

A distributed forward-backward algorithm for stochastic generalized Nash equilibrium seeking

Franci, Barbara; Grammatico, Sergio

DOI

[10.1109/TAC.2020.3047369](https://doi.org/10.1109/TAC.2020.3047369)

Publication date

2021

Document Version

Accepted author manuscript

Published in

IEEE Transactions on Automatic Control

Citation (APA)

Franci, B., & Grammatico, S. (2021). A distributed forward-backward algorithm for stochastic generalized Nash equilibrium seeking. *IEEE Transactions on Automatic Control*, 66(11), 5467-5473.
<https://doi.org/10.1109/TAC.2020.3047369>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

A distributed forward–backward algorithm for stochastic generalized Nash equilibrium seeking

Barbara Franci and Sergio Grammatico

Abstract—We consider the stochastic generalized Nash equilibrium problem (SGNEP) with expected-value cost functions. Inspired by Yi and Pavel (Automatica, 2019), we propose a distributed GNE seeking algorithm based on the preconditioned forward–backward operator splitting for SGNEP, where, at each iteration, the expected value of the pseudo-gradient is approximated via a number of random samples. Our main contribution is to show almost sure convergence of the proposed algorithm if the pseudo-gradient mapping is restricted (monotone and) cocoercive.

Index Terms—Stochastic generalized Nash equilibrium problems, variational inequalities, stochastic approximation.

I. INTRODUCTION

Generalized Nash equilibrium problems (GNEPs) have recently received strong attention from the multi-agent system and control community [1]–[4]. One reason for this interest is related to the possible applications that range from economics to engineering and operation research [2], [5]. In GNEPs, each agent seeks to minimize his own cost function under some *joint* feasibility constraints. Namely, both the cost function and the constraints depend on the strategies chosen by the other agents. Consequently, the search for a GNE is usually very difficult. A number of results are available concerning algorithms and methodologies to solve a GNEP [6]. In the deterministic case, many algorithms are available to find an equilibrium, both distributed or semi-decentralized [1], [7], [8]. Among the methodologies to reach an equilibrium, an effective approach is to seek for a solution of the associated variational inequality (VI) [6].

To recast a GNEP as a VI, the Karush-Kuhn-Tucker conditions can be considered to rewrite the problem as a monotone inclusion. The latter problem can then be solved via operator splitting techniques. Among others, we focus on the forward–backward (FB) splitting which leads to one of the simplest and computationally inexpensive algorithms available [9].

The downside of the FB scheme is that, when directly applied to GNEPs, it cannot be distributed. On the other hand, when considering a game-theoretic setup, it is desirable to consider distributed algorithms, in the sense that each agent should only know its local cost function and its local constraints. For this reason, *preconditioning* has been recently introduced in [1], see [10] for a preliminary extension of this method to the stochastic case.

A stochastic GNEP (SGNEP) is a GNEP where the cost functions are expected value functions [11]. Such problems arise when there is some uncertainty, expressed through a random variable with unknown distribution. SGNEPs are not studied as much as their deterministic counterpart, despite many practical problems must be modelled with uncertainty. Among others, in transportation systems, a possible source of uncertainty is the drivers perception of travel time [12]; in electricity markets, companies produce energy without knowing in advance the demand [13]. Moreover, any network Cournot game with

market capacity constraints and uncertain demand can be modelled as a SGNEP [14], [15]. Due to their wide applicability, SGNEPs have been studied by the control community as well [16]–[18].

Unfortunately, the pseudo-gradient is usually not directly accessible, for instance due to the need of excessive computations in performing the expected value. For this reason, in many situations, the search for an equilibrium relies on samples of the random variable. Essentially, there are two main methodologies available: sample average approximation (SAA) and stochastic approximation (SA). In the SAA approach, one replaces the expected value formulation with the average over an infinite number of samples of the random variable. In the SA scheme, each agent can sample only one realization of the random variable. This approach is less computationally expensive, but, not surprisingly, it usually requires stronger assumptions on the problem data [17], [19]. Alternatively, the variance reduced SA scheme considers the average over an increasing number of samples, which is possible when there is a huge amount of data available as in Monte Carlo simulations or machine learning [20].

One of the first formulations of a stochastic FB algorithm is in [21], under the assumption of strong monotonicity and Lipschitz continuity of the mapping involved. In [22] instead, convergence is proved under cocoercivity and uniform monotonicity. While [21] considers the SA scheme, the algorithm in [22] is independent on the chosen approximation scheme. To weaken the assumptions, algorithms more involved than the FB have been proposed in the literature. For instance, in a recent paper, [23], the authors propose a forward-backward-forward (FBF) algorithm that converges to a solution under the assumption of pseudomonotone pseudo-gradient mapping but it requires two costly evaluations of the pseudo-gradient mapping. Alternatively, under the same assumptions, one can consider the extragradient (EG) method proposed in [20] which takes two projection steps that can be slow. Therefore, taking weaker assumptions comes at the price of increasing computational complexity and slowness of the algorithms.

In this paper, we present the first preconditioned distributed FB algorithm for SGNEPs and we prove almost sure convergence under restricted cocoercivity of the pseudo-gradient mapping via the SA scheme with variance reduction. Our technical assumptions are weaker when compared to the current literature on FB algorithms, e.g., cocoercivity and uniform monotonicity [22] or even strong monotonicity [21], [24]. Moreover, compared to a direct application of the FBF and EG algorithms, ours shows faster convergence in terms of number of iterations and computational time.

A preliminary study related to this work is presented in [10]. In that paper, we consider a SGNEP and build a preconditioned FB algorithm with damping. The algorithm is guaranteed to reach a SGNE if the pseudo-gradient mapping is strongly monotone and its convergence follows directly from [22]. Here, we show that the assumption of uniform monotonicity taken in [22] can be dropped and that restricted cocoercivity is enough for the analysis and to ensure convergence.

