Personalized incentives as feedback design in generalized Nash equilibrium problems

Filippo Fabiani, Andrea Simonetto, and Paul J. Goulart

Abstract

We investigate both stationary and time-varying, nonmonotone generalized Nash equilibrium problems that exhibit symmetric interactions among the agents, which are known to be potential. As may happen in practical cases, however, we envision a scenario in which the formal expression of the underlying potential function is not available, and we design a semi-decentralized Nash equilibrium seeking algorithm. In the proposed two-layer scheme, a coordinator iteratively integrates possibly noisy and sporadic agents' feedback to learn the pseudo-gradients of the agents, and then design personalized incentives for them. On their side, the agents receive those personalized incentives, compute a solution to an extended game, and then return feedback measurements to the coordinator. In the stationary setting, our algorithm returns a Nash equilibrium in case the coordinator is endowed with standard learning policies, while it returns a Nash equilibrium up to a constant, yet adjustable, error in the time-varying case. As a motivating application, we consider the ride-hailing service provided by several competing companies with mobility as a service orchestration, necessary to both handle competition among firms and avoid traffic congestion.

Index Terms

At the time of the work, F. Fabiani was with the Department of Engineering Science, University of Oxford, OX1 3PJ, United Kingdom. Currently, he is with the IMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100 Lucca, Italy (`filippo.fabiani@imtlucca.it`). A. Simonetto is with the UMA, ENSTA Paris, Institut Polytechnique de Paris, 91120 Palaiseau, France (`andrea.simonetto@ensta-paris.fr`). P. J. Goulart is with the Department of Engineering Science, University of Oxford, OX1 3PJ, United Kingdom (`paul.goulart@eng.ox.ac.uk`).

Game theory, Time-varying optimization, Machine learning.

I. Introduction

Noncooperative game theory represents a contemporary and pervasive paradigm for the modelling and optimization of modern multi-agent systems, where agents are typically modelled as rational decision-makers that interact and selfishly compete for shared resources in a *stationary* environment. Here, the (generalized) Nash equilibrium solution concept [1] denotes a desired outcome of the game, which is typically self-learned by the agents through iterative procedures alternating distributed computation and communication steps [2]–[5].

Real-world scenarios, however, are rarely stationary. This fact, along with recent developments in machine learning and online optimization [6]–[12], has fostered the implementation of online multi-agent learning procedures, thus contributing in growing the interest for games where the population of agents ambitiously aim at tracking possibly *time-varying* Nash equilibria online. In this paper we focus on both static and time-varying *nonmonotone* generalized Nash equilibrium problems (GNEPs) that exhibit symmetric interactions among the agents, for which we design a semi-decentralized algorithm with convergence guarantees. The presence of symmetric interactions, indeed, brings numerous advantages to the GNEP that enjoys it. Among them, the underlying GNEP is known to be potential with associated potential function [13], [14], which implicitly entails the existence of a generalized Nash equilibrium (GNE). Nonetheless, such a potential function frequently enables for the design of equilibrium seeking algorithms with convergence guarantees (especially in nonconvex setting [15]–[18]). However, unless one has a deep knowledge on the main quantities characterizing the symmetric interactions of the GNEP at hand, finding the formal expression of the potential function is known to be a hard task [14, Ch. 2] [19]. Moreover, in many in-network operations that require some degree of coordination it is highly desirable that the parameters of the agents' cost function, which reflect local sensitive data, stay private. In particular, our work is widely motivated by the ride-hailing application with mobility as a service (MaaS) orchestration formally developed and described in §VI.

In this framework, the proposed two-layer algorithm then reads as follows: in the outer loop,

we endow a coordinator with an online learning procedure, aiming at iteratively integrating possibly noisy and sporadic agents' feedback to learn some of their private information, i.e., the pseudo-gradient mappings associated to the agents' cost functions. The reconstructed information is thereby exploited by the coordinator to design *personalized incentives* [11], [20]–[22] for the agents, which in turn compute a solution to an extended game, and then return (dis)satisfaction feedback measurements to the coordinator.

A. Related work

A recent research direction established the convergence of online distributed mirror descent-type algorithms in strictly monotone GNEPs [23], aggregative games with estimated information [24], or in price-based congestion control methods for generic noncooperative games [25]. Conversely, [26] focused on the prediction of the long-term outcome of a monotone Nash equilibrium problem (NEP), also extended to the case of delays in the communication protocol [27]. The convergence of no-regret learning policies with exponential weights in potential games within a (semi-)bandit framework was explored in [28], while [29] introduced an algorithm with sublinear Nash equilibrium regret under bandit feedback for time-varying matrix games, and [30] showed that, in case of a slowly-varying monotone NEP, the dynamic regret minimization allows the agents to track the sequence of equilibria. Unlike the problem setting considered in this paper, [31] dealt with the quadratic class of stationary nonmonotone GNEPs only where a feedback to the coordinator was provided at every iteration.

We remark that our primary goal is to devise an algorithm to solve the GNE seeking problem of an assigned noncooperative game. We do that by endowing a central coordinator with a learning procedure to reconstruct individual information of the agents, which is hence exploited to design personalized incentives, thus enabling for the equilibrium seeking. This is however misaligned with the overall goal of typical agents' utility and mechanism design, which mainly consists in designing games and payoff functions in which certain emergent behaviours coincide with a desirable outcome [32]–[34].

Finally, our semi-decentralized scheme may also be interpreted as a Stackelberg game in

which the leader does not control any decision variable, albeit aims at minimizing the unknown potential function on the basis of optimistic conjectures on the followers' strategies [35], [36].

B. Summary of contributions and paper organization

In contrast to the aforementioned literature, we consider nonmonotone GNEPs admitting symmetric interactions, for which we propose an online learning procedure based on *personalized incentives* [11], [20], [21]. The main contributions made can be summarized as follows:

- We design a semi-decentralized scheme that allows the agents to compute (or track in a neighbourhood) a GNE of a *nonmonotone* GNEP that admits an *unknown* potential function (§II–III). In particular:
  - In the static case, we show that the proposed algorithm converges to a GNE by exploiting the asymptotic consistency bounds characterizing typical learning procedures for the coordinator, such as least squares (LS) or Gaussian process (GP) (§IV);
  - In the time-varying setting, we show that the fixed point residual, our metric for assessing convergence, asymptotically behaves as $O(1)$, i.e., the proposed semi-decentralized scheme allows the agents to track a GNE in a neighbourhood of adjustable size (§V)
- Inspired by real data available online [37], we develop a mathematical model capturing intrinsic features of the competition arising among companies of different size participating in the ride-hailing market with MaaS orchestration. The developed model is then used as a case study to corroborate our theoretical results (§VI);

We also show that the design of the personalized incentives is key for the convergence of the algorithm, as they bring a twofold benefit: i) enabling the agents for the computation of a variational generalized Nash equilibrium (v-GNE) in the inner loop by acting as a convexification terms for their cost functions; and ii) boosting the convergence and/or lessening the tracking error through a fine tuning of few parameters. Since our convergence results strongly rely on the knowledge of a private information held by each agent, specifically, the Lipschitz constant of their cost functions, in Appendix E we provide a possible learning-based solution to address this

privacy issue, also accompanied by a dedicated analysis. The proofs of theoretical results are all deferred to Appendix A–D.

Notation

$\mathbb{N}$, $\mathbb{R}$ and $\mathbb{R}_{\geq 0}$ denote the set of natural, real and nonnegative real numbers, respectively. $\mathbb{S}^n$ is the space of $n \times n$ symmetric matrices. For vectors $v_1, \ldots, v_N \in \mathbb{R}^n$ and $\mathcal{I} = \{1, \ldots, N\}$, we denote $\boldsymbol{v} := (v_1^\top, \ldots, v_N^\top)^\top = \mathrm{col}((v_i)_{i \in \mathcal{I}})$ and $\boldsymbol{v}_{-i} := \mathrm{col}((v_j)_{j \in \mathcal{I} \setminus \{i\}})$. With a slight abuse of notation, we also use $\boldsymbol{v} = (v_i, \boldsymbol{v}_{-i})$. $\mathcal{C}^1$ is the class of continuously differentiable functions. The mapping $F : \mathbb{R}^n \to \mathbb{R}^n$ 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 $x, y \in \mathcal{X}$. If F is differentiable, $\mathrm{J}_F : \mathbb{R}^n \to \mathbb{R}^{n \times n}$ denotes its Jacobian matrix. Throughout the paper, variables with t as subscript do not explicitly depend on time, as opposed when t is an argument.

II. PROBLEM FORMULATION

We consider a noncooperative game $\Gamma := (\mathcal{I}, (\mathcal{X}_i)_{i \in \mathcal{I}}, (g_i)_{i \in \mathcal{I}})$, with N agents, indexed by the set $\mathcal{I} := \{1, \ldots, N\}$. Each agent $i \in \mathcal{I}$ controls a local variable $x_i \in \mathcal{X}_i \subseteq \mathbb{R}^{n_i}$ and, at every discrete time instant $t \in \mathbb{N}$, aims at solving the following time-varying optimization problem:

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in \mathcal{X}_i} & g_i(x_i, \boldsymbol{x}_{-i}; t) \\ \text{s.t.} & h_i(x_i) + \sum_{j \in \mathcal{I} \setminus \{i\}} h_j(x_j) \leq 0, \end{cases} \tag{1}$$

for some $g_i : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}$, $n = \sum_{i \in \mathcal{I}} n_i$, which denotes the private individual cost, whose value at time $t \in \mathbb{N}$ can be interpreted as the (dis)satisfaction of the i-th agent associated to the collective strategy $(x_i, \boldsymbol{x}_{-i})$. The collection of optimization problems in (1) amounts to a GNEP, where every $h_i : \mathbb{R}^{n_i} \to \mathbb{R}^m$ is a map stacking m coupling, yet locally separable, constraints among the agents. First, we define sets $\mathcal{X} := \prod_{i \in \mathcal{I}} \mathcal{X}_i$ and $\mathcal{X}_i(\boldsymbol{x}_{-i}) := \{x_i \in \mathcal{X}_i \mid h(x_i, \boldsymbol{x}_{-i}) \leq 0\}$, with $h(x_i, \boldsymbol{x}_{-i}) := h_i(x_i) + \sum_{j \in \mathcal{I} \setminus \{i\}} h_j(x_j)$, and then we introduce some standard assumptions:

Standing Assumption 1. *For each $i \in \mathcal{I}$, and for all $t \in \mathbb{N}$,*

    i) *The mapping $x_i \mapsto g_i(x_i, \boldsymbol{x}_{-i}; t)$ is of class $\mathcal{C}^1$ and has a ℓ_i-Lipschitz continuous gradient;*

    ii) *$\mathcal{X}_i$ is a nonempty, compact and convex set, $h_i : \mathbb{R}^{n_i} \to \mathbb{R}^m$ is a convex and of class $\mathcal{C}^1$ function.* $\qquad\square$

The feasible set of the time-varying GNEP Γ thus coincides with $\Omega := \{\boldsymbol{x} \in \mathcal{X} \mid h(\boldsymbol{x}) \leq 0\}$ [1, §3.2]. In the proposed time-varying context, we are then interested in designing an equilibrium seeking algorithm for the game Γ, according to the following popular definition of GNE:

Definition 1. (Generalized Nash equilibrium [1]) *For all $t \in \mathbb{N}$, $\boldsymbol{x}^\star(t) \in \Omega$ is a* GNE *of the game Γ if, for all $i \in \mathcal{I}$,*

$$g_i(x_i^\star(t), \boldsymbol{x}_{-i}^\star(t); t) \leq \inf_{y_i \in \mathcal{X}_i(\boldsymbol{x}_{-i}^\star(t))} g_i(y_i, \boldsymbol{x}_{-i}^\star(t); t). \tag{2}$$

$\qquad\square$

A collective vector of strategies $\boldsymbol{x}^\star(t) \in \Omega$ is therefore an equilibrium at time $t \in \mathbb{N}$ if no agent can decrease their objective function by changing unilaterally $\boldsymbol{x}_{-i}^\star(t)$ to any other feasible point. Note that Definition 1 assumes a local connotation if one simply focuses on a certain neighbourhood of $\boldsymbol{x}^\star(t)$ for which condition (2) holds true. Throughout the paper we will make use of the following assumption on the pseudo-gradient (or game) mapping $G : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}^n$, which is formally defined as $G(\boldsymbol{x}; t) := \mathrm{col}((\nabla_{x_i} g_i(x_i, \boldsymbol{x}_{-i}; t))_{i \in \mathcal{I}})$:

Standing Assumption 2. *For every $\boldsymbol{x} \in \Omega$ and $t \in \mathbb{N}$, $\mathrm{J}_G(\boldsymbol{x}; t) \in \mathbb{S}^n$.* $\qquad\square$

Roughly speaking, Standing Assumption 2 establishes that each pair of agents $(i, j) \in \mathcal{I}^2$ influences each other in an equivalent way. For the mapping G, this entails the existence of a differentiable, yet possibly *unknown*, function $\theta : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}$ such that $G(\boldsymbol{x}; t) = \nabla\theta(\boldsymbol{x}; t)$, for all $\boldsymbol{x} \in \Omega$ and $t \in \mathbb{N}$ [38, Th. 1.3.1], which coincides with an exact potential function [13], [14], [39] for Γ and can be characterized as stated next.

Lemma 1. *For all $t \in \mathbb{N}$, $\boldsymbol{x} \mapsto \nabla\theta(\boldsymbol{x}; t)$ is ℓ-Lipschitz continuous, while $\boldsymbol{x} \mapsto \theta(\boldsymbol{x}; t)$ is ℓ-weakly convex, i.e., $\boldsymbol{x} \mapsto \theta(\boldsymbol{x}; t) + \frac{\ell}{2}\|x\|^2$ is convex, with $\ell := \sum_{i \in \mathcal{I}} \ell_i$.* $\qquad\square$

Note that θ is a smooth function that in principle may be nonconvex. Let $\Theta(t) \coloneqq \mathrm{argmin}_{\boldsymbol{y} \in \Omega}\, \theta(\boldsymbol{y}; t)$ be the set of its (local and global) constrained minimizers, assumed to be nonempty, and $\Theta^{\mathrm{s}}(t)$ be the set of its constrained stationary points, with $\Theta(t) \subseteq \Theta^{\mathrm{s}}(t)$, for all $t \in \mathbb{N}$. We stress that the nonemptiness of $\Theta(t)$ guarantees the existence of an (at least local) GNE for Γ, since any $\boldsymbol{x}^{\star}(t) \in \Theta(t)$ satisfies the relation in (2). It is well-known in potential game theory, indeed, that any minimum point of the exact potential function θ coincides with a GNE, whose global or local nature depends on the point computed in $\Theta(t)$ – see, for instance, [16], [17], [40].

A. *Main challenges and technical considerations*

In the considered framework, we identify three main critical issues that rule out the possibility to compute a GNE for the GNEP Γ in (1) through standard arguments, thus fully motivating the design of a tailored learning procedure.

First, the time-varying nature of the optimization problems in (1) calls for an answer to the thorny question on whether there exist online learning policies that allow agents to track a Nash equilibrium over time (or to converge to one if the stage games stabilize). Even in the case of a potential game with known potential function, this is a challenging problem [28].

In addition, despite the symmetry of interactions among agents, we note that for all $t \in \mathbb{N}$ the mapping $\boldsymbol{x} \mapsto G(\boldsymbol{x}; t)$ may not be monotone, a key technical requirement for the most common solution algorithms for GNEPs available in the literature, which compute a GNE by relying on the (at least) monotonicity of the pseudo-gradient mapping [2]–[5].

Finally, we stress that Standing Assumption 2, albeit quite mild and practically satisfied in several real-world scenarios [41]–[44], is key to claim that the underlying GNEP is potential. However, unless one has a deep knowledge of the GNEP at hand, finding the formal expression of the potential function is known to be a hard task [14, Ch. 2]. Thus, we assume not to have an expression for $\theta(\boldsymbol{x}; t)$ that can be exploited directly for the equilibrium seeking algorithm design.

