t^*\| - \frac{\kappa_t}{\beta_t} e(t))^2) \geq \bar{\beta} (\sum_{t \in \mathcal{T}} \hat{\beta}_t (\|\Delta_t^*\| - \frac{\kappa_t}{\beta_t} e(t)))^2$. Then, after replacing this latter inequality into (10), with few algebraic manipulations we obtain

$$\sum_{t \in \mathcal{T}} \beta_t \left(\|\Delta_t^*\| - \frac{\kappa_t}{\beta_t} e(t) \right) \leq \sqrt{\sum_{t \in \mathcal{T}} \left(\beta_t \Delta_{\bar{t}} + \frac{\bar{\beta} \kappa_t^2}{\beta_t} e^2(t) \right)}.$$

We now bring the term $\sum_{t \in \mathcal{T}} \kappa_t e(t)$ in the RHS, and we note that the obtained inequality is still valid by premultiplying both sides by $1/T$, hence getting the average over the horizon of length T . Finally, we have that $(1/T) \sum_{t \in \mathcal{T}} \beta_t \|\Delta_t^*\| \geq (1/T) \underline{\beta} \sum_{t \in \mathcal{T}} \|\Delta_t^*\|$, with $\underline{\beta} := \min_{t \in \mathcal{T}} \beta_t$, which directly yields to the relation in (8) with probability $1 - \delta$. ■

We can hence upper bound the average value of the residual $\|\Delta_t^*\|$ over a certain horizon T with the sum of two terms: the first one depends on the initial distance from a global minimum for the unknown potential function θ , while the second one is mainly affected by the reconstruction error e .

Remarkably, the terms in the RHS of (8) can be made small in two ways: i) by choosing a small step-size ξ_t to make κ_t close to zero, or ii) by tuning the product $c_t \xi_t$ close to one, thus producing a large value of $\underline{\beta}$. This latter choice, however, requires an accurate trade-off in tuning the gain c_t and the step-size ξ_t , as larger values of $\underline{\beta}$ lead to larger values of the sub-optimal constant $\Delta_{\bar{t}}$. A possible choice to do not promote aggressive personalized actions establishes that the central coordinator may want to match the lower bound for c_t , while striving to counterbalance the choice for ξ_t to possibly obtain a faster convergence rate for Algorithm 1.

We now discuss the impact on the bound in (8) of the learning strategy \mathcal{L} endowing the central coordinator. By assuming that β_t is fixed over the outer iterations, i.e., $\beta_t = \beta$, we then know from Assumption 1 that there exists some \bar{t} such that $e(t) \leq e(\bar{t})$ for all $t \geq \bar{t}$. Therefore, it follows that the bound in (8) satisfies

$$\frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^*\| \leq O(1/\sqrt{T}) + O(1).$$

As T grows, it is evident that the term $O(1/\sqrt{T})$ vanishes, and $\frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^*\|$ is confined in a ball whose radius depends on the quantity of agents' feedback made available to perform the reconstruction in the first step of Algorithm 1 (specifically, on the learning strategy \mathcal{L}). In addition, we further stress that Assumption 1 is quite general and it holds true under mild conditions for LS and GP approaches to learning $Qx + q$:

- Note that $Qx + q$ can be always modelled as an affine function of the learning parameters η . Thus, setting up a LS approach to minimize the loss between the model parameters η and the agents' feedback turns out to be a quadratic program. Due to the large-scale properties of LS, the error term $e(t)$ behaves as a normal distribution, for which Assumption 1 holds true (see, e.g., [15, Lemma A.4]), and hence $\lim_{t \rightarrow \infty} e(t) = 0$.
- In a kernel-based methods, suppose $Qx + q$ is a sample path of a GP with zero mean and a certain kernel. Due to the large-scale property of such a regressor and under standard assumptions, also in this case Assumption 1 holds true (see [13]) and $\lim_{t \rightarrow \infty} e(t) = 0$.

We finally remark that, since $\sum_{t \in \mathcal{T}} e(t) = o(T)$, for the parametric/non-parametric procedures above $\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^*\| = 0$, which allows us to recover the result in the perfect reconstruction case described in §IV-A.




Fig. 1: Convergence behaviour of Algorithm 1 for different values of the step-size ξ , with perfect reconstruction of $G(\cdot)$.




Fig. 2: Average value of the residual $\|\Delta_t^*\|$ over the time horizon T , for different learning strategies of the coordinator.