Notation and preliminaries: $\mathbb{R}$ denotes the set of real numbers and $\mathbb{R} = \mathbb{R} \cup \{+\infty\}$. $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ denotes the standard inner product and $\|\cdot\|$ represents the associated Euclidean norm. We indicate a (symmetric and) positive definite matrix A , i.e., $x^T A x > 0$, with $A \succ 0$. Given a matrix $\Phi \succ 0$, we denote the Φ -induced inner product, $\langle x, y \rangle_\Phi = \langle \Phi x, y \rangle$. The associated Φ -induced norm, $\|\cdot\|_\Phi$,

The authors are in the Delft Center for System and Control, TU Delft, The Netherlands. E-mail addresses: {b.franci-1, s.grammatico}@tudelft.nl.

This work was partially supported by NWO under research projects OMEGA (613.001.702) and P2P-TALES (647.003.003), and by the ERC under research project COSMOS (802348).

is defined as $\|x\|_\Phi = \sqrt{\langle \Phi x, x \rangle}$. $A \otimes B$ indicates the Kronecker product between matrices A and B . $\mathbf{0}_m$ indicates the vector with m entries all equal to 0. Given N vectors $x_1, \dots, x_N \in \mathbb{R}^n$, $\mathbf{x} := \text{col}(x_1, \dots, x_N) = \begin{bmatrix} x_1^\top, \dots, x_N^\top \end{bmatrix}^\top$.

$J_F = (\text{Id} + F)^{-1}$ is the resolvent of the operator $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ where Id indicates the identity operator. The set of fixed points of F is $\text{fix}(F) := \{x \in \mathbb{R}^n \mid x \in F(x)\}$. For a closed set $C \subseteq \mathbb{R}^n$, the mapping $\text{proj}_C : \mathbb{R}^n \rightarrow C$ denotes the projection onto C , i.e., $\text{proj}_C(x) = \arg\min_{y \in C} \|y - x\|$. The residual mapping is, in general, defined as $\text{res}(x^k) = \|x^k - \text{proj}_C(x^k - F(x^k))\|$. Given a proper, lower semi-continuous, convex function g , the subdifferential is the operator $\partial g(x) := \{u \in \Omega : \forall y \in \Omega, \langle y - x, u \rangle + g(x) \leq g(y)\}$. The proximal operator is defined as $\text{prox}_g(v) := \arg\min_{u \in \Omega} \{g(u) + \frac{1}{2}\|u - v\|^2\} = J_{\partial g}(v)$. ι_C is the indicator function of C , i.e., $\iota_C(x) = 1$ if $x \in C$ and $\iota_C(x) = +\infty$ otherwise. The set-valued mapping $N_C : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denotes the normal cone operator for the set C , i.e., $N_C(x) = \emptyset$ if $x \notin C$, $\{v \in \mathbb{R}^n \mid \sup_{z \in C} v^\top(z - x) \leq 0\}$ otherwise.

We now recall some basic properties of operators [9]. A mapping $F : \text{dom } F \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$ is: ℓ -Lipschitz continuous if, for some $\ell > 0$, $\|F(x) - F(y)\| \leq \ell\|x - y\|$ for all $x, y \in \mathbb{R}^n$; (strictly) monotone if for all $x, y \in \text{dom}(F)$ ($x \neq y$), $\langle F(x) - F(y), x - y \rangle \geq (>) 0$; η -strongly monotone if, for $\eta > 0$, $\langle F(x) - F(y), x - y \rangle \geq \eta\|x - y\|^2$ for all $x, y \in \text{dom}(F)$; β -cocoercive with $\beta > 0$, if for all $x, y \in \text{dom}(F)$, $\langle F(x) - F(y), x - y \rangle \geq \beta\|F(x) - F(y)\|^2$; firmly non expansive if for all $x, y \in \text{dom}(F)$, $\|F(x) - F(y)\|^2 \leq \|x - y\|^2 - \|(\text{Id} - F)(x) - (\text{Id} - F)(y)\|^2$. We use the adjective “restricted” if a property holds for all $(x, y) \in \text{dom}(F) \times \text{fix}(F)$. We note that a strongly monotone and Lipschitz continuous mapping is also cocoercive and that a firmly nonexpansive operator is also cocoercive, hence monotone [9, Def. 4.1].

II. STOCHASTIC GENERALIZED NASH EQUILIBRIUM PROBLEMS

In this section we describe the stochastic generalized Nash equilibrium problem (SGNEP), i.e., a collection of optimizations problem where the cost functions are expected value functions and the agents are subject to coupling constraints.

We consider a set $\mathcal{I} = \{1, \dots, N\}$ of self-interested agents, each of them choosing its strategy $x_i \in \mathbb{R}^{n_i}$ from its local decision set $\Omega_i \subseteq \mathbb{R}^{n_i}$. We call $\mathbf{x}_{-i} = \text{col}((x_j)_{j \neq i})$ the decisions of all the agents with the exception of i and set $n = \sum_{i=1}^N n_i$. The aim of each agent is to minimize its local cost function within its feasible strategy set. The local cost function of agent i is defined as

$$\mathbb{J}_i(x_i, \mathbf{x}_{-i}) := \mathbb{E}_{\xi_i} [f_i(x_i, \mathbf{x}_{-i}, \xi_i(\omega))] + g_i(x_i), \quad (1)$$

for some measurable function $f_i : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$. The cost function presents the typical splitting in smooth and non-smooth parts where the latter is represented by the function $g_i : \mathbb{R}^{n_i} \rightarrow \bar{\mathbb{R}}$ which can model not only a local cost, but also local constraints via an additive indicator function, e.g., $g_i(x_i) = \tilde{g}_i(x_i) + \iota_{\Omega_i}(x_i)$.