To address these crucial issues, we design *personalized feedback functionals* $u_i : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}$
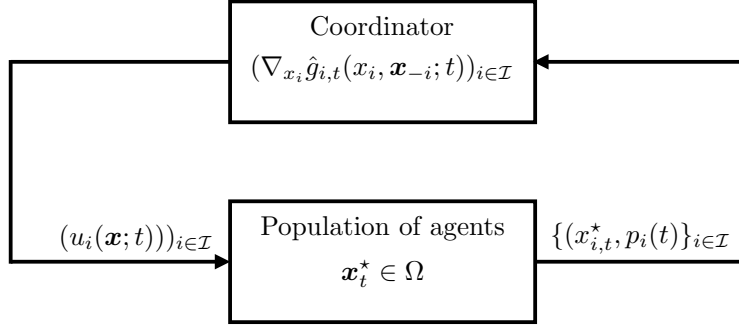
Fig. 1: Personalized incentives as feedback design to steer the population to a point guaranteeing the "minimum" (dis)satisfaction, according to the unknown function θ.

Algorithm 1: Two-layer semi-decentralized scheme

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

- (S0) Learn pseudo-gradients $(\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}^\star_{t-1}; t-1))_{i\in\mathcal{I}}$

- (S1) Design personalized incentives $(u_i(\boldsymbol{x}; t))_{i\in\mathcal{I}}$

∘ (S2) Compute a GNE of the extended game $\overline{\Gamma}$, $\boldsymbol{x}^\star_t \in \Omega$

- (S3) Retrieve noisy agents' feedback $\{(x^\star_{i,t}, p_i(t))\}_{i\in\mathcal{I}}$

---

in the spirit of [11], [20], [21], [31], which are then used as "control actions" in the semi-decentralized scheme depicted in Fig. 1. Specifically, our goal is to steer the noncooperative agents to track minimizers of the unknown, time-varying function θ, i.e., a GNE of the game Γ, according to Definition 1. Any $\boldsymbol{x}^\star(t) \in \Theta(t)$ can indeed be interpreted as a collective strategy that minimizes the (dis)satisfaction of the of agents, measured by the function θ.

III. LEARNING ALGORITHM WITH
PERSONALIZED INCENTIVES

A. *The two-layer algorithm*

The proposed approach is summarized in Algorithm 1, where 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 $\mathcal{I}$. Thus, in the outer loop a central entity aims at learning online the unknown, time-varying function θ (or its gradient mapping, $\nabla\theta$) by leveraging possibly noisy and sporadic agents' feedback on the private functions g_i's (S0). On the basis of the estimated $\hat{g}_{i,t}$, at item (S1) the coordinator designs personalized incentive functionals u_i, which are successively communicated to the noncooperative agents taking part to the game. These latter then face with an extended version of the GNEP Γ in (1) at item (S2), i.e., $\overline{\Gamma} := (\mathcal{I}, (\mathcal{X}_i)_{i\in\mathcal{I}}, (f_i)_{i\in\mathcal{I}})$, with $g_i(x_i, \boldsymbol{x}_{-i}; t) + u_i(x_i, \boldsymbol{x}_{-i}; t) =: f_i(x_i, \boldsymbol{x}_{-i}; t)$ in place of $g_i(x_i, \boldsymbol{x}_{-i}; t)$. Under a suitable choice of the personalized incentives, we will show that they act as regularization terms, as well as they trade-off convergence and robustness to the inexact knowledge of function θ and its gradient. Specifically, such incentives enable for the practical computation of an equilibrium of the extended game $\overline{\Gamma}$ at item (S2) through available solution algorithms for GNEPs [2], [4], [5], which typically require distributed computation and inner communication rounds among the noncooperative agents in $\mathcal{I}$.

Note that standard procedures in literature typically returns a v-GNE [1], [45], which coincides to any solution to the GNEP $\overline{\Gamma}$ that is also a solution to the associated variational inequality (VI), i.e., any vector $\boldsymbol{x}_t^\star \in \Omega$ such that, for all $t \in \mathbb{N}$,

$$(\boldsymbol{y} - \boldsymbol{x}_t^\star)^\top F(\boldsymbol{x}_t^\star; t) \geq 0, \;\; \text{for all } \boldsymbol{y} \in \Omega, \, t \in \mathbb{N}, \tag{3}$$

where the mapping $F : \mathbb{R}^n \times \mathbb{N} \to \mathbb{R}^n$ is formally defined as $F(\boldsymbol{x}; t) := \text{col}((\nabla_{x_i} f_i(x_i, \boldsymbol{x}_{-i}; t))_{i\in\mathcal{I}}) = G(\boldsymbol{x}; t) + U(\boldsymbol{x}; t)$, and $U(\boldsymbol{x}; t) := \text{col}((\nabla_{x_i} u_i(x_i, \boldsymbol{x}_{-i}; t))_{i\in\mathcal{I}})$. For these reasons, in referring to the computational step (S2), we tacitly assume that the agents compute a v-GNE of the extended game $\overline{\Gamma}$.

Finally, at item (S3) the agents communicate feedback measures and their equilibrium strategies, $\{(x_{i,t}^\star, p_i(t))\}_{i\in\mathcal{I}}$, with $p_i(t) \coloneqq g_i(\boldsymbol{x}_t^\star; t) + \varepsilon_{i,t}$, for some random variable ε_i, to the central entity, thus indicating to what extent the current equilibrium $\boldsymbol{x}_i^\star$ (dis)satisfies the entire population of agents.

Besides the time-varying nature of the original game Γ, introducing the personalized incentives $u_i(x_i, \boldsymbol{x}_{-i}; t)$ further modifies Γ, thus forcing the agents to deal with an extended game $\overline{\Gamma}$ at each step (S2) of Algorithm 1. The relation between the equilibria of Γ and $\overline{\Gamma}$ is hence not straightforward. What we prove later in the paper is that our approach allows the agents to compute (or track in a neighbourhood) a GNE of the original GNEP Γ, particularly that GNE coinciding with a minimum of the unknown potential function $\theta(\boldsymbol{x}; t)$. We can therefore claim that, as $t \to \infty$, Γ and $\overline{\Gamma}$ share at least one equilibrium, precisely that one coinciding with a minimum point of $\theta(\boldsymbol{x}; t)$. In different words, the learning procedure $\mathscr{L}$ characterizing the central coordinator, along with the parametric personalized incentives chosen, allows us to locally approximate the original nonmonotone GNEP Γ around a certain GNE trajectory.

B. Personalized incentives design

In view of Standing Assumption 2 we have $G(\boldsymbol{x}; t) = \nabla\theta(\boldsymbol{x}; t)$, and hence a natural approach to design the personalized incentives u_i seems to iteratively learn and point a descent direction for the unknown function θ, thus implicitly requiring one to estimate the pseudo-gradients $(\nabla_{x_i} g_i(x_i, \boldsymbol{x}_{-i}; t))_{i\in\mathcal{I}}$, at every $t \in \mathbb{N}$. Along the line of [20], [31], we assume the central coordinator being endowed with a learning procedure $\mathscr{L}$ such that, at every outer iteration $t \in \mathbb{N}$ (Algorithm 1, item (S0)), it integrates the most recent agents' feedback $\{p_i(t-1)\}_{i\in\mathcal{I}}$ to return an estimate of the pseudo-gradients, $(\nabla_{x_i} \hat{g}_{i,t-1}(\boldsymbol{x}_{t-1}^\star; t-1))_{i\in\mathcal{I}}$. A possible personalized incentive functional can hence be designed as

$$u_i(\boldsymbol{x}; t) = \tfrac{1}{2} c(t)\|x_i - x_i^+(t)\|^2, \text{ for all } i \in \mathcal{I}, \tag{4}$$

where $x_i^+(t) \coloneqq x_{i,t-1}^\star + \xi(t)\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}_{t-1}^\star; t-1)$, for some parameters $c(t)$, $\xi(t) \geq 0$, for all $t \in \mathbb{N}$. Unlike what one might expect, each $x_i^+(t)$ requires a positive sign for the gradient step $\xi(t)\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}_{t-1}^\star; t-1)$. However, note that this fact is not uncommon – see, e.g., the recent

Heavy Anchor method [5, Eq. (7)]. Moreover, we will also discuss later on how such a choice enables us to boost the convergence of Algorithm 1 or lessen the tracking error through a fine tuning of $\xi(t)$.

Thus, once the parametric form in (4) is fixed, we design suitable bounds for $c(t)$ and $\xi(t)$ in such a way that the sequence of GNE, $(x_t^\star)_{t\in\mathbb{N}}$, monotonically decreases $\theta(x;t)$ and converges to some point in $\Theta(t)$. As stressed in the previous section, the gain $c(t)$ is crucial to enable for the computation of a v-GNE at item (S2) in Algorithm 1, as formalized next:

Proposition 1. *Let $c(t) \geq 2\ell$ for all $t \in \mathbb{N}$. Then, with the personalized incentives in (4), the mapping $x \mapsto F(x;t)$ is ℓ-strongly monotone, for all $t \in \mathbb{N}$.*　□

Thus, at every $t \in \mathbb{N}$, in (S2) the population of agents computes the (unique, see [38, Th. 2.3.3]) v-GNE associated to the extended version of the GNEP in (1), $\overline{\Gamma}$. A key quantity for the convergence analysis of the proposed algorithm, both in the stationary and time-varying case, will be the fixed point residual $\Delta_t^\star := x_t^\star - x_{t-1}^\star$, whose norm "measures" the distance to the points in Θ^{s} when the function $\theta(x)$ is fixed in time.

Lemma 2. *Let $(x_t^\star)_{t\in\mathbb{N}}$ be the sequence of v-GNE generated by Algorithm 1 with $1 - c(t)\xi(t) > 0$, assume perfect reconstruction of the mapping $x \mapsto G(x)$, and that $\|\Delta_t^\star\| = 0$ for some $x_t^\star \in \Omega$. Then, $x_t^\star \in \Theta^{\mathrm{s}}$.*　□

As briefly seen, choosing and tuning appropriately the parameters $c(t)$ and $\xi(t)$ is important to drive the sequence of v-GNE $(x_t^\star)_{t\in\mathbb{N}}$ along a descent direction for the unknown θ, and ensuring $\|\Delta_t^\star\| \to 0$. Typically, we need $c(t) \geq 2\ell(t), 1 - c(t)\xi(t) > 0$, and we will see in §VI how different choices could boost convergence and performance. In case of imperfect reconstruction of $x \mapsto G(x)$, or in the time-varying setting, we also adopt the average value of $\|\Delta_t^\star\|$ over a certain horizon of length $T \in \mathbb{N}$, i.e., $\frac{1}{T}\sum_{t\in\mathcal{T}} \|\Delta_t^\star\|$, $\mathcal{T} := \{1, \ldots, T\}$, as a metric for the convergence of the sequence $(x_t^\star)_{t\in\mathbb{N}}$ generated by Algorithm 1 to the stationary point set.

We remark here that, on the one hand, finding the stationary points is the general goal in nonconvex setting [46], and on the other hand, since Algorithm 1 generates monononically

decreasing values for θ, the application of simple perturbation techniques (e.g., [47]) can ensure that the stationary points to which we converge are in practice constrained local minima for θ, namely points belonging to $\Theta \subseteq \Theta^s$, and therefore GNE of the GNEP Γ in (1), according to Definition 1. A technique that works well in practice is to use $G(\boldsymbol{x}) = \nabla\theta(\boldsymbol{x})$ to verify whether $\boldsymbol{x}$ is a minimum point of θ by small feasible perturbations, and if not, introduce that perturbation into (4).

Remark 1. *The bounds on the parameters $c(t)$ and $\xi(t)$ provided in the paper assume the knowledge of the constant of weak convexity of θ, $\ell = \sum_{i\in\mathcal{I}}\ell_i$. As long as the coordinator is endowed with a learning policy, however, one may include this additional condition in the learning process, thus obtaining bounds that depend on $\hat{\ell} \geq 0$, the estimate of ℓ. We discuss and elaborate more around this point in Appendix E.* $\qquad\square$

IV. THE STATIONARY CASE

We start by discussing the case in which each g_i in (1) is fixed in time, thus implying that $\theta(\boldsymbol{x};t) = \theta(\boldsymbol{x})$. First, we analyze the case of perfect reconstruction of the pseudo-gradient mappings $(\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}^\star_{t-1}))_{i\in\mathcal{I}}$ (§IV-A), and then we investigate their inexact estimate (§IV-B). Here, our result will be of the form $(1/T)\sum_{t\in\mathcal{T}}\|\Delta^\star_t\| = O(1)$ in case the reconstruction error is non-vanishing. Otherwise, $(1/T)\sum_{t\in\mathcal{T}}\|\Delta^\star_t\| = 0$ (§IV-C), thus recovering the results shown in §IV-A.

A. Online perfect reconstruction of the pseudo-gradients

In case the learning procedure $\mathscr{L}$ enables for $\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}^\star_{t-1}) = \nabla_{x_i}g_i(\boldsymbol{x}^\star_{t-1})$, $i \in \mathcal{I}$, by adopting the personalized incentives in (4) at every outer iteration $t \in \mathbb{N}$, we have the following result:

Lemma 3. *Let $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. Then, with the personalized incentives in (4), the vector $\Delta^\star_t$ is a descent direction for $\theta(\boldsymbol{x}^\star_{t-1})$, i.e., $\Delta^{\star\top}_t\nabla\theta(\boldsymbol{x}^\star_{t-1}) < 0$.* $\quad\square$

Then, if $c(t)$ (resp., $\xi(t)$) is large (small) enough, at every iteration $t \in \mathbb{N}$ of Algorithm 1, the personalized functionals in (4) allow to point a descent direction for the unknown (dis)satisfaction function θ. Next, we establish the convergence of the sequence of v-GNE generated by Algorithm 1.

Proposition 2. *Let $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. With the personalized incentives in (4), the sequence of* v-GNE $(\boldsymbol{x}_t^\star)_{t \in \mathbb{N}}$, *generated by Algorithm 1, converges to some point in* Θ^s. $\qquad\qquad\square$

By introducing $\alpha(t) := 1 - c(t)\xi(t)$, from the first step of the proof of Proposition 2 we have that $\theta(\boldsymbol{x}_t^\star) \leq \theta(\boldsymbol{x}_{t-1}^\star) - \ell(2 - \alpha(t))/2\alpha(t)\|\Delta_t^\star\|^2$, which points out that a fine tuning of the term $c(t)\xi(t)$ allows us to boost the convergence of Algorithm 1 to some point in Θ^s (also observed on a numerical example in [31, §V]). This essentially explains the choice for a positive sign in the gradient step of (4). However, due to the presence of noise in the agents' feedback $\{p_i(t-1)\}_{i \in \mathcal{I}}$, it seems unlikely that the online algorithm $\mathcal{L}$ is able to return a perfect reconstruction of $(\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}_{t-1}^\star))_{i \in \mathcal{I}}$, at least at the beginning of the procedure in Algorithm 1.

B. Inexact estimate of the pseudo-gradients

At every outer iteration $t \in \mathbb{N}$, we assume the coordinator has available $K \in \mathbb{N}$ agents' feedback $\{p(k)\}_{k \in \mathcal{K}}$, $\mathcal{K} := \{1, \ldots, K\}$, and $p(k) := \mathrm{col}((p_i(k))_{i \in \mathcal{I}}) \in \mathbb{R}^N$, to estimate the gradients $(\nabla_{x_i}g_i(x_i, \boldsymbol{x}_{-i}))_{i \in \mathcal{I}}$ (and hence the mapping $\boldsymbol{x} \mapsto G(\boldsymbol{x})$). The value of K reflects situations in which the coordinator gathered information before starting the procedure ($K \geq t$), or it obtains sporadic feedback from the agents ($K < t$). Without restriction, we make the following, standard assumption on the reconstructed mapping $\boldsymbol{x} \mapsto \hat{G}_t(\boldsymbol{x})$ directly, rather than on each single gradient [10], [12], [20].

Assumption 1. *For all $t \in \mathbb{N}$ and $\boldsymbol{x} \in \mathcal{X}$, $\hat{G}_t(\boldsymbol{x}) := G(\boldsymbol{x}) + \epsilon_t$, and, for any $\delta_1 \in (0, 1]$, there exists $\bar{t} < \infty$ and available K_1 agents' feedback such that*

$$\mathbb{P}\{\|\epsilon_t\| \leq \mathrm{e}(K) \mid \forall t \geq \bar{t}, \forall K \geq K_1\} \geq 1 - \delta_1,$$

for some nonincreasing function $\mathrm{e} : \mathbb{N} \to \mathbb{R}_{\geq 0}$ *such that* $\mathrm{e}(K) < \infty$, *for all $K \in \mathbb{N}$.* $\qquad\square$

With Assumption 1, the reconstruction error on $\boldsymbol{x} \mapsto G(\boldsymbol{x})$ made by $\mathscr{L}$ is bounded with high probability $1 - \delta_1$ by some function of the available K agents' feedback.