V. NUMERICAL SIMULATIONS

We verify our findings through numerical instances of the hypomonotone GNEP in (1) with parameters summarized in Tab. I–II. Both examples involves $N = 20$ agents controlling a scalar variable $n_i = 1$ constrained to the local set $\mathcal{X}_i = [0, 1]$. Each matrix A_i is constructed so that every pair of consecutive agents are coupled, i.e., $x_i + x_{i+1} \leq b_i$, for all $i \in \mathcal{I}$, while the shared resource term, b_i is randomly sampled following a uniform distribution on $(0, 1)$. In addition, every Q_i and $C_{i,j}$ is randomly generated from a normal distribution, while every q_i follows a uniform distribution on the interval $(-1, 1)$.

Once given the personalized feedback functional at every outer iteration $t \in \mathbb{N}$, the v-GNE is computed by means of the projection-type method described in [20, Ch. 12], thus partially neglecting the multi-agent nature of the problem addressed. Since the algorithm in [20, Ch. 12] is tailored for monotone VIs, we stress that its convergence is enabled

TABLE I: Simulation parameters – Figure 1

Parameter	Description	Value
ℓ	Constant of weak convexity	10.09
c_t	Personalized functional gain	20.18
P	Weight for the extragradient algorithm	1
l_F	Horizontal scale – GP method	50
σ_F	Vertical scale – GP method	100




Fig. 3: Sub-optimality (left y-axis) and tracking error (right y-axis) of the sequence of v-GNE generated by Algorithm 1 w.r.t. $x^* \in \Theta$, for different learning strategies for the coordinator.

by the incentives in (4), as $Q + Q^\top$ is not positive semi-definite (i.e., $G(\cdot)$ is not monotone), while $(Q + c_t I) + (Q + c_t I)^\top$ is (i.e., $\tilde{G}_t(\cdot)$ is strongly monotone, for $c_t \geq 2\ell$ – see Proposition 2).

In Fig. 1 is illustrated the convergence boost effect, or error reduction, of the step-size ξ in case of perfect and inexact reconstruction, while in Fig. 2 we compare the behaviour of the metric $(1/T) \sum_{t \in \mathcal{T}} \|\Delta_t^*\|$ when the coordinator is endowed with an LS-based learning strategy, and a GP-based one (covariance matrices computed by using a radial basis function kernel). Since the problem in (1) makes available the explicit expression of $\theta(\cdot)$ in (2), we can show the evolution of the potential function evaluated in the sequence of v-GNE generated by Algorithm 1 compared to a global minimum (Fig. 3 – left y-axis), as well as the distance between any computed point w.r.t. some point in Θ (Fig. 3 – right y-axis).

TABLE II: Simulation parameters – Figure 2 and 3

Parameter	Description	Value
ℓ	Constant of weak convexity	1.22
μ	Mean of the agents' feedback noise $\varepsilon_{i,t}$	0
σ^2	Variance of the agents' feedback noise $\varepsilon_{i,t}$	25
c_t	Personalized functional gain	2.44
ξ_t	Personalized functional step-size	0.41
P	Weight for the extragradient algorithm	1
l_F	Horizontal scale – GP method	50
σ_F	Vertical scale – GP method	100

VI. CONCLUSION

The design of suitable personalized incentives is key to compute GNE in quadratic, nonmonotone GNEP characterized by bilateral symmetric interactions among agents. The benefit of adopting such functionals are twofold: i) they serve as regularization terms of the agents' cost functions, thus enabling for the practical computation of a v-GNE at each outer iteration of the proposed, two-layer procedure, and ii) they provide a mean to achieve faster convergence rate. The proposed algorithm converges to a GNE by exploiting the consistency bounds characterizing standard learning procedures for the coordinator, such as LS or GP.

Future research directions may include extending the proposed approach to time-varying GNEPs, also exploring how to avoid to know the constant of weak convexity ℓ of the potential function, which represents a fundamental, albeit possibly unknown, parameter to establish convergence.