Assumption 1 (Local cost): For each $i \in \mathcal{I}$, the function g_i in (1) is lower semicontinuous and convex, and $\text{dom}(g_i) = \Omega_i \subseteq \mathbb{R}^{n_i}$ is nonempty, compact and convex. $\square$

The uncertainty in the cost function $\mathbb{J}_i$ in (1) is expressed through the random variable $\xi_i : \Xi \rightarrow \mathbb{R}^d$ where $(\Xi, \mathcal{F}, \mathbb{P})$ is the probability space and $\Xi = \Xi_1 \times \dots \times \Xi_N$. Therefore, the cost function depends on the local variable x_i , on the decision of the other agents $\mathbf{x}_{-i}$ and on the random variable $\xi_i(\omega)$. $\mathbb{E}_\xi$ represent the mathematical expectation with respect to the distribution of the random variable

$\xi^\dagger$. We assume that $\mathbb{E}[f_i(x, \xi_i)]$ is well defined for all the feasible $\mathbf{x} \in \mathcal{X}$.

Assumption 2 (Cost functions convexity): For each $i \in \mathcal{I}$ and $\mathbf{x}_{-i} \in \mathcal{X}_{-i}$ the function $\mathbb{J}_i(\cdot, \mathbf{x}_{-i})$ is convex and continuously differentiable. $\square$

Furthermore, we consider games with affine shared constraints $A\mathbf{x} \leq b$. Therefore, the feasible decision set of each agent $i \in \mathcal{I}$ is denoted by the set-valued mapping:

$$\mathcal{X}_i(\mathbf{x}_{-i}) := \left\{ y_i \in \Omega_i \mid A_i y_i \leq b - \sum_{j \neq i}^N A_j x_j \right\}, \quad (2)$$

where $A_i \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. The set Ω_i represents the local decision set for agent i , while the matrix A_i defines how agent i is involved in the coupling constraints. The collective feasible set can be written as

$$\mathcal{X} = \{y \in \Omega \mid Ay - b \leq \mathbf{0}_m\} \quad (3)$$

where $\Omega = \prod_{i=1}^N \Omega_i$ and $A = [A_1, \dots, A_N] \in \mathbb{R}^{m \times n}$. We suppose that there is no uncertainty in the constraints.

Assumption 3 (Constraint qualification): The set $\mathcal{X}$ satisfies Slater's constraint qualification. $\square$

The aim of each agent i , given the decision variables of the other agents $\mathbf{x}_{-i}$, is to choose a strategy x_i , that solves its local optimization problem, i.e.,

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in \Omega_i} & \mathbb{J}_i(x_i, \mathbf{x}_{-i}) \\ \text{s.t.} & A_i x_i \leq b - \sum_{j \neq i}^N A_j x_j. \end{cases} \quad (4)$$

When the optimization problems in (4) are simultaneously solved, the solution concept that we are seeking is that of stochastic generalized Nash equilibrium (SGNE).

Definition 1: A stochastic generalized Nash equilibrium is a collective strategy $\mathbf{x}^* \in \mathcal{X}$ such that, for all $i \in \mathcal{I}$,

$$\mathbb{J}_i(x_i^*, \mathbf{x}_{-i}^*) \leq \inf \{ \mathbb{J}_i(y, \mathbf{x}_{-i}^*) \mid y \in \mathcal{X}_i(\mathbf{x}_{-i}^*) \}. \quad \square$$

In other words, a SGNE is a set of strategies where no agent can decrease its objective function by unilaterally deviating from its decision. To guarantee the existence of a SGNE, we make further assumptions on the cost function.

Assumption 4 (Cost functions measurability): For each $i \in \mathcal{I}$ and for each $\xi \in \Xi$, the function $f_i(\cdot, \mathbf{x}_{-i}, \xi)$ is convex, Lipschitz continuous, and continuously differentiable. The function $f_i(x_i, \mathbf{x}_{-i}, \cdot)$ is measurable and for each $\mathbf{x}_{-i}$, the Lipschitz constant $\ell_i(\mathbf{x}_{-i}, \xi)$ is integrable in ξ . $\square$

While, under Assumptions 1–4, existence of SGNE of the game is guaranteed by [16, §3.1], uniqueness does not hold in general [16, §3.2].

Within all the possible Nash equilibria, we focus on those that corresponds to the solution set of an appropriate variational inequality. To this aim, let us define the (pseudo) gradient mapping as

$$\mathbb{F}(\mathbf{x}) = \text{col}((\mathbb{E}[\nabla_{x_i} f_i(x_i, \mathbf{x}_{-i}, \xi_i)])_{i \in \mathcal{I}}), \quad (5)$$

where the possibility to exchange the expected value and the gradient is guaranteed by Assumption 4. Then, the associated SVI reads as

$$\langle \mathbb{F}(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle + \sum_{i \in \mathcal{I}} \{g_i(x_i) - g_i(x_i^*)\} \geq 0, \forall \mathbf{x} \in \mathcal{X}. \quad (6)$$

Remark 1: Under Assumptions 1–4, the solution set of SVI($\mathcal{X}, \mathbb{F}$) is non empty and compact, i.e., $\text{SOL}(\mathcal{X}, \mathbb{F}) \neq \emptyset$ [25, Cor. 2.2.5].

When Assumptions 1–4 hold, any solution of SVI($\mathcal{X}, \mathbb{F}$) in (6) is a SGNE of the game in (4), while vice versa does not hold in general. Indeed, there may be Nash equilibria that are not solution of the SVI [26, Prop. 12.7]. The SGNE that are also solution of the associated SVI are called variational equilibria (v-SGNE). Therefore,

[†]From now on, we use ξ instead of $\xi(\omega)$ and $\mathbb{E}$ instead of $\mathbb{E}_\xi$.