Remark 2. *Assumption 1 is reasonable for various learning strategies. Consider for instance a scalar* LS *estimator: take $e(K) = e$, i.e., a constant, then $\delta_1 \propto \mathrm{e}^{-\bar{t}}$, while for a function $e(K) \propto 1/t^{\vartheta}, \vartheta \in [0, 1/2)$, then $\delta_1 \propto \Gamma(1/\varsigma, \bar{t}^{\varsigma})$, with $\varsigma = 1 - 2\vartheta \in (0, 1]$, and $\Gamma(\cdot, \cdot)$ being the upper incomplete Gamma function. The latter is finite for our choice of parameters, and goes to zero as $\bar{t} \to \infty$. In both cases Assumption 1 is verified. See also [21, Lemma A.4] for the derivations and further extensions.* □

After defining quantities $\kappa(t) \coloneqq (1 - \alpha(t))/2\alpha(t)$ and $\beta(t) \coloneqq \ell(2 - \alpha(t))/2\alpha(t)$, we have the following result.

Lemma 4. *Let Assumption 1 hold true for some fixed $\delta_1 \in (0, 1]$, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$ for all $t \in \mathbb{N}$. Then, with the personalized incentives in (4), for all $t \geq \bar{t}$ we have*

$$\Delta_t^{\star^{\top}} \nabla \theta(\boldsymbol{x}_{t-1}^{\star}) \leq \frac{\alpha(t)\kappa^2(t)}{\ell} \mathrm{e}^2(K - 1) - \left(\sqrt{\frac{\ell}{\alpha(t)}} \|\Delta_t^{\star}\| - \kappa(t)\sqrt{\frac{\alpha(t)}{\ell}} \, \mathrm{e}(K - 1) \right)^2, \tag{5}$$

with probability $1 - \delta_1$, for some $K \geq K_1$. □

In case of inexact estimate of the pseudo-gradients, the vector $\Delta_t^{\star}$ is not guaranteed to be a descent direction for the unknown function θ. In fact, the term $\mathrm{e}^2(K - 1)$ rules out the possibility that the LHS in (5) is strictly negative, albeit it can be made arbitrarily small through $\kappa(t)$ by an appropriate choice of the step-size $\xi(t)$. As in §IV-A, the following bound characterizes the sequence of v-GNE generated by Algorithm 1.

Theorem 1. *Let Assumption 1 hold true for some fixed $\delta_1 \in (0, 1]$, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. Moreover, let some $T \in \mathbb{N}$ be fixed, $\mathcal{T} \coloneqq \{\bar{t} + 1, \ldots, T + \bar{t}\}$ and, for any global minimizer $\boldsymbol{x}^{\star} \in \Theta$, $\Delta_{\bar{t}} \coloneqq \theta(\boldsymbol{x}_{\bar{t}}^{\star}) - \theta(\boldsymbol{x}^{\star})$. Then, with the personalized incentives in (4), the sequence of v-GNE $(\boldsymbol{x}_t^{\star})_{t \in \mathcal{T}}$, generated by Algorithm 1, satisfies the following relation with probability $1 - \delta_1$*

$$\frac{1}{T}\sum_{t \in \mathcal{T}} \|\Delta_t^{\star}\| \leq \frac{1}{T\underline{\beta}}\sqrt{\sum_{t \in \mathcal{T}}\left(\beta(t)\Delta_{\bar{t}} + \frac{\bar{\beta}\kappa^2(t)}{\beta(t)}\mathrm{e}^2(q(t)) \right)} + \frac{1}{T\underline{\beta}}\sum_{t \in \mathcal{T}}\kappa(t)\mathrm{e}(q(t)). \tag{6}$$

Here, $\bar{\beta} := \sum_{t \in \mathcal{T}} \beta(t)$, $\underline{\beta} := \min_{t \in \mathcal{T}} \beta(t)$, *and* $q(t) \geq K_1$ *is the number of available agents'* *feedback at the t-th outer iteration, $t \in \mathcal{T}$.* $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Roughly speaking, Theorem 1 establishes that, with arbitrarily high probability, the average value of the residual $\|\Delta_t^\star\|$ over a certain horizon T is bounded by the sum of two terms, which depend on the initial distance from a minimum for the unknown function θ, and the reconstruction error $\mathrm{e}(\cdot)$. Note that the terms in the RHS can be made small by either choosing a small step-size $\xi(t)$, in order to make $\kappa(t)$ close to zero, or tuning the product $c(t)\xi(t)$ close to one, thus leading to a large $\underline{\beta}$. This latter choice, however, would increase the term involving the sub-optimal constant $\Delta_{\bar{t}}$, thus requiring an accurate trade-off in tuning the gain $c(t)$ and the step-size $\xi(t)$. In the stationary case, to foster not exceedingly aggressive personalized actions the coordinator may then want to match the lower bound for $c(t)$, while striking a balance in choosing $\xi(t)$ to possibly boost the convergence of Algorithm 1.

For simplicity, let us now assume that $\beta(t)$ is a constant term. From Assumption 1, $\mathrm{e}(q(t)) \leq \mathrm{e}(K_1)$, and hence $\frac{1}{T}\sum_{t \in \mathcal{T}}\|\Delta_t^\star\| \leq O(1/\sqrt{T}) + O(1)$. We note that, as T grows, $O(1/\sqrt{T})$ vanishes, and the average of $\|\Delta_t^\star\|$ stays in a ball whose radius depends on the number of agents' feedback $q(t)$ made available to perform (S0) in Algorithm 1 and, specifically, on the learning strategy $\mathscr{L}$. Next, we analyze the bound above under the lens of different learning procedures.

C. Specifying the learning strategy $\mathscr{L}$

By requiring that the reconstruction error is bounded in probability, Assumption 1 is quite general and it holds true under standard assumptions for LS and GP approaches to learning G. In particular, we have the following:

- In parametric learning, if $G(\boldsymbol{x})$ is modelled as an affine function of the learning parameters η's, then setting up an LS approach to minimize the loss between the model parameters η and the agents' feedback leads to a convex quadratic program. Due to the large-scale properties of LS (under standard assumptions), the error term $\mathrm{e}(q(t))$ behaves as a normal distribution,

for which Assumption 1 holds true (see [21, Lemma A.4]), and $\lim_{K \to \infty} \mathrm{e}(K) = 0$.

- In non-parametric learning, suppose $G(\boldsymbol{x})$ is a sample path of a GP with zero mean and a certain kernel. Due to the large-scale property of such regressor and under standard assumptions, also in this case Assumption 1 holds true (see [11]) and $\lim_{K \to \infty} \mathrm{e}(K) = 0$.

Note that, in general, $q(t) \propto t$. Therefore, since $\sum_{t \in \mathcal{T}} \mathrm{e}(q(t)) = o(T)$, for the cases above we obtain $\lim_{T \to \infty} \frac{1}{T} \sum_{t \in \mathcal{T}} \|\Delta_t^\star\| = 0$, thus implying that, for T large enough, the average of $\|\Delta_t^\star\|$ converges to 0. If, in addition, there exists some $\bar{K} \geq 0$ such that $\mathrm{e}(K) = 0$ for all $K \geq \bar{K}$, then one is allowed to recover exactly the results obtained for the perfect reconstruction case shown in §IV-A.

V. THE TIME-VARYING CASE

We now investigate the GNEP in (1) in case the local cost function of each agent g_i varies in time, thus implying that also the function θ is non-stationary. Our goal is still to design the parameters defining the personalized incentives to track a time-varying GNE that minimizes the (dis)satisfaction function, i.e., some $\boldsymbol{x}^\star(t) \in \Theta(t)$, both in case of perfect (§V-A) and inexact reconstruction (§V-B) of the pseudo-gradient mapping.

To start, we make the following typical assumptions in the literature on online optimization [6], [8], [12].

Assumption 2. *For all $t \in \mathbb{N}$ and $\boldsymbol{x} \in \Omega$, it holds that*

i) $|\theta(\boldsymbol{x}; t) - \theta(\boldsymbol{x}; t-1)| \leq \mathrm{e}_\theta$, *for $0 \leq \mathrm{e}_\theta < \infty$;*

ii) *For all $i \in \mathcal{I}$, $\|\nabla_{x_i} g_i(\boldsymbol{x}; t) - \nabla_{x_i} g_i(\boldsymbol{x}; t-1)\| \leq \mathrm{e}_{\nabla_i}$, for $0 \leq \mathrm{e}_{\nabla_i} < \infty$;*

iii) $\|\boldsymbol{x}^\star(t) - \boldsymbol{x}^\star(t-1)\| \leq \mathrm{e}_\delta$, *for $0 \leq \mathrm{e}_\delta < \infty$.* $\qquad\square$

Assumptions 2 i) and ii) essentially bound the variation in time of both the unknown function θ and the pseudo-gradient mappings, while Assumptions 2 iii) guarantees the boundedness of the distance between two consecutive minima such that $\boldsymbol{x}^\star(t) \in \Theta(t)$ and $\boldsymbol{x}^\star(t-1) \in \Theta(t-1)$.

Note that, with these standard assumptions in place, an asymptotic error term of the form of $O(1)$ is inevitable [6], [12], [48], [49].

Lemma 5. *Let Assumption 2 ii) hold true. For all $t \in \mathbb{N}$, $\|\nabla\theta(\boldsymbol{x}; t) - \nabla\theta(\boldsymbol{x}; t-1)\| \leq \mathrm{e}_\nabla$, with $\mathrm{e}_\nabla := \sum_{i \in \mathcal{I}} \mathrm{e}_{\nabla_i}$.* □

A. Online perfect reconstruction of the pseudo-gradients

In case the learning procedure $\mathcal{L}$ allows for $\nabla_{x_i}\hat{g}_{i,t-1}(\boldsymbol{x}^\star_{t-1}; t-1) = \nabla_{x_i}g_i(\boldsymbol{x}^\star_{t-1}; t-1)$, for all $i \in \mathcal{I}$ and $t \in \mathbb{N}$, we have the following ancillary results:

Lemma 6. *Let Assumption 2 ii) hold true, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. Then, with the personalized incentives in (4), for all $t \in \mathbb{N}$ we have*

$$\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}^\star_{t-1}; t-1) \leq \tfrac{1}{4\alpha(t)\ell}\mathrm{e}_\nabla^2 - \left(\sqrt{\tfrac{\ell}{\alpha(t)}}\|\Delta_t^\star\| - \tfrac{1}{2}\sqrt{\tfrac{1}{\alpha(t)\ell}}\,\mathrm{e}_\nabla\right)^2. \tag{7}$$

□

As in the stationary case in §IV-B, the vector $\Delta_t^\star$ is not guaranteed to be a descent direction for the unknown mapping $\boldsymbol{x} \mapsto \theta(\boldsymbol{x}; t-1)$ in the sense of Lemma 3. In fact, the error e_∇, introduced because of the time-varying nature of the pseudo-gradients, excludes that the LHS in (5) is strictly negative. The following bound characterizes the sequence of v-GNE originating from Algorithm 1 in case $\mathcal{L}$ allows for a perfect reconstruction of the time-varying mapping $\boldsymbol{x} \mapsto G(\boldsymbol{x}; t)$.

Theorem 2. *Let Assumption 2 hold true, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. Moreover, let some $T \in \mathbb{N}$ be fixed, $\mathcal{T} := \{1, \dots, T\}$ and, for any global minimizer $\boldsymbol{x}^\star(0) \in \Theta(0)$, $\Delta_0 := |\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(0); 0)|$. Then, with the personalized incentives in (4), the sequence of v-GNE $(\boldsymbol{x}_t^\star)_{t \in \mathcal{T}}$, generated by Algorithm 1, satisfies the following relation*

$$\tfrac{1}{T}\sum_{t \in \mathcal{T}}\|\Delta_t^\star\| \leq \tfrac{\mathrm{e}_\nabla}{T\underline{\beta}}\sum_{t \in \mathcal{T}}\tfrac{1}{2\alpha(t)} + \tfrac{1}{T\underline{\beta}}\sqrt{\sum_{t \in \mathcal{T}}\left(\beta(t)\left(\Delta_0 + T\phi\right) + \tfrac{\bar{\beta}}{4\alpha^2(t)\beta(t)}\mathrm{e}_\nabla^2\right)}. \tag{8}$$

with $\underline{\beta} := \min_{t \in \mathcal{T}}\beta(t)$, $\phi := 2\mathrm{e}_\theta + \tfrac{\ell}{2}\mathrm{e}_\delta^2$ and $\bar{\beta} := \sum_{t \in \mathcal{T}}\beta(t)$. □

Theorem 2 says that the average of the residual $\|\Delta_t^\star\|$ over the horizon T is bounded by the sum of two terms, which depend on the initial sub-optimality of a computed v-GNE compared to a minimum for the unknown function θ, and several bounds on the variations in time of θ, G and constrained minima postulated in Assumption 2 and Lemma 5. In this case, the coordinator may reduce the error in the RHS by properly tuning the product $c(t)\xi(t)$ close to one, thus leading to a large $\underline{\beta}$, and hence possibly boosting the convergence of Algorithm 1. In fact, if the parameter $\beta(t)$ is fixed in time, we obtain $\frac{1}{T}\sum_{t\in\mathcal{T}}\|\Delta_t^\star\| \leq O(1)$. This inequality ensures that $\|\Delta_t^\star\|$ will be always contained into a ball of constant radius, whose value can be adjusted through $c(t)$ and $\xi(t)$.

B. Inexact estimate of the pseudo-gradients

As in §IV-B, we consider the case in which, due to possibly noisy agents' feedback, the learning procedure $\mathscr{L}$ does not allow a perfect reconstruction of each time-varying gradient g_i, $i \in \mathcal{I}$. First, we postulate the time-varying counterpart of Assumption 1, and then we provide a preliminary result.

Assumption 3. *For all $t \in \mathbb{N}$ and $\boldsymbol{x} \in \mathcal{X}$, $\hat{G}_t(\boldsymbol{x};t) := G(\boldsymbol{x};t) + \epsilon_t$, and, for any $\delta_1 \in (0,1]$, there exists $\bar{t} < \infty$ and available K_1 agents' feedback such that*

$$\mathbb{P}\{\|\epsilon_t\| \leq \mathrm{e}(K) \mid \forall t \geq \bar{t}, \forall K \geq K_1\} \geq 1 - \delta_1,$$

for some nonincreasing function $\mathrm{e} : \mathbb{N} \to \mathbb{R}_{\geq 0}$ such that $\mathrm{e}(K) < \infty$, for all $K \in \mathbb{N}$. □

Lemma 7. *Let Assumption 2 and 3 hold true for some fixed $\delta_1 \in (0,1]$, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. With the personalized incentives in (4), for all $t \geq \bar{t}$ we have*

$$\Delta_t^{\star^\top} \nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1) \leq \frac{1}{4\alpha(t)\ell}\sigma^2(K,t) - \left(\sqrt{\tfrac{\ell}{\alpha(t)}}\|\Delta_t^\star\| - \frac{1}{2\alpha(t)}\sqrt{\tfrac{\alpha(t)}{\ell}}\sigma(K,t)\right)^2, \qquad (9)$$

where $\sigma(K,t) := \mathrm{e}_\nabla + (1-\alpha(t))\mathrm{e}(K-1)$, with probability $1-\delta_1$, for some $K \geq K_1$. □

Along the same line drawn for the stationary case with inexact reconstruction, we now provide the following bound on the sequence of v-GNE, $(\boldsymbol{x}_t^\star)_{t\in\mathcal{T}}$, generated by Algorithm 1. Note the slight abuse of notation in defining $\Delta_{\bar{t}}$, which is different from the one in Theorem 1.