REFERENCES

- [1] Q. Zhu and T. Başar, "A multi-resolution large population game framework for smart grid demand response management," in *International Conference on Network Games, Control and Optimization (NetGCoop 2011)*. IEEE, 2011, pp. 1–8.
- [2] C. Cenedese, F. Fabiani, M. Cucuzzella, J. M. Scherpen, M. Cao, and S. Grammatico, "Charging plug-in electric vehicles as a mixed-integer aggregative game," in *2019 IEEE 58th Conference on Decision and Control (CDC)*. IEEE, 2019, pp. 4904–4909.
- [3] J. Neel, R. M. Buehrer, B. H. Reed, and R. P. Gilles, "Game theoretic analysis of a network of cognitive radios," in *The 2002 45th Midwest Symposium on Circuits and Systems, 2002. MWSCAS-2002.*, vol. 3. IEEE, 2002, pp. III–III.
- [4] E. Altman, Y. Hayel, and H. Kameda, "Evolutionary dynamics and potential games in non-cooperative routing," in *2007 5th International Symposium on Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks and Workshops*. IEEE, 2007, pp. 1–5.
- [5] F. Zhang and N. E. Leonard, "Cooperative filters and control for cooperative exploration," *IEEE Transactions on Automatic Control*, vol. 55, no. 3, pp. 650–663, 2010.
- [6] F. Fabiani, D. Fenucci, and A. Caiti, "A distributed passivity approach to AUV teams control in cooperating potential games," *Ocean Engineering*, vol. 157, pp. 152–163, 2018.
- [7] Q. D. Lã, Y. H. Chew, and B.-H. Soong, *Potential Game Theory*. Springer Science & Business Media, 2016.
- [8] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *4OR*, vol. 5, no. 3, pp. 173–210, 2007.
- [9] T. Heikkinen, "A potential game approach to distributed power control and scheduling," *Computer Networks*, vol. 50, no. 13, pp. 2295–2311, 2006.
- [10] F. Fabiani and S. Grammatico, "Multi-vehicle automated driving as a generalized mixed-integer potential game," *IEEE Transactions on Intelligent Transportation Systems*, vol. 21, no. 3, pp. 1064–1073, 2019.
- [11] N. Li and J. R. Marden, "Designing games for distributed optimization," *IEEE Journal of Selected Topics in Signal Processing*, vol. 7, no. 2, pp. 230–242, 2013.
- [12] D. Gadjov and L. Pavel, "On the exact convergence to Nash equilibrium in hypomonotone regimes under full and partial-information," *arXiv preprint arXiv:2104.11096*, 2021.
- [13] A. Simonetto, E. Dall'Anese, J. Monteil, and A. Bernstein, "Personalized optimization with user's feedback," *Automatica*, vol. 131, p. 109767, 2021.
- [14] A. M. Ospina, A. Simonetto, and E. Dall'Anese, "Personalized demand response via shape-constrained online learning," in *Proceedings of the IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, 2020.
- [15] I. Notarnicola, A. Simonetto, F. Farina, and G. Notarstefano, "Distributed personalized gradient tracking with convex parametric models," *IEEE Transactions on Automatic Control*, 2022.

- [16] I. V. Konnov, M. Ali, and E. Mazurkevich, "Regularization of non-monotone variational inequalities," *Applied Mathematics and Optimization*, vol. 53, no. 3, pp. 311–330, 2006.
- [17] H. Yin, U. V. Shanbhag, and P. G. Mehta, "Nash equilibrium problems with scaled congestion costs and shared constraints," *IEEE Transactions on Automatic Control*, vol. 56, no. 7, pp. 1702–1708, 2011.
- [18] I. V. Konnov, "On penalty methods for non monotone equilibrium problems," *Journal of Global Optimization*, vol. 59, no. 1, pp. 131–138, 2014.
- [19] S. Lucidi, M. Passacantando, and F. Rinaldi, "Solving non-monotone equilibrium problems via a DIRECT-type approach," *Journal of Global Optimization*, pp. 1–27, 2022.
- [20] F. Facchinei and J. S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.
- [21] F. Facchinei, V. Piccialli, and M. Sciandrone, "Decomposition algorithms for generalized potential games," *Computational Optimization and Applications*, vol. 50, no. 2, pp. 237–262, 2011.
- [22] F. Fabiani and A. Caiti, "Nash equilibrium seeking in potential games with double-integrator agents," in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 548–553.
- [23] D. Davis and D. Drusvyatskiy, "Stochastic model-based minimization of weakly convex functions," *SIAM Journal on Optimization*, vol. 29, no. 1, pp. 207–239, 2019.
- [24] F. Salehisadaghiani and L. Pavel, "Distributed Nash equilibrium seeking: A gossip-based algorithm," *Automatica*, vol. 72, pp. 209–216, 2016.
- [25] M. Ye and G. Hu, "Distributed Nash equilibrium seeking by a consensus based approach," *IEEE Transactions on Automatic Control*, vol. 62, no. 9, pp. 4811–4818, 2017.
- [26] C. Jin, R. Ge, P. Netrapalli, S. M. Kakade, and M. I. Jordan, "How to escape saddle points efficiently," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, 2017, pp. 1724–1732.
- [27] D. P. Bertsekas, "Nonlinear programming," *Journal of the Operational Research Society*, vol. 48, no. 3, pp. 334–334, 1997.
- [28] R. Dixit, A. S. Bedi, R. Tripathi, and K. Rajawat, "Online learning with inexact proximal online gradient descent algorithms," *IEEE Transactions on Signal Processing*, vol. 67, no. 5, pp. 1338–1352, 2019.
- [29] R. T. Rockafellar, *Convex analysis*. Princeton university press, 1970, vol. 36.