Algorithm 1 Preconditioned Stochastic Forward-BackwardInitialization: $x_i^0 \in \Omega_i$, $\lambda_i^0 \in \mathbb{R}_{\geq 0}^m$, and $z_i^0 \in \mathbb{R}^m$.Iteration k : Agent i (1) Receives x_j^k for all $j \in \mathcal{N}_i^f$, λ_j^k for $j \in \mathcal{N}_i^\lambda$ then updates:

$$\begin{aligned} x_i^{k+1} &= \text{prox}_{g_i} [x_i^k - \alpha_i (\hat{F}_i(x_i^k, \mathbf{x}_{-i}^k, \xi_i^k) - A_i^\top \lambda_i^k)] \\ z_i^{k+1} &= z_i^k - \nu_i \sum_{j \in \mathcal{N}_i^\lambda} w_{i,j} (\lambda_i^k - \lambda_j^k) \end{aligned}$$

(2) Receives $z_{j,k+1}$ for all $j \in \mathcal{N}_i^\lambda$ then updates:

$$\begin{aligned} \lambda_i^{k+1} &= \text{proj}_{\mathbb{R}_+^m} \left[\lambda_i^k + \sigma_i \left(A_i (2x_i^{k+1} - x_i^k) - b_i \right) \right. \\ &\quad \left. + \sigma_i \sum_{j \in \mathcal{N}_i^\lambda} w_{i,j} \left(2(z_i^{k+1} - z_j^{k+1}) - (z_i^k - z_j^k) \right) \right. \\ &\quad \left. - \sigma_i \sum_{j \in \mathcal{N}_i^\lambda} w_{i,j} (\lambda_i^k - \lambda_j^k) \right] \end{aligned}$$

a v-SGNE of the game in (4), is the solution of the SVI($\mathcal{X}, \mathbb{F}$) in (6) where $\mathbb{F}$ is described in (5) and $\mathcal{X}$ is defined in (3).

Now, we recast the SGNEP as a monotone inclusion. For each agent $i \in \mathcal{I}$, let the Lagrangian function be $\mathcal{L}_i(\mathbf{x}, \lambda_i) := \mathbb{J}_i(x_i, \mathbf{x}_{-i}) + g_i(x_i) + \lambda_i^\top (A\mathbf{x} - b)$, where $\lambda_i \in \mathbb{R}_{\geq 0}^m$ is the dual variable associated with the coupling constraints. The collective decision $\mathbf{x}^*$ is a v-SGNE of the game in (4) if and only if the following Karush-Kuhn-Tucker (KKT) conditions are satisfied:

$$\forall i \in \mathcal{I} : \begin{cases} 0 \in \mathbb{E}[\nabla_{x_i} f_i(x_i^*, \mathbf{x}_{-i}^*, \xi_i)] + N_{\Omega_i}(x_i^*) + A_i^\top \lambda_i, \\ 0 \in (A\mathbf{x}^* - b) + N_{\mathbb{R}_{\geq 0}^m}(\lambda^*). \end{cases} \quad (7)$$

In (7), $\lambda_i = \lambda$ for all $i \in \mathcal{I}$, namely, all the agents agree on the same dual variable, see [27, Thm. 3.1], [28, Thm. 3.1].

III. STOCHASTIC PRECONDITIONED FORWARD-BACKWARD ALGORITHM

In the following sections, we design a stochastic counterpart of [1, Alg. 1] (Algorithm 1). First, we describe the preconditioning procedure that leads to the distributed iterations presented in Algorithm 1. Therein, x_i^k , z_i^k and λ_i^k are the state variables of agent i at iteration k , while ξ_i^k is a vector of i.i.d. random variables that agent i may use to approximate its pseudo-gradient mapping with a random approximation $\hat{F}_i$, defined later in this section.

Remark 2: If the local cost function g_i is the indicator function, we can use the projection on the local feasible set Ω_i , instead of the proximal operator [9, Ex. 12.25]. $\square$

We suppose that each agent i only knows its local data, i.e., Ω_i , A_i and b_i . Moreover, each agent has access to a pool of samples of the random variable and is able to compute, given the strategies of the other agents $\mathbf{x}_{-i}$, $\mathbb{E}[\nabla_{x_i} f_i(x_i, \mathbf{x}_{-i}, \xi)]$ (or an approximation, as introduced later in this section). Since the cost function is affected by the other agents strategies, we call $\mathcal{N}_i^f$ the set of agents interacting with i . Specifically, $j \in \mathcal{N}_i^f$ if the function $f_i(x_i, \mathbf{x}_{-i})$ explicitly depends on x_j .

We let the local copy of the dual variable be shared through the dual variables graph, $\mathcal{G}^\lambda = (\mathcal{I}, \mathcal{E}^\lambda)$. Along with the dual variable, the agents share on $\mathcal{G}^\lambda$ a copy of the auxiliary variables $z_i \in \mathbb{R}^m$. The role of $\mathbf{z} = \text{col}(z_1, \dots, z_N)$ is to force consensus, since this is the configuration that we are seeking. The set of edges $\mathcal{E}^\lambda$ is given by: $(i, j) \in \mathcal{E}^\lambda$ if player i can receive $\{\lambda_j, z_j\}$ from player j . The neighbouring agents in $\mathcal{G}^\lambda$ form a set $\mathcal{N}_i^\lambda = \{j \in \mathcal{I} :$

$(i, j) \in \mathcal{E}^\lambda\}$ for all $i \in \mathcal{I}$. In this way, each agent control his own decision variable, a local copy of the dual variable and of the auxiliary variable, and has access to the other agents variables through the graphs.

Assumption 5 (Graph connectivity): The dual-variable graph $\mathcal{G}^\lambda$ is undirected and connected. $\square$

We call $W \in \mathbb{R}^{N \times N}$ the weighted adjacency matrix of $\mathcal{G}^\lambda$. Then, by letting $d_i = \sum_{j=1}^N w_{i,j}$ and $D = \text{diag}\{d_1, \dots, d_N\}$, the associated Laplacian is the matrix $L = D - W$. Moreover, it follows from Assumption 5 that $L = L^\top$.