Theorem 3. *Let Assumption 2 and 3 hold true for some fixed $\delta_1 \in (0,1]$, $c(t) \geq 2\ell$ and $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$. Moreover, let some $T \in \mathbb{N}$ be fixed, $\mathcal{T} := \{\bar{t}+1, \ldots, \bar{t}+T\}$ and, for any global minimizer $\boldsymbol{x}^\star(\bar{t}) \in \Theta(\bar{t})$, $\Delta_{\bar{t}} := |\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(\bar{t}); \bar{t})|$. Then, with the personalized incentives in (4), the sequence of v-GNE $(\boldsymbol{x}_t^\star)_{t \in \mathcal{T}}$, generated by Algorithm 1, satisfies the following relation*

$$\frac{1}{T}\sum_{t \in \mathcal{T}}\|\Delta_t^\star\| \leq \frac{1}{T\underline{\beta}}\sum_{t \in \mathcal{T}}\frac{1}{2\alpha(t)}\sigma(K,t) + \frac{1}{T\underline{\beta}}\sqrt{\sum_{t \in \mathcal{T}}\left(\beta(t)\left(\Delta_{\bar{t}} + T\phi\right) + \frac{\bar{\beta}}{4\alpha^2(t)\beta(t)}\sigma^2(K,t)\right)}, \quad (10)$$

with probability $1 - \delta_1$, where $\underline{\beta} := \min_{t \in \mathcal{T}}\beta(t)$, $\bar{\beta} := \sum_{t \in \mathcal{T}}\beta(t)$ and $q(t) \geq K_1$ is the number of agents' feedback at the t-th outer iteration. $\quad\square$

Also in this case, the average of the residual $\|\Delta_t^\star\|$ over the horizon T is bounded by the sum of two terms, which depend, among the others, on the reconstruction error of the mapping G and its variations in time. We note that the bound in the RHS of (4) can be adjusted through an accurate choice of the gain $c(t)$ and the step-size $\xi(t)$. Specifically, choosing a small $\xi(t)$ reduces the reconstruction error, hidden in the variable σ, while setting $c(t)\xi(t)$ close to one induces a large value for $\underline{\beta}$ (and for $\beta(t)$ as well), thus possibly eliminating the second term under the square root of (10), and the one outside.

For simplicity, let us now suppose that the parameters $c(t)$ and $\xi(t)$ of the personalized incentives in (4) are fixed in time, namely $\beta(t)$ is a constant term. From (10), we note that $\frac{1}{T}\sum_{t \in \mathcal{T}}\|\Delta_t^\star\| \leq O(1)$. Due to the time-varying nature of the problem in question, also in this case the average residual $\|\Delta_t^\star\|$ can not vanish as T grows, albeit the radius of the error ball can be reduced through a fine tuning of $c(t)$ and $\xi(t)$.

C. Specifying the time-varying learning strategy $\mathscr{L}$

In a time-varying setting, one cannot expect $\mathrm{e}(K)$ to vanish in general, since the time variations in θ are not supposed to be asymptotically vanishing [6], [12]. Popular learning approaches include LS with forgetting factors [50] and time-varying GP [7], for which we have $\lim_{K \to \infty}\mathrm{e}(K) = O(1)$.

VI. Ride-hailing with MaaS orchestration

With the growing business related to ride-hailing, a MaaS coordination platform appears indispensable to contrast the traffic congestion due to the increasing number of vehicles dispatched on the road, while facilitating the competition among service providers [51]–[53]. In this section we develop a mathematical model capturing intrinsic features of the problem considered, and then use it to verify our theoretical findings.

A. Mathematical model description

We consider a scenario in which N companies compete to put the most vehicles (a capped local resource, $0 \leq x_i \leq \bar{x}_i$) on the road to attract the most customers. During the day, each company aims at maximizing its (time-varying) profit P_i, which is implicitly related to how many cars it could currently put on the road to meet customer needs, properly discounted to account for, e.g., refusals rates or the time of the day, t. To this end, bigger companies can be naïvely induced to dispatch as many cars as they own. However, this may cause traffic congestion, thus reducing the quality of the service provided, and therefore lessen what the company can charge for each ride. In fact, by leveraging their own experience, those big companies may estimate how many cars actually get customers on top of the available x_i as, e.g., a *concave* function $a_i(t)x_i - b_i(t)x_i^2 \leq x_i$, with $a_i, b_i \geq 0$ tuned accordingly. Therefore, assuming the same fare $r_i(t)$ applies per average trip to each costumer, the profit function of the i-th lead company can read as $P_i(x_i; t) = r_i(t)(a_i(t)x_i - b_i(t)x_i^2)$. On the other hand, the strategies of smaller companies are typically less affected by traffic congestions, since the quality service is generally worse in the sense that they can dispatch a little number of cars on the road. In this case, their direct experience may suggest that the number of cars that actually get customers on the available x_i can be modelled as a *convex* function $a_i(t)x_i + b_i(t)x_i^2 \leq x_i$, thus reflecting the fact that the larger the number of deployed cars, the larger the possibility to cover enough space to be attractive. The overall time-varying profit hence reads as $P_i(x_i; t) = r_i(t)(a_i(t)x_i + b_i(t)x_i^2)$.

In addition to the profit, however, the companies also incur in costs that have to be minimized and vary during the day, such as gas consumed or miles travelled, here represented by some

$d_i(t) \geq 0$. By assuming, for instance, the same cost associated to each vehicle per average trip, we model the overall cost as $C_i(x_i, \boldsymbol{x}_{-i}; t) = d_i(t)x_i + \sum_{j \in \mathcal{I} \setminus \{i\}} w_{i,j}(x_i - x_j)^2$, where $(x_i - x_j)^2$ enables for competition among equally-sized companies, weighted with $w_{i,j} = w_{j,i} \geq 0$ to preserve symmetries.

After splitting the hours of a day in intervals, enumerated in the set $\mathcal{N}$ according to the estimated average travel time of each costumer with no shared trips from [37] (about 15 minutes), at every $t \in \mathcal{N}$ each firm $i \in \mathcal{I}$ aims at solving the following mutually inter-dependent optimization problem:

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in [\underline{x}_i, \bar{x}_i]} & C_i(x_i, \boldsymbol{x}_{-i}; t) - P_i(x_i; t) \\ \\ \text{s.t.} & A\boldsymbol{x} \leq q, \end{cases} \tag{11}$$

where A and q collect constraints so that the total number of vehicles involved in the service is capped, i.e., $\mathbf{1}^\top \boldsymbol{x} \in [\underline{x}, \bar{x}]$, for some $0 \leq \underline{x} \leq \bar{x}$, and the fact that firm i wants to improve its service compared to j, thus requiring $x_i \geq x_j$.

In the proposed scenario where the firms exhibit symmetries in the mixed convex-concave cost functions[1], the time-varying, unknown exact potential function for the GNEP in (11) is

$$\theta(\boldsymbol{x}; t) := \sum_{i \in \mathcal{I}} \left(-P_i(x_i; t) + \sum_{j \in \mathcal{I}, j < i} C_i(x_i, \boldsymbol{x}_{-i}; t) \right),$$

and the MaaS platform aims at coordinating the whole ride-hailing service while avoiding traffic congestion. This can be achieved, for instance, by imposing extra fees, incentives or restrictions to the companies, possibly according to their size and turnover. However, note that the parameters a_i and b_i affecting the cost function of each firm, which are hence key to drive its strategy, can not be disclosed to the MaaS platform, since they represent sensitive information, as opposed to the incurred cost $C_i(x_i; t)$ that can be estimated directly, as it depends on mileage, fuel consumption and number of deployed cars. These data are indeed not sensitive, and therefore can be either publicly available as in [37] (e.g., ride-hailing trips per day, average vehicles per day, daily trips per vehicle, minutes per trip, fare-box per trip, etc.), or may be estimated from them. Thus, a

---

[1]The adopted convex-concave costs can be thought as proxy for the actual objective that has to be minimized, which however can be complicated by considering, e.g., sum of exponential or logarithmic functions parametrized in the sensitive private quantities the coordinator aims at learning.

TABLE I: Simulation parameters

Parameter	Description	Value
$\ell(t)$	Constant of weak convexity	$[0.38, 5.2]$
$a_i(t)$	Discount parameter – profit function	$[0.92, 0.94]$
$b_i(t)$	Discount parameter – profit function	$[-0.41, 6.6] \times 10^{-4}$
$r_i(t)$	Trip fare	$[14, 30]$ [$]
$d_i(t)$	Trip cost	$[4, 10]$ [$]
P	Weight for the extragradient algorithm	I
l_F	Horizontal scale – GP method	2×10^3
σ_F	Vertical scale – GP method	10^4
γ	Discount factor – GP method	0.8

possible strategy requires the MaaS platform to learn those time-varying parameters by leveraging feedback collected from users, e.g., on the price they are charged $r_i(t)$, and then design tailored incentives for the coordination. The fact that only a subset of deployed vehicles gets customers intrinsically represents the noise in the costumers' feedback, as the profit of each firm does not naïvely coincide with the applied fare $r_i(t)$ times the number of vehicles $x_i(t)$.

B. Numerical simulations

The open data collected in New York City in April 2019 [37] provide us information on $N = 5$ main companies: Yellow taxi, Uber, Lyft, Juno and Via. We stress that the quality of the ride-hailing service, measured as the total number of vehicles deployed on the road, coincides with an integer variable, i.e., $n_i = 1$, for all $i \in \mathcal{I} = \{1, \ldots, 5\}$, thus leading to a mixed-integer setting. However, since the fleet dimension of each firm we are considering is in the order of few thousands of vehicles (i.e., Juno and Via), or tens of thousands for bigger companies (Uber, Lyft, Yellow taxi), we consider a relaxed version by treating x_i as a scalar continuous variable, and then rounding its value [51]. For this reason, we roughly estimate a round-off error in the order of 10^{-3} for any GNE computed at item (S2) in Algorithm 1 through an extragradient type method [54] (thus neglecting the multi-agent nature of the inner loop).

To learn the unknown terms characterizing each pseudo-gradient associated to the cost function
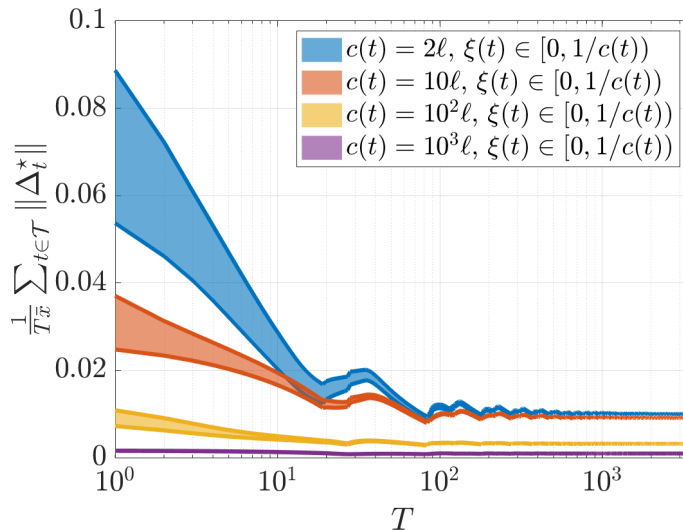
Fig. 2: Residual $\Delta_t^\star$ over the time horizon T, normalized by the total number of deployable cars at any t, for different values of personalized incentive feedback gains $c(t)$ and $\xi(t)$ in (4).

of each company, i.e., the time-varying parameters a_i and b_i, we assume the MaaS platform being endowed with a GP learning algorithm. Further numerical examples with other regression strategies can be found in the preliminary work [31], where we considered a specific class of stationary nonmonotone GNEPs only.

The main parameters adopted to run numerical experiments are summarized in Table I, where $d_i(t)$ and $r_i(t)$ are time-varying functions designed to capture the costumer requests variability over an entire week (see the dotted green line in Fig. 4 for instance). We stress that the information on the costumer demands is not directly exploited anywhere, except for the design of functions $d_i(t)$ and $r_i(t)$, i.e., the numerical results are obtained without artificially imposing the service requests as a reference trajectory. According to the data available at [37], in April 2019 Yellow taxi represented approximately the 27% of the whole car-based mobility market, whereas among the ride-hailing firms Uber impacted for the 72% on the market, thus representing the leading company with potentially no competitors, Lyft for the 19.7%, Via and Juno for the 4.8% and 3.4%, respectively. Thus, by comparing the vehicle fleet size characterizing each firm, it turns out that the weight matrix $W := [w_{i,j}]_{i \in \mathcal{I}, j \in \mathcal{I}}$ links Yellow taxi and Lyft, and Via and Juno only. These companies, indeed, compete each other to provide at least a comparable service for the
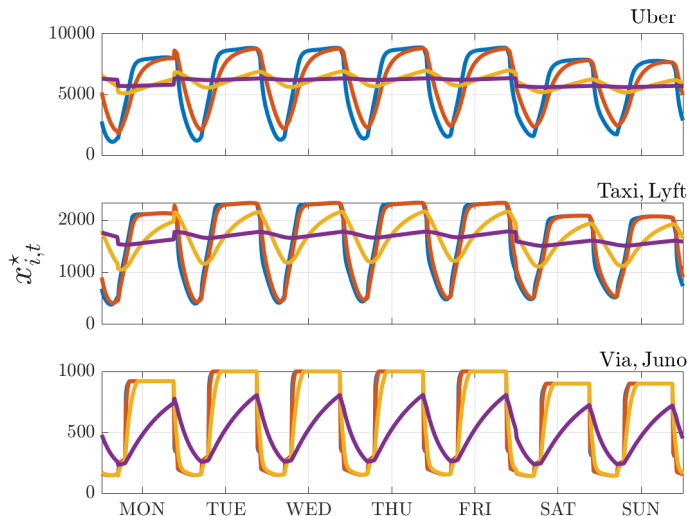
Fig. 3: Optimal strategies pursued by the five companies over a week, coinciding with the number of vehicles deployed at every time interval. The colour line follows the legend in Fig. 2.

sake of reputation.

In Fig. 2 is reported the average value of the residual $\Delta_t^\star$ over the time horizon T, normalized by its magnitude. This latter, indeed, was identified as a candidate metric to assess convergence of the proposed algorithm, with different values of personalized incentive gains $c(t)$ and $\xi(t)$. Note that increasing $\xi(t)$ noticeably reduces the residual $\Delta_t^\star$ for short time horizon T, while a larger $c(t)$ seems providing a smaller asymptotic error. Given the time variability of the costumer demands, however, this may result in firms' optimal strategies leading to an overall unsatisfactorily service, as evidenced by Fig. 3 and, more prominently, by Fig. 4. In fact, for $c(t) = 2\ell(t)$ (or even $10\ell(t)$, blue and brown lines), the MaaS platform manages to accomplish the task of serving hundreds of thousands of costumers per week, whereas for larger values of $c(t)$ the companies do not capture the variability of the demand (yellow and violet lines), adopting almost constant strategies over the week and hence experiencing higher costs, as reported in Fig. 5. This represents the trade-off that one has to strike between performance of Algorithm 1 and expected/desired behaviour of the companies and, more importantly, of the overall service provided. For this reason, the MaaS platform is a key tool enabling for competition among companies in ride-hailing mobility while, at the same time, guaranteeing a certain degree of service by satisfying costumer requests.
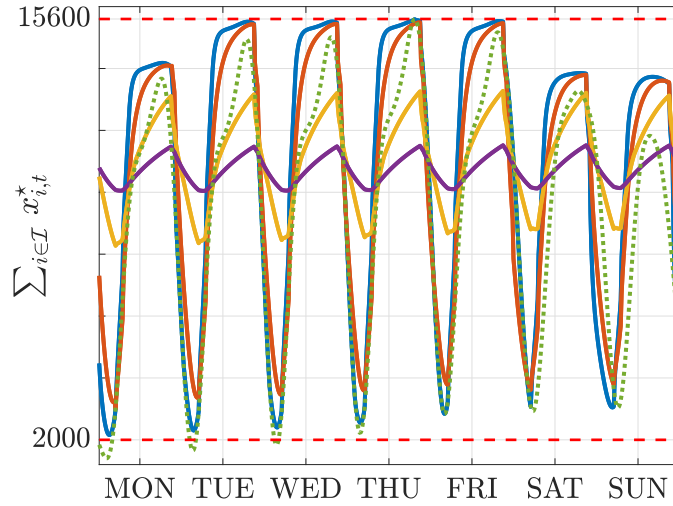
Fig. 4: Overall service provided by the five firms over a whole week. The colour line follows the legend in Fig. 2, while the dotted green line corresponds to the costumer requests obtained from data after a spline interpolation. The dashed red lines coincide with $\underline{x}$ and $\bar{x}$, respectively.
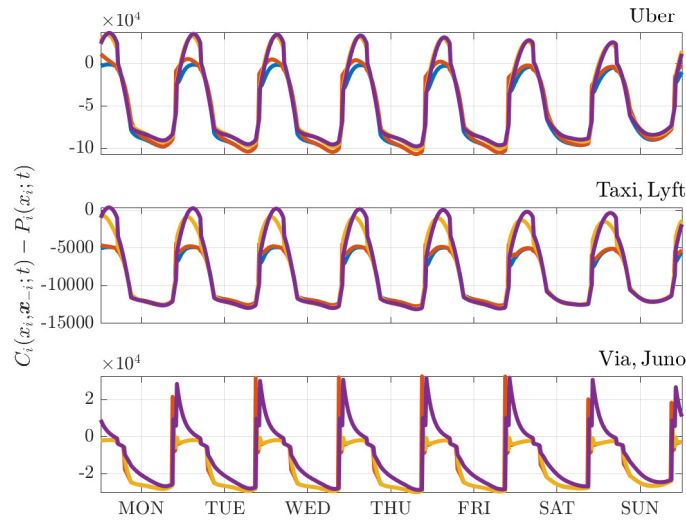

Fig. 5: Profit/cost experienced by the companies over a week. The colour line follows the legend in Fig. 2.

We finally give further motivation to the importance of a MaaS orchestration system as the one proposed in this paper in terms of potential impact to alleviate traffic congestion. In [37], indeed,
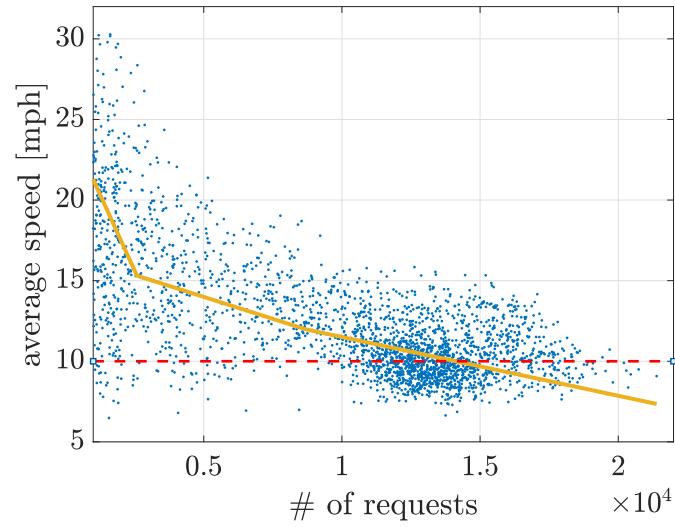
Fig. 6: Piecewise-affine interpolation linking the number of costumer requests and average speed per trip.
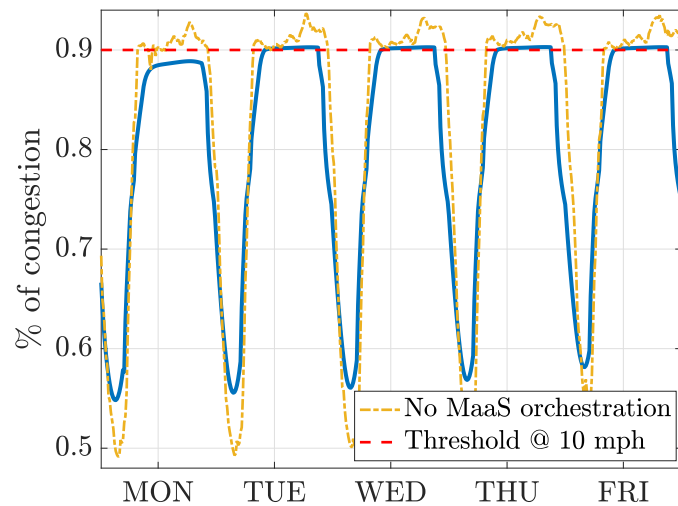
Fig. 7: Comparison between the traffic congestion caused by ridehailing service with and without MaaS orchestration.

is shown that each vehicle contributes to two rides per hour on average, while for the rest of the time is idle. By considering an average trip length of 15 minutes, this potentially leads to more than 50% of cars deployed on the road than necessary (i.e., in accordance with the costumer

requests), thus heavily affecting the urban traffic, especially in big cities such as New York. Then, by leveraging available data, we first extrapolate the relation among number of costumer requests and average speed per trip, as described in Fig. 6 where we have adopted a piece-wise affine function to characterize this relation. Successively, by assuming a conservative estimate of up to the 25% of vehicles dispatched more than the actual number of costumer requests, we evaluate the effect of the MaaS platform on the traffic congestion in Fig. 7. By focusing on weekdays, in which the traffic congestion produces the greatest costumer discomfort, the MaaS orchestration seems to represent a crucial component to avoid deploying an excessive number of cars on the road, thus avoiding to reach and surpass the threshold at 90% corresponding to the critical speed per trip, identified as 10 miles per hour.

VII. Conclusion and Outlook

We have shown that a suitable design of parametric personalized incentives is crucial to compute (or track in a neighbourhood) GNE in nonmonotone GNEP characterized by symmetric interactions among agents, both in static and time-varying setting. First, the designed functionals act as regularization terms of the agents' cost functions, thus allowing the agents for the practical computation of a v-GNE at each outer iteration of the proposed two-layer algorithm. Then, they provide a mean to boost the convergence of the algorithm in the stationary case, or to adjust the asymptotic error obtained in a time-varying setting. In the static case, the proposed algorithm converges to a GNE by exploiting the asymptotic consistency bounds characterizing standard learning procedures for the coordinator, such as LS or GP , while in a time-varying setting our metric for assessing convergence asymptotically behaves as $O(1)$, and hence the semi-decentralized scheme allows the agents to track a GNE in a neighbourhood of adjustable size.

Future research directions may include for instance the design of other possible parametric personalized incentives, as well as non-parametric ones, in which the pseudo-gradient mappings are directly learned through, e.g., neural networks. The case of possibly approximate computation of a v-GNE at item `(S2)` may be explored, also working towards the relaxation of the symmetric interaction requirement.

APPENDIX

A. Proofs of §II

Proof of Lemma 1: Let some $t \in \mathbb{N}$ be fixed. Then, in view of Standing Assumption 1 and 2, for any $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^n$, the following chain of inequalities hold $\|\nabla\theta(\boldsymbol{x};t) - \nabla\theta(\boldsymbol{y};t)\| = \|G(\boldsymbol{x};t) - G(\boldsymbol{y};t)\| = \|\mathrm{col}((\nabla_{x_i} g_i(x_i, \boldsymbol{x}_{-i};t) - \nabla_{y_i} g_i(y_i, \boldsymbol{y}_{-i};t))_{i\in\mathcal{I}})\| \leq \sum_{i\in\mathcal{I}} \|\nabla_{x_i} g_i(x_i, \boldsymbol{x}_{-i};t) - \nabla_{y_i} g_i(y_i, \boldsymbol{y}_{-i};t)\| \leq \sum_{i\in\mathcal{I}} \ell_i \|\boldsymbol{x} - \boldsymbol{y}\| =: \ell\|\boldsymbol{x} - \boldsymbol{y}\|$. Finally, the fact that $\boldsymbol{x} \mapsto \theta(\boldsymbol{x};t)$ is a $\mathcal{C}^1$-smooth function with ℓ-Lipschitz continuous gradient directly entails that, for any $\boldsymbol{x} \in \Omega$, the auxiliary function $\psi(\boldsymbol{x};t) := \theta(\boldsymbol{x};t) + \frac{\ell}{2}\|\boldsymbol{x}\|^2$ is convex, which in turn implies the ℓ-weak convexity of $\boldsymbol{x} \mapsto \theta(\boldsymbol{x};t)$ [9], [55]. ■

B. Proofs of §III

Proof of Proposition 1: With the personalized incentive functionals in (4), $F(\boldsymbol{x};t) = G(\boldsymbol{x};t) + c(t)(\boldsymbol{x} - \boldsymbol{x}_{t-1}^\star - \xi(t)\hat{G}_{t-1}(\boldsymbol{x}_{t-1}^\star))$. Thus, for any $\boldsymbol{x}, \boldsymbol{y} \in \Omega$, and $t \in \mathbb{N}$, we have that: $(\boldsymbol{x} - \boldsymbol{y})^\top(F(\boldsymbol{x};t) - F(\boldsymbol{y};t)) = (\boldsymbol{x} - \boldsymbol{y})^\top(G(\boldsymbol{x};t) + c(t)(\boldsymbol{x} - \boldsymbol{x}_{t-1}^\star - \xi(t)\hat{G}_{t-1}(\boldsymbol{x}_{t-1}^\star)) - G(\boldsymbol{y};t) - c(t)(\boldsymbol{y} - \boldsymbol{x}_{t-1}^\star - \xi(t)\hat{G}_{t-1}(\boldsymbol{x}_{t-1}^\star))) = (\boldsymbol{x} - \boldsymbol{y})^\top(G(\boldsymbol{x};t) + c(t)\boldsymbol{x} - G(\boldsymbol{y};t) - c(t)\boldsymbol{y}) = (\boldsymbol{x} - \boldsymbol{y})^\top(\nabla\theta(\boldsymbol{x};t) + c(t)\boldsymbol{x} - \nabla\theta(\boldsymbol{y};t) - c(t)\boldsymbol{y})$, where this last equality follows from Standing Assumption 2. Then, let us introduce the auxiliary function $\psi(\boldsymbol{x};t) := \theta(\boldsymbol{x};t) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$. Note that, in case $c(t) \geq 2\ell$ for all $t \in \mathbb{N}$, $\boldsymbol{x} \mapsto \psi(\boldsymbol{x};t)$ is ℓ-strongly convex in view of the ℓ-weak convexity of θ proved in Lemma 1, which yields to

$$(\boldsymbol{x} - \boldsymbol{y})^\top(F(\boldsymbol{x};t) - F(\boldsymbol{y};t))$$
$$= (\boldsymbol{x} - \boldsymbol{y})^\top(\nabla\theta(\boldsymbol{x};t) + c(t)\boldsymbol{x} - \nabla\theta(\boldsymbol{y};t) - c(t)\boldsymbol{y})$$
$$= (\boldsymbol{x} - \boldsymbol{y})^\top(\nabla\psi(\boldsymbol{x};t) - \nabla\psi(\boldsymbol{y};t)) \geq \ell\|\boldsymbol{x} - \boldsymbol{y}\|^2,$$

i.e., the definition of strongly monotone mapping. ■

Proof of Lemma 2: Consider the stationary points of $\theta(\boldsymbol{x})$ as the ones satisfying the first-order optimality conditions: $w\nabla\theta(\boldsymbol{x}) + \mathcal{N}_\Omega(\boldsymbol{x}) \ni \boldsymbol{0}$, where $w > 0$ is any scalar scaling of the cost θ, and

$\mathcal{N}_\Omega$ is the normal cone operator of the feasible set Ω. At every $t \in \mathbb{N}$, $\boldsymbol{x}_t^\star$ with perfect reconstruction is the solution to $\nabla\theta(\boldsymbol{x}_t^\star) + c(t)(\boldsymbol{x}_t^\star - \boldsymbol{x}_{t-1}^\star - \xi(t)\nabla\theta(\boldsymbol{x}_{t-1}^\star)) + \mathcal{N}_\Omega(\boldsymbol{x}_t^\star) \ni \boldsymbol{0}$, If $\|\Delta_t^\star\| = 0$ then $\boldsymbol{x}_t^\star = \boldsymbol{x}_{t-1}^\star$, which means that $\boldsymbol{x}_t^\star$ is the solution to $(1 - c(t)\xi(t))\nabla\theta(\boldsymbol{x}_t^\star) + \mathcal{N}_\Omega(\boldsymbol{x}_t^\star) \ni \boldsymbol{0}$, which satisfies the first-order optimality conditions for $\theta(\boldsymbol{x})$ and it is therefore one of its stationary points. ∎

C. Proofs of §IV

Proof of Lemma 3: In view of item (S2), at every outer iteration $t \in \mathbb{N}$ we have by (3) that $\boldsymbol{x}_t^\star$ satisfies $(\boldsymbol{y} - \boldsymbol{x}_t^\star)^\top F(\boldsymbol{x}_t^\star; t) \geq 0$ for all $\boldsymbol{y} \in \Omega$, and therefore, since $\boldsymbol{x}_{t-1}^\star \in \Omega$ as it is a v-GNE at $t-1$, $\Delta_t^{\star\top} F(\boldsymbol{x}_t^\star; t) \leq 0$. Thus, by adding and subtracting the term $\Delta_t^{\star\top} G(\boldsymbol{x}_{t-1}^\star)$, we obtain $\Delta_t^{\star\top} G(\boldsymbol{x}_{t-1}^\star) \leq \Delta_t^{\star\top}(G(\boldsymbol{x}_{t-1}^\star) - (G(\boldsymbol{x}_t^\star) + U(\boldsymbol{x}_t^\star; t)))$. With the personalized functionals in (4) and the perfect estimate of the pseudo-gradients, we have that $U(\boldsymbol{x}_t^\star; t) = c(t)(\Delta_t^\star - \xi(t)G(\boldsymbol{x}_{t-1}^\star))$. Then, in view of Standing Assumption 2, it follows that

$$(1 - c(t)\xi(t))\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq \Delta_t^{\star\top}(\nabla\theta(\boldsymbol{x}_{t-1}^\star) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla\theta(\boldsymbol{x}_t^\star) - c(t)\boldsymbol{x}_t^\star).$$

Now, by defining $\alpha(t) := 1 - c(t)\xi(t)$ for all $t \in \mathbb{N}$, let us introduce the auxiliary function $\psi(\boldsymbol{x}; t) := \theta(\boldsymbol{x}) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$. Note that, in case $c(t) \geq 2\ell$ for all $t \in \mathbb{N}$, $\boldsymbol{x} \mapsto \psi(\boldsymbol{x}; t)$ is ℓ-strongly convex in view of the ℓ-weak convexity of θ, proved in Lemma 1. As such, $\Delta_t^{\star\top}(-\nabla\psi(\boldsymbol{x}_{t-1}^\star; t) + \nabla\psi(\boldsymbol{x}_t^\star; t)) \geq \ell\|\Delta_t^\star\|^2$. Therefore, we have $\alpha(t)\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq \Delta_t^{\star\top}(\nabla\psi(\boldsymbol{x}_{t-1}^\star; t) - \nabla\psi(\boldsymbol{x}_t^\star; t)) \leq -\ell\|\Delta_t^\star\|^2$. Finally, we obtain the desired result, namely $\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq -(\ell/\alpha(t))\|\Delta_t^\star\|^2 < 0$, by imposing $\alpha(t) = 1 - c(t)\xi(t) > 0$, i.e., $\xi_i < 1/c(t) \leq 1/2\ell$, for all $t \in \mathbb{N}$. ∎

Proof of Proposition 2: By combining the descent lemma [56, Prop. A.24] and Lemma 3, the sequence $(\boldsymbol{x}_t^\star)_{t\in\mathbb{N}}$ satisfies

$$\theta(\boldsymbol{x}_t^\star) \leq \theta(\boldsymbol{x}_{t-1}^\star) + (\Delta_t^\star)^\top\nabla\theta(\boldsymbol{x}_{t-1}^\star) + \frac{\ell}{2}\|\Delta_t^\star\|^2$$

$$\leq \theta(\boldsymbol{x}_{t-1}^\star) - \ell\frac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2.$$

Thus, by imposing $(2 - \alpha(t))/2\alpha(t) > 0$, which entails that $0 \leq \xi(t) < 1/c(t) \leq 1/2\ell$, then the sequence $(\theta(\boldsymbol{x}_t^\star))_{t\in\mathbb{N}}$ shall converge to a finite value, as $\theta(\boldsymbol{x}_t^\star) \to -\infty$ can not happen in view of the compactness of $\mathcal{X}$ (Standing Assumption 1). Therefore, by the continuity of θ, the

convergence of $(\theta(\boldsymbol{x}_t^\star))_{t\in\mathbb{N}}$ entails that $\lim_{t\to\infty}\|\Delta_t^\star\| = 0$, and hence by Lemma 2, the bounded sequence of (feasible by definition) points $(\boldsymbol{x}_t^\star)_{t\in\mathbb{N}} \to \boldsymbol{x} \in \Theta^{\mathrm{s}}$. ∎