Rewriting the KKT conditions in (7) in compact form as

$$0 \in \mathcal{T}(\mathbf{x}, \boldsymbol{\lambda}) := \begin{bmatrix} G(\mathbf{x}) + \mathbb{F}(\mathbf{x}) + A^\top \boldsymbol{\lambda} \\ N_{\mathbb{R}_{\geq 0}^m}(\boldsymbol{\lambda}) - (A\mathbf{x} - b) \end{bmatrix} \quad (8)$$

where $\mathcal{T} : \mathcal{X} \times \mathbb{R}_{\geq 0}^m \rightrightarrows \mathbb{R}^n \times \mathbb{R}^m$ is a set-valued mapping and $G(\mathbf{x}) = \text{col}(\partial g_1(x_1), \dots, \partial g_N(x_N))$, it follows that the v-SGNEs correspond to the zeros of the mapping $\mathcal{T}$.

In the remaining part of this section, we split $\mathcal{T}$ into the sum of two operators $\mathcal{A}$ and $\mathcal{B}$ that satisfy specific properties. The advantage of this technique is that the zeros of the mapping $\mathcal{A} + \mathcal{B}$ correspond to the fixed point of a specific operator depending on both $\mathcal{A}$ and $\mathcal{B}$, as exploited in [1], [7]. Such a scheme is known as forward backward (FB) splitting [9, §26.5]. In fact, it holds that, for any matrix $\Phi \succ 0$, $\boldsymbol{\omega} \in \text{zer}(\mathcal{A} + \mathcal{B})$ if and only if

$$\boldsymbol{\omega} = (\text{Id} + \Phi^{-1}\mathcal{B})^{-1} \circ (\text{Id} - \Phi^{-1}\mathcal{A})(\boldsymbol{\omega}).$$

Specifically, the operator $\mathcal{T}$ in (8) can be written as a summation of the two operators

$$\begin{aligned} \mathcal{A} : \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{bmatrix} &\mapsto \begin{bmatrix} \mathbb{F}(\mathbf{x}) \\ b \end{bmatrix} \\ \mathcal{B} : \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{bmatrix} &\mapsto \begin{bmatrix} G(\mathbf{x}) \\ N_{\mathbb{R}_{\geq 0}^m}(\boldsymbol{\lambda}) \end{bmatrix} + \begin{bmatrix} 0 & A^\top \\ -A & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \boldsymbol{\lambda} \end{bmatrix}. \end{aligned} \quad (9)$$

Therefore, finding a v-SGNE translates in finding a pair $(\mathbf{x}^*, \boldsymbol{\lambda}^*) \in \mathcal{X} \times \mathbb{R}_{\geq 0}^m$ such that $(\mathbf{x}^*, \boldsymbol{\lambda}^*) \in \text{zer}(\mathcal{A} + \mathcal{B})$.

To impose consensus on the dual variables, the authors in [1] proposed the Laplacian constraint $L\boldsymbol{\lambda} = 0$. This is why, to preserve monotonicity we expand the two operators $\mathcal{A}$ and $\mathcal{B}$ in (9) introducing the auxiliary variable $\mathbf{z} \in \mathbb{R}^{Nm}$. Let $L \in \mathbb{R}^{N \times N}$ be the Laplacian of $\mathcal{G}^\lambda$ and set $\mathbf{L} = L \otimes \text{Id}_m \in \mathbb{R}^{Nm \times Nm}$. Let us define $\mathbf{A} = \text{diag}\{A_1, \dots, A_N\} \in \mathbb{R}^{Nm \times n}$ and $\boldsymbol{\lambda} = \text{col}(\lambda_1, \dots, \lambda_N) \in \mathbb{R}^{Nm}$; similarly we define $\mathbf{b}$ of suitable dimension. Then, let us define the extended operators

$$\begin{aligned} \bar{\mathcal{A}} : \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \\ \boldsymbol{\lambda} \end{bmatrix} &\mapsto \begin{bmatrix} \mathbb{F}(\mathbf{x}) \\ 0 \\ \mathbf{b} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \mathbf{L}\boldsymbol{\lambda} \end{bmatrix} \\ \bar{\mathcal{B}} : \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \\ \boldsymbol{\lambda} \end{bmatrix} &\mapsto \begin{bmatrix} G(\mathbf{x}) \\ \mathbf{0} \\ N_{\mathbb{R}_{\geq 0}^m}(\boldsymbol{\lambda}) \end{bmatrix} + \begin{bmatrix} 0 & 0 & \mathbf{A}^\top \\ 0 & 0 & \mathbf{L} \\ -\mathbf{A} & -\mathbf{L} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{z} \\ \boldsymbol{\lambda} \end{bmatrix}. \end{aligned} \quad (10)$$

To ensure that the zeros of $\bar{\mathcal{A}} + \bar{\mathcal{B}}$ correspond to the zeros of the operator $\mathcal{T}$ in (8), we show that the following result holds.

Lemma 1: Let Assumptions 1-5 hold and consider the operators $\mathcal{A}$ and $\mathcal{B}$, $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$ in (9) and (10) respectively. Then, the following hold.

- (i) For any $\text{col}(\mathbf{x}^*, \mathbf{z}^*, \boldsymbol{\lambda}^*) \in \text{zer}(\bar{\mathcal{A}} + \bar{\mathcal{B}})$, $\mathbf{x}^*$ is a v-SGNE of game in (4), i.e., $\mathbf{x}^*$ solves the SVI($\mathcal{X}, \mathbb{F}$) in (6). Moreover $\boldsymbol{\lambda}^* = \mathbf{1}_N \otimes \lambda^*$, and $(\mathbf{x}^*, \boldsymbol{\lambda}^*)$ satisfy the KKT condition in (7) i.e., $\text{col}(\mathbf{x}^*, \boldsymbol{\lambda}^*) \in \text{zer}(\mathcal{A} + \mathcal{B})$;
- (ii) $\text{zer}(\mathcal{A} + \mathcal{B}) \neq \emptyset$ and $\text{zer}(\bar{\mathcal{A}} + \bar{\mathcal{B}}) \neq \emptyset$.