Proof of Lemma 4: By mimicking the same steps at the beginning of the proof of Lemma 3, in view of Assumption 1 we obtain, for all $t \geq \bar{t}$, $\alpha(t)\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq \Delta_t^{\star\top}(\nabla\theta(\boldsymbol{x}_{t-1}^\star) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla\theta(\boldsymbol{x}_t^\star) - c(t)\boldsymbol{x}_t^\star) + c(t)\xi(t)\Delta_t^{\star\top}\epsilon_{t-1}$. In case $c(t) \geq 2\ell > 0$ for all $t \in \mathbb{N}$, the first term in the RHS is always upper bound by $-\ell\|\Delta_t^\star\|^2$ in view of the ℓ-strongly convexity of the auxiliary function $\theta(\boldsymbol{x}) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$. On the other hand, the second term attains its maximum positive module when the two vectors $\Delta_t^\star$ and ϵ_{t-1} are aligned (since $\xi(t) \geq 0$), thus leading to the following chain of inequalities

$$\alpha(t)\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq -\ell\|\Delta_t^\star\|^2 + c(t)\xi(t)\|\Delta_t^\star\|\|\epsilon_{t-1}\|$$

$$\leq -\ell\|\Delta_t^\star\|^2 + c(t)\xi(t)\|\Delta_t^\star\|\mathrm{e}(K-1),$$

where the last relation follows from Assumption 1, for any $K \geq K_1$, with probability $1 - \delta_1$. Thus, for $\alpha(t) > 0$, i.e., $\xi(t) < 1/c(t)$, and $c(t)\xi(t) = 1 - \alpha(t)$, we obtain

$$\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq -\tfrac{\ell}{\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1-\alpha(t)}{\alpha(t)}\|\Delta_t^\star\|\mathrm{e}(K-1). \tag{12}$$

Finally, adding and subtracting the term $((1-\alpha(t))^2/4\alpha(t)\ell)\mathrm{e}^2(K-1)$ in the RHS directly yields to

$$\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq -\left(\sqrt{\tfrac{\ell}{\alpha(t)}}\|\Delta_t^\star\| - \tfrac{1-\alpha(t)}{2\alpha(t)}\sqrt{\tfrac{\alpha(t)}{\ell}}\,\mathrm{e}(K-1)\right)^2 + \tfrac{(1-\alpha(t))^2}{4\alpha(t)\ell}\mathrm{e}^2(K-1),$$

with probability $1 - \delta_1$. The proof concludes by substituting $\kappa(t) := (1 - \alpha(t))/2\alpha(t)$ in the inequality above. ∎

Proof of Theorem 1: As in the proof of Proposition 2, we start by combining the descent lemma [56, Prop. A.24] and Lemma 4 (specifically, the relation in (12)). In view of Assumption 1, for all $t \geq \bar{t}$ and $K \geq K_1$, we obtain

$$\theta(\boldsymbol{x}_t^\star) \leq \theta(\boldsymbol{x}_{t-1}^\star) - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1-\alpha(t)}{\alpha(t)}\|\Delta_t^\star\|\mathrm{e}(K-1)$$

with probability $1 - \delta_1$. By defining $\beta(t) := \ell(2 - \alpha(t))/2\alpha(t)$, which is guaranteed to be strictly greater than zero in case $\xi(t) \in [0, 1/c(t))$, for all $t \in \mathbb{N}$ (and hence for $t \geq \bar{t}$), and focusing on the second and third terms in the RHS of this latter inequality, we can complete the square by

adding and subtracting $((1-\alpha(t))^2/4\alpha^2(t)\beta(t))\mathrm{e}^2(K-1)$. After few algebraic manipulations, we obtain the following inequality

$$\theta(\boldsymbol{x}_t^\star) \leq \theta(\boldsymbol{x}_{t-1}^\star) - \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1-\alpha(t)}{2\alpha(t)\beta(t)}\mathrm{e}(K-1)\right)^2 + \tfrac{(1-\alpha(t))^2}{4\alpha^2(t)\beta(t)}\mathrm{e}^2(K-1). \tag{13}$$

From now on, we follow the proofline introduced in [9, Th. 3.1], [55, Th. 1], as in (13) we have obtained a typical structure characterizing the sample complexity analysis for sample-based stochastic algorithms (see, for instance, [9, Eq.(1.3)]). Specifically, by fixing some $T \in \mathbb{N}$ and summing up the inequality above over $t \in \mathcal{T} := \{\bar{t}+1, \ldots, T+\bar{t}\}$, in view of the possibly sporadic communication between the agents and the central coordinator, let $q(t)$ be a scalar indicating the number of agents' feedback adopted by the learning procedure $\mathscr{L}$ at the t-th outer iteration. In particular, by assuming to initialize Algorithm 1 with $q(0) = K_1$ agents' feedback, $q(t) = q(t-1)+1$ if, at the t-th iteration, the agents communicate $p(t)$ to the coordinator, $q(t) = q(t-1)$ otherwise. Then, for any global minimizer $\boldsymbol{x}^\star \in \Theta$, let $\kappa(t) := (1-\alpha(t))/2\alpha(t)$. With probability $1-\delta_1$ we obtain $\theta(\boldsymbol{x}^\star) \leq \theta(\boldsymbol{x}_T^\star) \leq \theta(\boldsymbol{x}_{\bar{t}}^\star) - \sum_{t\in\mathcal{T}}\beta(t)(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t)))^2 + \sum_{t\in\mathcal{T}}\tfrac{\kappa^2(t)}{\beta(t)}\mathrm{e}^2(q(t))$. After moving the term $\theta(\boldsymbol{x}^\star)$ to the RHS, the term with $\beta(t)$ to the LHS, and introducing $\Delta_{\bar{t}} := \theta(\boldsymbol{x}_{\bar{t}}^\star) - \theta(\boldsymbol{x}^\star) \geq 0$, then:

$$\sum_{t\in\mathcal{T}}\beta(t)\left(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t))\right)^2 \leq \Delta_{\bar{t}} + \sum_{t\in\mathcal{T}}\tfrac{\kappa^2(t)}{\beta(t)}\mathrm{e}^2(q(t)). \tag{14}$$

Since $\beta(t) > 0$ for all $t \in \mathcal{T}$, we can leverage the Jensen's inequality [57, Th. 3.4] on the convex function $(\cdot)^2$ to lower bound the summation in the LHS. To this end, we normalize the coefficient 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 $\bar{\beta}(\sum_{t\in\mathcal{T}}\hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t)))^2) \geq \bar{\beta}(\sum_{t\in\mathcal{T}}\hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t))))^2$. Then, replacing this latter inequality into (14),

$$\bar{\beta}\left(\sum_{t\in\mathcal{T}}\hat{\beta}(t)\left(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t))\right)\right)^2 = \tfrac{1}{\bar{\beta}}\left(\sum_{t\in\mathcal{T}}\beta(t)\left(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t))\right)\right)^2$$

$$\leq \Delta_{\bar{t}} + \sum_{t\in\mathcal{T}}\tfrac{\kappa^2(t)}{\beta(t)}\mathrm{e}^2(q(t)),$$

and performing few algebraic manipulations directly yields to,

$$\sum_{t\in\mathcal{T}}\beta(t)\left(\|\Delta_t^\star\| - \tfrac{\kappa(t)}{\beta(t)}\mathrm{e}(q(t))\right) \leq \sqrt{\sum_{t\in\mathcal{T}}\left(\beta(t)\Delta_{\bar{t}} + \tfrac{\bar{\beta}\kappa^2(t)}{\beta(t)}\mathrm{e}^2(q(t))\right)}.$$

By bringing the term $-\sum_{t\in\mathcal{T}}\kappa(t)\mathrm{e}(q(t))$ in the RHS, note that the obtained inequality still holds true if we premultiply both sides by $1/T$, thus obtaining the average over T. Moreover, we have that $(1/T)\sum_{t\in\mathcal{T}}\beta(t)\|\Delta_t^\star\| \geq (1/T)\underline{\beta}\sum_{t\in\mathcal{T}}\|\Delta_t^\star\|$, with $\underline{\beta} := \min_{t\in\mathcal{T}}\beta(t)$, which directly yields to the relation in (6) with probability $1 - \delta_1$, hence concluding the proof. ∎

D. Proofs of §V

Proof of Lemma 5: Let some $t \in \mathbb{N}$ be fixed. Then, in view of Standing Assumption 1, 2, and Assumption 2 ii), for any $\boldsymbol{x} \in \mathbb{R}^n$ the following chain of inequalities hold: $\|\nabla\theta(\boldsymbol{x};t) - \nabla\theta(\boldsymbol{x};t-1)\| = \|G(\boldsymbol{x};t) - G(\boldsymbol{x};t-1)\| \leq \sum_{i\in\mathcal{I}}\|\nabla_{x_i}g_i(\boldsymbol{x};t) - \nabla_{x_i}g_i(\boldsymbol{x};t-1)\| \leq \sum_{i\in\mathcal{I}}\mathrm{e}_{\nabla_i} =: \mathrm{e}_\nabla$. ∎

Proof of Lemma 6: In view of item (S2) in Algorithm 1, at every outer iteration $t \in \mathbb{N}$ we have by equation (3) that $\boldsymbol{x}_t^\star$ satisfies $(\boldsymbol{y} - \boldsymbol{x}_t^\star)^\top F(\boldsymbol{x}_t^\star;t) \geq 0$ for all $\boldsymbol{y} \in \Omega$, and therefore, since $\boldsymbol{x}_{t-1}^\star \in \Omega$ as it is a v-GNE at $t-1$, $\Delta_t^{\star\top}F(\boldsymbol{x}_t^\star;t) \leq 0$. Thus, by adding and subtracting the term $\Delta_t^{\star\top}G(\boldsymbol{x}_{t-1}^\star;t-1)$, we obtain $\Delta_t^{\star\top}G(\boldsymbol{x}_{t-1}^\star;t-1) \leq \Delta_t^{\star\top}(G(\boldsymbol{x}_{t-1}^\star;t-1)-(G(\boldsymbol{x}_t^\star;t)+U(\boldsymbol{x}_t^\star;t)))$. With the personalized functionals in (4) and the perfect estimate of the pseudo-gradients, we have that $U(\boldsymbol{x}_t^\star;t) = c(t)(\Delta^\star + \alpha(t)G(\boldsymbol{x}_{t-1}^\star;t-1))$. Then, in view of Standing Assumption 2, it follows that $\alpha(t)\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1) \leq \Delta_t^{\star\top}(\nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla\theta(\boldsymbol{x}_t^\star;t) - \xi(t)\boldsymbol{x}_t^\star)$. Now, by summing and subtracting the term $\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_t^\star;t-1)$ in the RHS, we obtain

$$\alpha(t)\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1) \leq \Delta_t^{\star\top}(\nabla\theta(\boldsymbol{x}_t^\star;t) - \nabla\theta(\boldsymbol{x}_t^\star;t-1))$$
$$- \Delta_t^{\star\top}(\nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1)+c(t)\boldsymbol{x}_{t-1}^\star-\nabla\theta(\boldsymbol{x}_t^\star;t-1)-c(t)\boldsymbol{x}_t^\star).$$

As in the proof of Lemma 3, the second term in the RHS is always upper bounded by $-\ell\|\Delta_t^\star\|^2$ once introduced the auxiliary function $\psi(\boldsymbol{x};t) := \theta(\boldsymbol{x};t) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$ and chosen $c(t)$ in such a way that $\boldsymbol{x} \mapsto \psi(\boldsymbol{x};t)$ is ℓ-strongly convex, i.e., $c(t) \geq 2\ell$ for all $t \in \mathbb{N}$, in view of the ℓ-weak convexity of $\boldsymbol{x} \mapsto \theta(\boldsymbol{x};t)$. On the other hand, the first term attains its maximum positive module when the two vectors $\Delta_t^\star$ and $(\nabla\theta(\boldsymbol{x}_t^\star;t) - \nabla\theta(\boldsymbol{x}_t^\star;t-1))$ are aligned, thus leading to the following relation implied by Lemma 5:

$$\Delta_t^{\star\top}\nabla\theta(\boldsymbol{x}_{t-1}^\star;t-1) \leq -\tfrac{\ell}{\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|\mathrm{e}_\nabla, \tag{15}$$

where we have imposed that $\alpha(t) = 1 - c(t)\xi(t) > 0$, i.e., $\xi(t) < 1/c(t)$. Finally, adding and subtracting the term $(1/4\alpha(t)\ell)e_\nabla^2$ allows us to complete the square in the RHS, thus yielding to (7), and hence concluding the proof. ∎

Proof of Theorem 2: We start by combining the descent lemma [56, Prop. A.24], Lemma 5 and 6 (specifically, the relation in (15)). For all $t \in \mathbb{N}$, we obtain

$$\theta(\boldsymbol{x}_t^\star; t-1) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) + (\Delta_t^\star)^\top \nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) + \tfrac{\ell}{2}\|\Delta_t^\star\|^2 \tag{16}$$
$$\leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|e_\nabla.$$

Adding and subtracting the term $\theta(\boldsymbol{x}_t^\star; t)$ yields to

$$\theta(\boldsymbol{x}_t^\star; t) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) + \theta(\boldsymbol{x}_t^\star; t) - \theta(\boldsymbol{x}_t^\star; t-1) - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|e_\nabla$$

$$\leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) + |\theta(\boldsymbol{x}_t^\star; t) - \theta(\boldsymbol{x}_t^\star; t-1)| - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\boldsymbol{x}_t^\star - \boldsymbol{x}_{t-1}^\star\|^2 + \tfrac{1}{\alpha(t)}\|\boldsymbol{x}_t^\star - \boldsymbol{x}_{t-1}^\star\|e_\nabla$$

$$\leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|e_\nabla + e_\theta,$$

where the latter inequality follows in view of Assumption 2 i). By defining $\beta(t) := \ell(2 - \alpha(t))/2\alpha(t)$, which is strictly positive if $\xi(t) \in [0,1)$ for all $t \in \mathbb{N}$, and focusing on the second and third terms in the RHS of this latter inequality, we can complete the square by adding and subtracting $(4\alpha^2(t)\beta(t))^{-1}e_\nabla^2$. After few algebraic manipulations, we obtain the following inequality

$$\theta(\boldsymbol{x}_t^\star; t) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) - \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}e_\nabla\right)^2 + \tfrac{1}{4\alpha^2(t)\beta(t)}e_\nabla^2 + e_\theta. \tag{17}$$

From now on, we follow the proofline introduced in [9, Th. 3.1], [55, Th. 1], as in (13) we have obtained a typical structure characterizing the sample complexity analysis for sample-based stochastic algorithms (see, for instance, [9, Eq.(1.3)]). Specifically, by fixing some $T \in \mathbb{N}$ and summing up the inequality above over $t \in \mathcal{T} := \{1, \ldots, T\}$, for any global minimizer $\boldsymbol{x}^\star(T) \in \Theta(T)$, we obtain

$$\theta(\boldsymbol{x}^\star(T); T) \leq \theta(\boldsymbol{x}_T^\star; T) \leq \theta(\boldsymbol{x}_0^\star; 0) + e_\nabla^2 \sum_{t\in\mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)} + Te_\theta - \sum_{t\in\mathcal{T}} \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}e_\nabla\right)^2.$$

After moving $\theta(\boldsymbol{x}^\star(T); T)$ in the RHS and $\sum_{t\in\mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}e_\nabla)^2$ in the LHS, we obtain: $\sum_{t\in\mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}e_\nabla)^2 \leq |\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(T); T)| + e_\nabla^2 \sum_{t\in\mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)} + Te_\theta.$

Moreover, to upper bound the term $|\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(T); T)|$, which explicitly depends on the considered time horizon T, we exploit the following chain of inequalities:

$$|\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(T); T)| \leq |\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(0); 0)| + \sum_{t \in \mathcal{T}} |\theta(\boldsymbol{x}^\star(t); t) - \theta(\boldsymbol{x}^\star(t); t-1)|$$
$$+ \sum_{t \in \mathcal{T}} |\theta(\boldsymbol{x}^\star(t); t-1) - \theta(\boldsymbol{x}^\star(t-1); t-1)|,$$

where the inequality is obtained after adding and subtracting the terms $\sum_{t \in \mathcal{T}} (\theta(\boldsymbol{x}^\star(t); t-1) + \theta(\boldsymbol{x}^\star(t-1); t-1))$ in the absolute value, and upper bounding each resulting summation. In view of Assumption 2 i), the first summation in the RHS of the this latter inequality is upper bound by $T\mathrm{e}_\theta$, while the second one by $(\ell/2)T\mathrm{e}_\delta^2$. This follows by combining i) the ℓ-weak convexity of the mapping $\boldsymbol{x} \mapsto \theta(\boldsymbol{x}; t-1)$ (Lemma 1), ii) [9, Lemma 2.1] with the optimality condition $\nabla\theta(\boldsymbol{x}^\star(t-1); t-1) = 0$, and iii) the bound postulated in Assumption 2 iii). Then, by defining $\Delta_0 := |\theta(\boldsymbol{x}_0^\star; 0) - \theta(\boldsymbol{x}^\star(0); 0)|$, which does not depend on the length of the horizon T, we obtain, with $\phi := 2\mathrm{e}_\theta + \frac{\ell}{2}\mathrm{e}_\delta^2$,

$$\sum_{t \in \mathcal{T}} \beta(t) \left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla \right)^2 \leq \Delta_0 + T\phi + \mathrm{e}_\nabla^2 \sum_{t \in \mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)}. \tag{18}$$

Since $\beta(t) > 0$ for all $t \in \mathcal{T}$, we can leverage the Jensen's inequality [57, Th. 3.4] on the convex function $(\cdot)^2$ to lower bound the summation in the LHS. To this end, we normalize the coefficient 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 $\bar{\beta}(\sum_{t \in \mathcal{T}} \hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla)^2) \geq \bar{\beta}(\sum_{t \in \mathcal{T}} \hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla))^2$. Then, replacing this latter inequality into (18), $\bar{\beta}(\sum_{t \in \mathcal{T}} \hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla))^2 = \tfrac{1}{\bar{\beta}}(\sum_{t \in \mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla))^2 \leq \Delta_0 + T\phi + \mathrm{e}_\nabla^2 \sum_{t \in \mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)}$, and performing few algebraic manipulations directly yields to,

$$\sum_{t \in \mathcal{T}} \beta(t) \left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\mathrm{e}_\nabla \right) \leq \sqrt{\sum_{t \in \mathcal{T}} \left(\tfrac{\bar{\beta}}{4\alpha^2(t)\beta(t)}\mathrm{e}_\nabla^2 + \beta(t)(\Delta_0 + T\phi) \right)}.$$

By bringing the term $-\mathrm{e}_\nabla \sum_{t \in \mathcal{T}} (1/2\alpha(t))$ in the RHS, note that the obtained inequality still holds true if we premultiply both sides by $1/T$. Additionally, we have $(1/T)\sum_{t \in \mathcal{T}} \beta(t)\|\Delta_t^\star\| \geq (1/T)\underline{\beta}\sum_{t \in \mathcal{T}} \|\Delta_t^\star\|$, with $\underline{\beta} := \min_{t \in \mathcal{T}} \beta(t)$. This directly yields to the relation in (8). ∎

Proof of Lemma 7: By replicating the initial steps as in the proof of Lemma 6, due to Assumption 3 we have, for all $t \geq \bar{t}$,

$$\alpha(t){\Delta_t^\star}^\top \nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) \leq c(t)\xi(t){\Delta_t^\star}^\top \epsilon_{t-1} - {\Delta_t^\star}^\top (\nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla\theta(\boldsymbol{x}_t^\star; t) - c(t)\boldsymbol{x}_t^\star).$$

Summing and subtracting the term ${\Delta_t^\star}^\top \nabla\theta(\boldsymbol{x}_t^\star; t-1)$ in the RHS yields to $\alpha(t){\Delta_t^\star}^\top \nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) \leq {\Delta_t^\star}^\top (\nabla\theta(\boldsymbol{x}_t^\star; t) - \nabla\theta(\boldsymbol{x}_t^\star; t-1)) + c(t)\xi(t){\Delta_t^\star}^\top \epsilon_{t-1} - {\Delta_t^\star}^\top (\nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla\theta(\boldsymbol{x}_t^\star; t-1) - c(t)\boldsymbol{x}_t^\star)$. In case $c(t) \geq 2\ell > 0$ for all $t \in \mathbb{N}$, the third term in the RHS is always upper bound by $-\ell\|\Delta_t^\star\|^2$ in view of the ℓ-strongly convexity of the time-varying auxiliary function $\theta(\boldsymbol{x}; t) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$. Moreover, the first (resp., second) term attains its maximum positive module when the two vectors $\Delta_t^\star$ and $(\nabla\theta(\boldsymbol{x}_t^\star; t) - \nabla\theta(\boldsymbol{x}_t^\star; t-1))$ ($\Delta_t^\star$ and ϵ_{t-1}) are aligned, thus leading to the following relation

$$\alpha(t){\Delta_t^\star}^\top \nabla\theta(\boldsymbol{x}_{t-1}^\star) \leq -\ell\|\Delta_t^\star\|^2 + c(t)\xi(t)\|\Delta_t^\star\|\|\epsilon_{t-1}\| + \|\Delta_t^\star\|\|\nabla\theta(\boldsymbol{x}_t^\star; t) - \nabla\theta(\boldsymbol{x}_t^\star; t-1)\|,$$

once assumed that $\xi(t) \geq 0$. Following Lemma 6, the second term is upper bound by $\|\Delta_t^\star\|\mathrm{e}_\nabla$, while, in view of Assumption 3, for any $t \geq \bar{t}$ and $K \geq K_1$, it holds that $c(t)\xi(t)\|\epsilon_{t-1}\| \leq (1 - \alpha(t))\mathrm{e}(K-1)$ with probability $1 - \delta_1$. For all $t \geq \bar{t}$, imposing $\alpha(t) > 0$, i.e., $\xi(t) < 1/c(t)$, yields to

$$\tag{19} {\Delta_t^\star}^\top \nabla\theta(\boldsymbol{x}_{t-1}^\star; t-1) \leq -\tfrac{\ell}{\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|(\mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1)),$$

defining $\sigma(K, t) := \mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1)$ and completing the square in the RHS by adding and subtracting $(1/4\alpha(t)\ell)\sigma^2(K, t)$ yields to (9). ∎

Proof of Theorem 3: Also in this case, we combine the descent lemma [56, Prop. A.24] and (19) from Lemma 7. In view of Assumption 3, for all $t \geq \bar{t}$ and $K \geq K_1$, we obtain

$$\theta(\boldsymbol{x}_t^\star; t-1) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) - \ell\tfrac{2-\alpha(t)}{2\alpha(t)}\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|(\mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1)),$$

with probability $1 - \delta_1$. By introducing $\beta(t) := \ell(2 - \alpha(t))/2\alpha(t)$, as in the proof of Theorem 2 we first add and subtract the term $\theta(\boldsymbol{x}_t^\star; t)$ in the LHS to obtain $\theta(\boldsymbol{x}_t^\star; t) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) + |\theta(\boldsymbol{x}_t^\star; t) - \theta(\boldsymbol{x}_t^\star; t-1)| - \beta(t)\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|(\mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1)) \leq \theta(\boldsymbol{x}_{t-1}^\star; t-1) - \beta(t)\|\Delta_t^\star\|^2 + \tfrac{1}{\alpha(t)}\|\Delta_t^\star\|(\mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1)) + \mathrm{e}_\theta$. By focusing on the second and third terms in the RHS of this latter inequality, we can complete the square by adding and subtracting $(1/4\alpha^2(t)\beta(t))(\mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K-1))^2$, which is always possible if $\xi(t) \in [0, 1/c(t))$ for all

$t \in \mathbb{N}$. Therefore, it follows that also $\beta(t) > 0$, for all $t \in \mathbb{N}$ (and hence for $t \geq \bar{t}$). After few algebraic manipulations, by introducing $\sigma(K, t) := \mathrm{e}_\nabla + (1 - \alpha(t))\mathrm{e}(K - 1)$ we obtain what follows

$$\theta(\boldsymbol{x}_t^\star; t) \leq \theta(\boldsymbol{x}_{t-1}^\star; t - 1) + \frac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K, t) + \mathrm{e}_\theta - \beta(t)\left(\|\Delta_t^\star\| - \frac{1}{2\alpha(t)\beta(t)}\sigma(K, t)\right)^2. \quad (20)$$

From now on, we follow the proofline introduced in [9, Th. 3.1], [55, Th. 1], as in (20) we have obtained a typical structure characterizing the sample complexity analysis for sample-based stochastic algorithms (see, for instance, [9, Eq.(1.3)]). Specifically, by fixing some $T \in \mathbb{N}$ and summing up the inequality above over $t \in \mathcal{T} := \{\bar{t} + 1, \ldots, \bar{t} + T\}$, in view of the possibly sporadic communication between the agents and the central coordinator, let $q(t)$ be the subscript indicating the number of samples adopted by the learning procedure $\mathscr{L}$ at the t-th outer iteration. In particular, by assuming to initialize Algorithm 1 with $q(0) = K_1$ samples, $q(t) = q(t-1) + 1$ if, at the t-th iteration, the agents communicate p_t to the coordinator, $q(t) = q(t-1)$ otherwise. Note that the existence of some $K \in \mathbb{N}$ is guaranteed in view of Assumption 3. Then, for any global minimizer $\boldsymbol{x}^\star(T) \in \Theta(T)$, with probability $1 - \delta_1$ we obtain $\theta(\boldsymbol{x}^\star(T); T) \leq$ $\theta(\boldsymbol{x}_T^\star; T) \leq \theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \sum_{t \in \mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \frac{1}{2\alpha(t)\beta(t)}\sigma(K, t))^2 + \sum_{t \in \mathcal{T}} \frac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K, t) + T\mathrm{e}_\theta$. Then, by moving $-\sum_{t \in \mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \frac{1}{2\alpha(t)\beta(t)}\sigma(K, t))^2$ in the LHS and $\theta(\boldsymbol{x}^\star(T); T)$ in the RHS of this latter inequality, we obtain $\sum_{t \in \mathcal{T}} \beta(t)(\|\Delta_t^\star\| - \frac{1}{2\alpha(t)\beta(t)}\sigma(K, t))^2 \leq \theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(T); T) + \sum_{t \in \mathcal{T}} \frac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K, t) + T\mathrm{e}_\theta \leq |\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(T); T)| + \sum_{t \in \mathcal{T}} \frac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K, t) + T\mathrm{e}_\theta$. Moreover, to upper bound the term $|\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(T); T)|$, which explicitly depends on the horizon T, we exploit the following inequality obtained by adding and subtracting $\sum_{t \in \mathcal{T}}(\theta(\boldsymbol{x}^\star(t); t - 1) + \theta(\boldsymbol{x}^\star(t - 1); t - 1))$ in the absolute value:

$$|\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(T); T)| \leq |\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(\bar{t}); \bar{t})| + \sum_{t \in \mathcal{T}} |\theta(\boldsymbol{x}^\star(t); t) - \theta(\boldsymbol{x}^\star(t); t - 1)|$$

$$+ \sum_{t \in \mathcal{T}} |\theta(\boldsymbol{x}^\star(t); t - 1) - \theta(\boldsymbol{x}^\star(t - 1); t - 1)|.$$

In view of Assumption 2 i), the first summation in the RHS of the this latter inequality is upper bound by $T\mathrm{e}_\theta$, while the second one by $(\ell/2)T\mathrm{e}_\delta^2$. This follows by combining i) the ℓ-weak convexity of the mapping $\boldsymbol{x} \mapsto \theta(\boldsymbol{x}; t - 1)$ (Lemma 1), ii) [9, Lemma 2.1] with the optimality condition $\nabla\theta(\boldsymbol{x}^\star(t - 1); t - 1) = 0$, and iii) the bound postulated in Assumption 2 iii). Then, by

defining $\Delta_{\bar{t}} := |\theta(\boldsymbol{x}_{\bar{t}}^\star; \bar{t}) - \theta(\boldsymbol{x}^\star(\bar{t}); \bar{t})|$, which does not depend on the length of the horizon T, we obtain

$$\sum_{t\in\mathcal{T}} \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t)\right)^2 \leq \Delta_{\bar{t}} + T\phi + \sum_{t\in\mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K,t), \tag{21}$$

with $\phi := 2\mathrm{e}_\theta + \tfrac{\ell}{2}\mathrm{e}_\delta^2$. Since $\beta(t) > 0$ for all $t \in \mathcal{T}$, we can leverage the Jensen's inequality [57, Th. 3.4] on the convex function $(\cdot)^2$ to lower bound the summation in the LHS. To this end, we normalize the coefficient 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 $\bar{\beta}(\sum_{t\in\mathcal{T}} \hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t))^2) \geq \bar{\beta}(\sum_{t\in\mathcal{T}} \hat{\beta}(t)(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t)))^2$. Then, replacing this latter inequality into (21),

$$\bar{\beta}\left(\sum_{t\in\mathcal{T}} \hat{\beta}(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t)\right)\right)^2 = \tfrac{1}{\bar{\beta}}\left(\sum_{t\in\mathcal{T}} \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t)\right)\right)^2$$

$$\leq \Delta_{\bar{t}} + T\phi + \sum_{t\in\mathcal{T}} \tfrac{1}{4\alpha^2(t)\beta(t)}\sigma^2(K,t),$$

and performing few algebraic manipulations directly yields to,

$$\sum_{t\in\mathcal{T}} \beta(t)\left(\|\Delta_t^\star\| - \tfrac{1}{2\alpha(t)\beta(t)}\sigma(K,t)\right) \leq \sqrt{\sum_{t\in\mathcal{T}} \left(\beta(t)\left(\Delta_{\bar{t}} + T\phi\right) + \tfrac{\bar{\beta}}{4\alpha^2(t)\beta(t)}\sigma^2(K,t)\right)}.$$

By bringing the term $-\sum_{t\in\mathcal{T}}(1/2\alpha(t))\sigma(K,t)$ in the RHS, note that the obtained inequality still holds true if we premultiply both sides by $1/T$, thus obtaining the average over T in the LHS. Moreover, we have that $(1/T)\sum_{t\in\mathcal{T}} \beta(t)\|\Delta_t^\star\| \geq (1/T)\underline{\beta}\sum_{t\in\mathcal{T}} \|\Delta_t^\star\|$, with $\underline{\beta} := \min_{t\in\mathcal{T}}\beta(t)$, which directly yields to the relation in (10) with probability $1 - \delta_1$. $\blacksquare$

E. On learning the Lipschitz constant ℓ

The convergence results shown in the paper require one to know the constant of weak convexity $\ell := \sum_{i\in\mathcal{I}} \ell_i$ characterizing the unknown (dis)satisfaction function θ. Given the noncooperative nature of the problem considered, however, it seems reasonable that each Lipschitz constant ℓ_i represents a private information held by each agent. On the other hand, as long as the coordinator is endowed with a learning policy $\mathscr{L}$ to estimate $\nabla\theta$, one may wonder if it is also possible to learn the associated Lipschitz constant ℓ. Therefore, given $K \in \mathbb{N}$ agents' feedback, in this appendix we assume to have available some estimate $\hat{\ell}_K > 0$ of ℓ, and then we address this

privacy issue in the stationary setting (the time-varying one is similar and it is hence omitted in the interest of the paper length).