Proof: See Appendix I. ■

For the forthcoming convergence analysis, we take the following assumption.

Assumption 6 (Restricted cocoercivity): $\mathbb{F}$ is restricted β -cocoercive, with $\beta > 0$. □

Remark 3: The gradient of a convex function is cocoercive, while the gradient of a strictly (strongly) convex function is strictly (strongly) monotone [9]. In a noncooperative game, convexity of the cost functions $\mathbb{J}_i$ in the local variables x_i (Assumption 2) does not imply monotonicity, nor cocoercivity (unless we restrict to the so-called jointly convex case [6]). □

Then, in light of Assumption 6, the two operators $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$ in (10) have the following properties (in the Φ -induced norm).

Lemma 2: Let Assumptions 1-6 hold and let $\Phi \succ 0$. The operators $\bar{\mathcal{A}}$ and $\bar{\mathcal{B}}$ in (10) have the following properties:

- (i) $\bar{\mathcal{A}}$ is θ -cocoercive where $0 < \theta \leq \min\{\frac{1}{2d^*}, \beta\}$ and d^* is the maximum weighted degree of $\mathcal{G}^\lambda$;
- (ii) The operator $\bar{\mathcal{B}}$ is maximally monotone;
- (iii) $\Phi^{-1}\bar{\mathcal{A}}$ is $\theta\gamma$ -cocoercive where $\gamma = \frac{1}{|\Phi^{-1}|}$;
- (iv) $\Phi^{-1}\bar{\mathcal{B}}$ is maximally monotone.

Proof: See Appendix I. ■

Since the expected value can be hard to compute, as the distribution of the random variable is unknown, we take an approximation of the pseudo-gradient via the stochastic approximation (SA) scheme with variance reduction. We note that the preconditioning can be done independently of the approximation scheme.

We assume that the agents have access to an increasing number S^k of samples of the random variable ξ and to be able to compute an approximation of $\mathbb{F}(\mathbf{x})$ of the form

$$\hat{F}(\mathbf{x}, \xi) = \text{col} \left(\frac{1}{S^k} \sum_{t=1}^{S^k} \nabla_{x_1} f_1(\mathbf{x}, \xi_1^{(t)}), \dots, \frac{1}{S^k} \sum_{t=1}^{S^k} \nabla_{x_N} f_N(\mathbf{x}, \xi_N^{(t)}) \right). \quad (11)$$

where $\xi = \text{col}(\xi_1, \dots, \xi_N)$ and for all $i \in \mathcal{I}$, $\xi_i = \text{col}(\xi_i^{(1)}, \dots, \xi_i^{(S^k)})$ is an i.i.d. sequence of random variables sampled from $\mathbb{P}$. The random variables should be i.i.d. to avoid correlation between the variables involved (see [20], [29] for more details). Approximations of the form in (11) are common in Monte-Carlo simulation approaches, machine learning and computational statistics.

Then, with the SA of the pseudogradient in (11), we approximate $\bar{\mathcal{A}}$ in (10) with

$$\hat{\mathcal{A}} : \begin{bmatrix} (\mathbf{x}, \xi) \\ \mathbf{z} \\ \boldsymbol{\lambda} \end{bmatrix} \mapsto \begin{bmatrix} \hat{F}(\mathbf{x}, \xi) \\ 0 \\ \mathbf{b} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \mathbf{L}\boldsymbol{\lambda} \end{bmatrix}. \quad (12)$$

Given any $\Phi \succ 0$, our fixed point problem reads as

$$\boldsymbol{\omega} = (\text{Id} + \Phi^{-1}\bar{\mathcal{B}})^{-1} \circ (\text{Id} - \Phi^{-1}\bar{\mathcal{A}})(\boldsymbol{\omega}) \quad (13)$$

and it suggests the stochastic FB algorithm

$$\boldsymbol{\omega}^{k+1} = (\text{Id} + \Phi^{-1}\bar{\mathcal{B}})^{-1} \circ (\text{Id} - \Phi^{-1}\hat{\mathcal{A}})(\boldsymbol{\omega}^k, \xi^k). \quad (14)$$

where $(\text{Id} + \Phi^{-1}\bar{\mathcal{B}})^{-1}$ represents the backward step and $(\text{Id} - \Phi^{-1}\hat{\mathcal{A}})$ is the forward one. We remark that $\hat{\mathcal{A}}$ is a sampled and computable, approximation of the operator $\bar{\mathcal{A}}$ in (10).

By expanding (14) and solving for $\mathbf{x}_k$, $\mathbf{z}_k$ and $\boldsymbol{\lambda}_k$, we obtain the distributed FB steps in Algorithm 1 with $\hat{\mathcal{A}}$ as in (12), $\bar{\mathcal{B}}$ as in (10) and

$$\Phi = \begin{bmatrix} \alpha^{-1} & 0 & -\mathbf{A}^\top \\ 0 & \nu^{-1} & -\mathbf{L} \\ -\mathbf{A} & -\mathbf{L} & \sigma^{-1} \end{bmatrix}. \quad (15)$$

In (15), $\alpha^{-1} = \text{diag}\{\alpha_1^{-1} \mathbf{I}_{n_1}, \dots, \alpha_N^{-1} \mathbf{I}_{n_N}\} \in \mathbb{R}^{n \times n}$ and similarly σ^{-1} and ν^{-1} are block diagonal matrices of suitable dimensions. We note that Φ is symmetric and such that $\boldsymbol{\omega}^k$ is easy to be computed. Moreover, with Φ as in (15), we overcome the problem to distribute the computation of the resolvent of $\bar{\mathcal{B}}$, the first operator in (14) [1].

IV. CONVERGENCE ANALYSIS

Since we use the approximation in (11), for all $k \geq 0$, let us introduce the approximation error

$$\epsilon^k = \hat{F}(\mathbf{x}^k, \xi^k) - \mathbb{F}(\mathbf{x}^k). \quad (16)$$

and consequently $\varepsilon^k = \hat{\mathcal{A}}(\boldsymbol{\omega}^k, \xi^k) - \bar{\mathcal{A}}(\boldsymbol{\omega}^k) = \text{col}(\epsilon^k, 0, 0)$, with $\hat{\mathcal{A}}$ as in (12). The following assumption is widely used in the stochastic framework [17], [20].

Assumption 7 (Zero mean error): For $k \in \mathbb{N}$, $\mathbb{E}[\epsilon^k | \mathcal{F}^k] = 0$ a.s. □

To guarantee that Φ is positive definite and in turn to obtain convergence, the step sizes sequence can be taken constant but it should satisfy some bounds [1, Lem. 6].

Assumption 8 (Bounded step sizes): For any agent $i \in \mathcal{I}$ and an arbitrary $\gamma > 0$, the step sizes are such that:

$$\begin{aligned} 0 < \alpha_i &\leq (\gamma + \max_{j \in \{1, \dots, n_i\}} \sum_{k=1}^m |[A_i^\top]_{jk}|)^{-1} \\ 0 < \nu_i &\leq (\gamma + 2d_i)^{-1} \\ 0 < \sigma_i &\leq (\gamma + 2d_i + \max_{j \in \{1, \dots, m\}} \sum_{k=1}^{n_i} |[A_i]_{jk}|)^{-1} \end{aligned}$$

where $[A_i^\top]_{jk}$ indicates the entry (j, k) of the matrix $A_i^\top$. Moreover, $\|\Phi^{-1}\| < 2\theta$, where θ is the cocoercivity constant of $\bar{\mathcal{A}}$ as in Lemma 2. □

The number of samples to be taken must satisfy some conditions as well.

Assumption 9 (Increasing batch size): The batch size sequence $(S^k)_{k \geq 1}$ is such that $S^k \geq c(k + k_0)^{a+1}$, for some $c, k_0, a > 0$. □ This assumption implies that $1/S^k$ is summable, which is usual in variance reduced schemes [20].

Assumption 10 (Variance reduction): There exist $p \geq 2$, $\sigma_1 \geq 0$, and a measurable locally bounded function $\sigma : \text{SOL}(\mathcal{X}, \mathbb{F}) \rightarrow \mathbb{R}$ such that for all $(\mathbf{x}, \mathbf{x}^*) \in \mathbb{R}^n \times \text{SOL}(\mathcal{X}, \mathbb{F})$

$$\mathbb{E} \left[\|\hat{F}(\mathbf{x}, \xi) - \mathbb{F}(\mathbf{x})\|^p \right]^{\frac{1}{p}} \leq \sigma(\mathbf{x}^*) + \sigma_1 \|\mathbf{x} - \mathbf{x}^*\|.$$

where $\hat{F}(\mathbf{x}, \xi) = \text{col}((\nabla_{x_i} f_i(\mathbf{x}, \xi_i))_{i \in \mathbb{N}})$ is a single approximation of the pseudogradient given one realization $\xi = \text{col}(\xi_1, \dots, \xi_N)$. □

Remark 4: For simplicity of presentation, let us consider a stronger condition than Assumption 10, namely, for all $\mathbf{x} \in \mathcal{X}$

$$\mathbb{E} \left[\|\hat{F}(\mathbf{x}, \xi) - \mathbb{F}(\mathbf{x})\|^2 \right] \leq \sigma^2 \quad (17)$$

for some $\sigma > 0$. In the literature, (17) is known as uniform bounded variance. Assumption 10 is more natural when the feasible set is unbounded and it is always satisfied when the mapping $\mathbb{F}$ is Caratheodory and Lipschitz continuous [23, Ex. 3.1]. Since we are in a game theoretic setup, our feasible set is bounded, we can use (17) as a variance control assumption.

We emphasize that all the following results hold also in the more general case given by Assumption 10 and using the L_p norm for any $p \geq 2$. We refer to [20] for a more detailed insight on this general case. □

We are now ready to state our main convergence result.

Theorem 1: Let Assumptions 1–10 hold. Then, the sequence $(\mathbf{x}^k)_{k \in \mathbb{N}}$ generated by Algorithm 1 converges a.s. to a v-SGNE of the game in (4).

Proof: See Appendix III. ■

V. COMPARATIVE NUMERICAL SIMULATIONS

Let us propose a number of numerical simulations to corroborate the theoretical analysis. We also compare our algorithm with the forward-backward-forward (FBF) [23] and the extragradient (EG) [20] algorithms. To obtain the two algorithms, let us rewrite the operator $\tilde{B}$ in (10) as $\tilde{B} = \mathcal{C} + \mathcal{D}$, where $\mathcal{C}$ contains the local constraints and $\mathcal{D}$ is the skew symmetric matrix. Let us also call $\mathcal{H} = \tilde{A} + \mathcal{D}$. Then, in compact form, the FBF algorithm generates two sequences $(\mathbf{u}^k, \mathbf{v}^k)_{k \geq 0}$ as follows:

$$\begin{aligned} \mathbf{u}^k &= J_{\Psi^{-1}\mathcal{C}}(\mathbf{v}^k - \Psi^{-1}\mathcal{H}\mathbf{v}^k) \\ \mathbf{v}^{k+1} &= \mathbf{u}^k + \Psi^{-1}(\mathcal{H}\mathbf{v}^k - \mathcal{H}\mathbf{u}^k). \end{aligned} \quad (18)$$

In (18), Ψ is a block-diagonal matrix with the step sizes:

$$\Psi = \text{diag}(\alpha^{-1}, \nu^{-1}, \sigma^{-1}), \quad (19)$$

where α , ν and σ are diagonal matrices of suitable dimensions. The convergence of the stochastic FBF with is guaranteed by [23, Thm. 4.5].