Lemma 8. *Let Assumption 1 hold true for some fixed $\delta_1 \in (0, 1]$, $c(t) \geq \gamma \hat{\ell}_K$, for some design parameter $\gamma > 0$, chosen so that $\gamma \hat{\ell}_{K-1} - \ell > 0$, and $K \geq K_1$, and $\xi(t) \in [0, 1/c(t))$ for all $t \in \mathbb{N}$. Then, for all $t \geq \bar{t}$, with the personalized incentives in* (4)*, we have*

$$\Delta_t^{\star \top} \nabla \theta(\boldsymbol{x}_{t-1}^\star) \leq \frac{\alpha(t)\kappa^2(t)}{\gamma \hat{\ell}_{K-1} - \ell} \mathrm{e}^2(K-1) - \left(\sqrt{\frac{\gamma \hat{\ell}_{K-1} - \ell}{\alpha(t)}} \|\Delta_t^\star\| - \kappa(t)\sqrt{\frac{\alpha(t)}{\gamma \hat{\ell}_{K-1} - \ell}} \mathrm{e}(K-1) \right)^2, \quad (22)$$

with probability $1 - \delta_1$. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

Proof. The first part of the proof replicates the one of Lemma 4, albeit in this case we can not rely on the ℓ-strongly convexity of the auxiliary function $\psi(\boldsymbol{x}; t) := \theta(\boldsymbol{x}) + \frac{c(t)}{2}\|\boldsymbol{x}\|^2$, since the Lipschitz constant ℓ is assumed to be unknown. In view of [9, Lemma 2.1], we have that $\Delta_t^{\star \top}(\nabla \theta(\boldsymbol{x}_{t-1}^\star) + c(t)\boldsymbol{x}_{t-1}^\star - \nabla \theta(\boldsymbol{x}_t^\star) - c(t)\boldsymbol{x}_t^\star) = \Delta_t^{\star \top}(\nabla \psi(\boldsymbol{x}_{t-1}^\star; t) - \nabla \psi(\boldsymbol{x}_t^\star; t)) \leq -(c(t) - \ell)\|\Delta_t^\star\|^2$. Then, by letting $c(t) \geq \gamma \hat{\ell}_{K-1}$ for all $t \geq \bar{t}$, $K \geq K_1$, relying on Assumption 1 and following the same reasoning for the two vectors $\Delta_t^\star$ and ϵ_{t-1} as in the proof of Lemma 4, directly yields

$$\Delta_t^{\star \top} \nabla \theta(\boldsymbol{x}_{t-1}^\star) \leq -\frac{\gamma \hat{\ell}_{K-1} - \ell}{\alpha(t)}\|\Delta_t^\star\|^2 + \frac{1 - \alpha(t)}{\alpha(t)}\|\Delta_t^\star\|\mathrm{e}(K-1),$$

with probability at least $1 - \delta_1$. Adding and subtracting $(\alpha(t)\kappa^2(t)/(\gamma \hat{\ell}_{K-1} - \ell))\mathrm{e}^2(K-1)$ to complete the square in the RHS, with γ such that $\gamma \hat{\ell}_{K-1} - \ell > 0$, leads to (22). $\qquad\blacksquare$

Two considerations are now in order: i) $\mathrm{e}^2(K-1)$ prevents $\Delta_t^\star$ from being a descent direction for θ, and ii) the condition $\gamma \hat{\ell}_{K-1} - \ell > 0$ requires one to implicitly overestimate the unknown ℓ. To overcome this latter issue, in the spirit of Assumption 1 we postulate the following condition:

Assumption 4. *For any $\delta_2 \in (0, 1]$, there exists some $K_2 < \infty$ such that $\mathbb{P}\{\hat{\ell}_K/\ell \in [\rho_{\min}, \rho_{\max}] \mid \forall K \geq K_2\} \geq 1 - \delta_2$, for some $0 < \rho_{\min} \leq 1 \leq \rho_{\max}$.* $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

With arbitrary probability $1 - \delta_2$, the ratio of $\hat{\ell}_K > 0$ and $\ell > 0$ is thus upper and lower bounded by some known terms, $\rho_{\min}$ and $\rho_{\max}$. As Assumptions 1 and 3, Assumption 4 is mild for a number of learning strategies, as LS [21].

Without restrictions, we henceforward set $\delta_1 = \delta_2 = \delta/2$, thus requiring to satisfy Assumption 1 and 4 with the same probability $1 - \delta/2$, and their intersection with probability $1 - \delta$ (via the union bound).

The following bound characterizes the sequence of v-GNE originated by Algorithm 1 in case the Lipschitz constant of $\nabla\theta$, ℓ, is not available a-priori.

Theorem 4. *Let Assumption 1–4 hold true for some fixed $\delta \in (0, 1]$, $\gamma \in (1/2\rho_{\min}, 1/\rho_{\min}]$, $c(t) \geq \gamma\hat{\ell}_{K-1}$ for some $K \geq \bar{K} := \max\{K_1, K_2\}$, and $\xi(t) \in [2(1 - \gamma\rho_{\min})/c(t), 1/c(t))$, for all $t \in \mathbb{N}$. Moreover, let some $T \in \mathbb{N}$ be fixed, $\mathcal{T} := \{\bar{t} + 1, \dots, T + \bar{t}\}$ and, for any global minimizer $\boldsymbol{x}^\star \in \Theta$, $\Delta_{\bar{t}} := \theta(\boldsymbol{x}_{\bar{t}}^\star) - \theta(\boldsymbol{x}^\star)$. Then, with the personalized incentives in (4), the sequence of v-GNE $(\boldsymbol{x}_t^\star)_{t \in \mathcal{T}}$, generated by Algorithm 1, satisfies the following relation*

$$\frac{1}{T}\sum_{t \in \mathcal{T}}\|\Delta_t^\star\| \leq \frac{1}{T\underline{\beta}}\sum_{t \in \mathcal{T}}\kappa(t)\mathrm{e}(q(t)) + \frac{1}{T\underline{\beta}}\sqrt{\sum_{t \in \mathcal{T}}\left(\beta_{q(t)}(t)\Delta_{\bar{t}} + \frac{\bar{\beta}\kappa^2(t)}{\beta_{q(t)}(t)}\mathrm{e}^2(q(t))\right)} \tag{23}$$

with probability $1 - \delta$, where $\beta_{q(t)}(t) := (2\gamma\hat{\ell}_{q(t)} - \ell(2 + \alpha(t)))/2\alpha(t)$, $\bar{\beta} := \sum_{t \in \mathcal{T}}\beta_{q(t)}(t)$, $\underline{\beta} := \min_{t \in \mathcal{T}}\beta_{q(t)}(t)$, and $q(t) \geq \bar{K}$ is the number of available agents' feedback at the t-th outer iteration, $t \in \mathcal{T}$. $\qquad\square$

Proof. The proof mimics the same steps of the one of Theorem 1. However, to recover the traditional structure characterizing the sample complexity analysis for sample-based stochastic algorithms as in, e.g., [9, Eq.(1.3)], we shall guarantee first that $\beta_{K-1}(t) := (2\gamma\hat{\ell}_{K-1} - \ell(2 + \alpha(t)))/2\alpha(t)$, which follows by combining the descent lemma and (22), is strictly positive, for all $t \geq \bar{t}$, $K \geq \bar{K} := \max\{K_1, K_2\}$. Therefore, $\beta_{K-1}(t) > 0$ if and only if $\alpha(t) < 2(\gamma(\hat{\ell}_{K-1}/\ell) - 1)$ and $\xi(t) < 1/c(t)$. Note that, due to Assumption 4, the former inequality is always satisfied if $\alpha(t) < 2(\gamma\rho_{\min} - 1)$, as $\hat{\ell}_{K-1}/\ell \geq \rho_{\min}$ with probability $1 - \delta$ for all $K \geq \bar{K} \geq K_2$. Thus, in view of the definition of $\alpha(t)$, we have that $\xi(t) > 2(1 - \gamma\rho_{\min})/c(t)$, for all $t \geq \bar{t}$, which amounts to a nonnegative lower bound in case $\gamma \leq 1/\rho_{\min}$, whereas it shall verify $\gamma > 1/2\rho_{\min}$ to meet $\xi(t) < 1/c(t)$. From now on, the proof is a verbatim copy of the one of Theorem 1 with $\beta_{K-1}(t)$ instead of $\beta(t)$, while in summing up over $\mathcal{T}$ we account for the sporadic communication among agents and central coordinator by considering $\beta_{q(t)}(t)$. $\qquad\blacksquare$

The considerations following Theorem 1, i.e., the parameter tuning for the personalized functionals, big-O analysis and learning strategy $\mathscr{L}$ apply also to this case *mutatis mutandis*.

REFERENCES

[1] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *4OR*, vol. 5, no. 3, pp. 173–210, 2007.

[2] F. Salehisadaghiani and L. Pavel, "Distributed Nash equilibrium seeking: A gossip-based algorithm," *Automatica*, vol. 72, pp. 209–216, 2016.

[3] F. Salehisadaghiani, W. Shi, and L. Pavel, "Distributed Nash equilibrium seeking under partial-decision information via the alternating direction method of multipliers," *Automatica*, vol. 103, pp. 27–35, 2019.

[4] 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.

[5] D. Gadjov and L. Pavel, "On the exact convergence to Nash equilibrium in hypomonotone regimes under full and partial-decision information," *IEEE Transactions on Automatic Control*, pp. 1–15, 2022.

[6] A. Jadbabaie, A. Rakhlin, S. Shahrampour, and K. Sridharan, "Online optimization: Competing with dynamic comparators," in *Artificial Intelligence and Statistics*. PMLR, 2015, pp. 398–406.

[7] I. Bogunovic, J. Scarlett, and V. Cevher, "Time-varying Gaussian process bandit optimization," in *Artificial Intelligence and Statistics*. PMLR, 2016, pp. 314–323.

[8] S. Shahrampour and A. Jadbabaie, "Distributed online optimization in dynamic environments using mirror descent," *IEEE Transactions on Automatic Control*, vol. 63, no. 3, pp. 714–725, 2017.

[9] 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.

[10] 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.

[11] A. Simonetto, E. Dall'Anese, J. Monteil, and A. Bernstein, "Personalized optimization with user's feedback," *Automatica*, vol. 131, p. 109767, 2021.

[12] E. Dall'Anese, A. Simonetto, S. Becker, and L. Madden, "Optimization and learning with information streams: Time-varying algorithms and applications," *IEEE Signal Processing Magazine*, vol. 37, no. 3, pp. 71–83, 2020.

[13] T. Ui, "A Shapley value representation of potential games," *Games and Economic Behavior*, vol. 31, no. 1, pp. 121–135, 2000.

[14] Q. D. Lã, Y. H. Chew, and B.-H. Soong, *Potential Game Theory*. Springer Science & Business Media, 2016.

[15] T. Heikkinen, "A potential game approach to distributed power control and scheduling," *Computer Networks*, vol. 50, no. 13, pp. 2295–2311, 2006.

[16] 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.

[17] 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.

[18] 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.

[19] B. F. Hobbs and J.-S. Pang, "Nash-Cournot equilibria in electric power markets with piecewise linear demand functions and joint constraints," *Operations Research*, vol. 55, no. 1, pp. 113–127, 2007.

[20] 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.

[21] I. Notarnicola, A. Simonetto, F. Farina, and G. Notarstefano, "Distributed personalized gradient tracking with convex parametric models," *IEEE Transactions on Automatic Control*, vol. 68, no. 1, pp. 588–595, 2023.

[22] A. M. Ospina, A. Simonetto, and E. Dall'Anese, "Time-varying optimization of networked systems with human preferences," *IEEE Transactions on Control of Network Systems*, vol. 10, no. 1, pp. 503–515, 2023.

[23] E. Tampubolon and H. Boche, "Coordinated online learning for multiagent systems with coupled constraints and perturbed utility observations," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5080–5095, 2021.

[24] ——, "On the convergence of online mirror ascent for aggregative games with approximated aggregates," in *2019 IEEE 20th International Workshop on Signal Processing Advances in Wireless Communications (SPAWC)*. IEEE, 2019, pp. 1–5.

[25] ——, "Robust pricing mechanism for resource sustainability under privacy constraint in competitive online learning multi-agent systems," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 8733–8737.

[26] P. Mertikopoulos and Z. Zhou, "Learning in games with continuous action sets and unknown payoff functions," *Mathematical Programming*, vol. 173, no. 1, pp. 465–507, 2019.

[27] Z. Zhou, P. Mertikopoulos, N. Bambos, P. Glynn, and C. Tomlin, "Multi-agent online learning with imperfect information," 2018, (Under review).

[28] J. Cohen, A. Héliou, and P. Mertikopoulos, "Learning with bandit feedback in potential games," in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, 2017, pp. 6372–6381.

[29] A. R. Cardoso, J. Abernethy, H. Wang, and H. Xu, "Competing against Nash equilibria in adversarially changing zero-sum games," in *International Conference on Machine Learning*. PMLR, 2019, pp. 921–930.

[30] B. Duvocelle, P. Mertikopoulos, M. Staudigl, and D. Vermeulen, "Multiagent online learning in time-varying games," *Mathematics of Operations Research*, 2022.

[31] F. Fabiani, A. Simonetto, and P. J. Goulart, "Learning equilibria with personalized incentives in a class of nonmonotone games," in *2022 European Control Conference (ECC)*. IEEE, 2022, pp. 2179–2184.

[32] 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.

[33] J. R. Marden and J. S. Shamma, "Game theory and control," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 1, pp. 105–134, 2018.

[34] D. Paccagnan, R. Chandan, and J. R. Marden, "Utility and mechanism design in multi-agent systems: An overview," *Annual Reviews in Control*, 2022.

[35] A. A. Kulkarni and U. V. Shanbhag, "An existence result for hierarchical Stackelberg v/s Stackelberg games," *IEEE Transactions on Automatic Control*, vol. 60, no. 12, pp. 3379–3384, 2015.

[36] F. Fabiani, M. A. Tajeddini, H. Kebriaei, and S. Grammatico, "Local Stackelberg equilibrium seeking in generalized aggregative games," *IEEE Transactions on Automatic Control*, vol. 67, no. 2, pp. 965–970, 2022.

[37] NYC OpenData. (2021) For-hire vehicle base aggregate report. [Online] https://opendata.cityofnewyork.us/.

[38] F. Facchinei and J. S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.

[39] 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.

[40] T. Roughgarden, "Algorithmic game theory," *Communications of the ACM*, vol. 53, no. 7, p. 78–86, jul 2010.

[41] R. W. Rosenthal, "A class of games possessing pure-strategy Nash equilibria," *International Journal of Game Theory*, vol. 2, no. 1, pp. 65–67, 1973.

[42] P. A. Kattuman, R. J. Green, and J. W. Bialek, "Allocating electricity transmission costs through tracing: A game-theoretic rationale," *Operations Research Letters*, vol. 32, no. 2, pp. 114–120, 2004.

[43] 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.

[44] 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.

[45] E. Cavazzuti, M. Pappalardo, and M. Passacantando, "Nash equilibria, variational inequalities, and dynamical systems," *Journal of Optimization Theory and Applications*, vol. 114, no. 3, pp. 491–506, 2002.

[46] G. Scutari, F. Facchinei, and L. Lampariello, "Parallel and distributed methods for constrained nonconvex optimization - Part I: Theory," *IEEE Transactions on Signal Processing*, vol. 65, no. 8, pp. 1929–1944, 2017.

[47] 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.

[48] A. Mokhtari, S. Shahrampour, A. Jadbabaie, and A. Ribeiro, "Online optimization in dynamic environments: Improved regret rates for strongly convex problems," in *IEEE Conference on Decision and Control*, 2016, pp. 7195–7201.

[49] Y. Li, G. Qu, and N. Li, "Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit," *IEEE Transactions on Automatic Control*, vol. 66, no. 10, pp. 4761–4768, 2021.

[50] G. Mateos and G. B. Giannakis, "Distributed recursive least-squares for consensus-based in-network adaptive estimation," *IEEE Transactions on Signal Processing*, pp. 4583–4588, 2009.

[51] V. Pandey, J. Monteil, C. Gambella, and A. Simonetto, "On the needs for MaaS platforms to handle competition in ridesharing mobility," *Transportation Research Part C: Emerging Technologies*, vol. 108, pp. 269–288, 2019.

[52] M. Diao, H. Kong, and J. Zhao, "Impacts of transportation network companies on urban mobility," *Nature Sustainability*, vol. 4, no. 6, pp. 494–500, 2021.

[53] F. Fabiani and B. Franci, "A stochastic generalized Nash equilibrium model for platforms competition in the ride-hail market," in *2022 IEEE 61st Conference on Decision and Control (CDC)*, 2022, pp. 4455–4460.

[54] M. V. Solodov and P. Tseng, "Modified projection-type methods for monotone variational inequalities," *SIAM Journal on Control and Optimization*, vol. 34, no. 5, pp. 1814–1830, 1996.

[55] V. Mai and M. Johansson, "Convergence of a stochastic gradient method with momentum for non-smooth non-convex optimization," in *International Conference on Machine Learning*. PMLR, 2020, pp. 6630–6639.

[56] D. P. Bertsekas, "Nonlinear programming," *Journal of the Operational Research Society*, vol. 48, no. 3, pp. 334–334, 1997.

[57] R. T. Rockafellar, *Convex analysis*. Princeton university press, 1970, vol. 36.