Analogously, we can write the stochastic distributed EG algorithm in compact form as:

$$\begin{aligned} \mathbf{u}^k &= J_{\Psi^{-1}\mathcal{C}}(\mathbf{v}^k - \Psi^{-1}\mathcal{H}\mathbf{v}^k) \\ \mathbf{v}^{k+1} &= J_{\Psi^{-1}\mathcal{C}}(\mathbf{v}^k - \Psi^{-1}\mathcal{H}\mathbf{u}^k). \end{aligned} \quad (20)$$

In this case, convergence is guaranteed by [20, Thm. 3.18].

As a case study, we consider the electricity market problem proposed in [11], which can be casted as a network Cournot game with markets capacity constraints [1], [18].

We consider a set of $N = 20$ generators (companies) that sell energy in a set of $m = 7$ locations (markets). The random variable ξ represents the uncertainty in the demand. Each company decides the quantity x_i of energy to deliver in the n_i markets it is connected with and its has a local constraint of the form $0 < x_i \leq \gamma_i$, where each component of γ_i is randomly drawn from $[1, 1.5]$. This constraint can be seen as the capacity limit of generator i . Each market has a bounded capacity b_j randomly drawn from $[0.5, 1]$. The collective constraints are given by $A\mathbf{x} \leq \mathbf{b}$ where $A = [A_1, \dots, A_N]$. Each A_i specifies in which market a company i participates. This information can be retrieved from the graph in [1, Fig. 1]. Each company has a local cost function related to the production of electricity which is $c_i(x_i) = \pi_i \sum_{j=1}^{n_i} ([x_i]_j)^2 + q_i^\top x_i$, where $[x_i]_j$ indicates the j component of x_i . π_i is randomly drawn from $[1, 8]$ and each component of q_i is randomly drawn from $[0.1, 0.6]$. The cost function is not uncertain as we suppose that the companies are able to compute their own cost of production. Moreover, c_i is strongly convex with Lipschitz continuous gradient. The prices of the locations are collected in $P: \mathbb{R}^m \times \Xi \rightarrow \mathbb{R}^m$. The uncertain variable appears in this functional. The price $P(\mathbf{x}, \xi) = \bar{P} - D(\xi)A\mathbf{x}$ is taken as an affine function and each component of $\bar{P} = \text{col}(\bar{P}_1, \dots, \bar{P}_7)$ is randomly drawn from $[2, 4]$. The uncertainty appears in the quantities $D(\xi) = \text{diag}\{d_1(\xi_1), \dots, d_7(\xi_7)\}$ that concern the total supply for each market. The entries of $D(\xi)$ are taken with a normal distribution with mean 0.8 and finite variance. The cost function of each agent is then given by $\mathbb{J}_i(x_i, x_{-i}, \xi_i) = c_i(x_i) - \mathbb{E}[P(\mathbf{x}, \xi)^\top A_i x_i]$. Since $c_i(x_i)$ is strongly convex with Lipschitz continuous gradient and the prices are linear, the pseudo-gradient mapping is strongly monotone. Following [1], the dual variables graph is a cycle graph with the addition of the edges (2, 15) and (6, 13).

We simulate the FB, FBF and EG algorithms to draw a comparison using the SA scheme with variance reduction. The parameters α , ν and σ are taken to be the largest possible to ensure convergence. The plot in Fig. 1 shows the relative distance from the unique v-SGNE $\mathbf{x}^*$. The plot in Fig. 2 shows the computational time needed

to reach a solution. From the figures, our algorithm is comparable in terms of number of iterations to the FBF algorithm but it is the least computationally expensive.




Fig. 1. Relative distance of the primal variable from the solution.




Fig. 2. Computation time (achieved on Matlab R2019a with a 2,3 GHz Intel Core i5 and 8 GB LPDDR3 RAM).

VI. CONCLUSION

The preconditioned forward-backward operator splitting is applicable to stochastic generalized Nash equilibrium problems to design distributed equilibrium seeking algorithms. Since the exact expected value is hard to compute, the sample average approximation can be used to ensure almost sure convergence. Convergence holds under restricted cocoercivity of the pseudo-gradient mapping. This assumption is among the weakest known in the literature and, remarkably, match well with the deterministic problem setup.

APPENDIX I PROOFS OF SECTION III

Proof: [Proof of Lemma 1] The proof of (i) follows similarly to [1, Thm. 2]. Concerning (ii), given Assumptions 1–4, the game in (4) has at least one solution $\mathbf{x}^*$, therefore, there exists a $\lambda^* \in \mathbb{R}_{\geq 0}^m$ such that the KKT conditions in (7) are satisfied [28, Thm.

3.1]. It follows that $\text{zer}(\mathcal{A} + \mathcal{B}) \neq \emptyset$. The existence of $\mathbf{z}^*$ such that $\text{col}(\mathbf{x}^*, \mathbf{z}^*, \boldsymbol{\lambda}^*) \in \text{zer}(\bar{\mathcal{A}} + \bar{\mathcal{B}})$ follows from properties of the normal cone and of the Laplacian matrix as a consequence of Assumption 5 [1, Thm. 2]. ■

Proof: [Proof of Lemma 2] First we note that $\|L\| \geq 2d^*$ and that by the Baillon-Haddad Theorem the Laplacian operator is $\frac{1}{2d^*}$ -cocoercive. Then Statement (i) follows by this computation:

$$\begin{aligned} & \langle \bar{\mathcal{A}}(\omega_1) - \bar{\mathcal{A}}(\omega_2), \omega_1 - \omega_2 \rangle \\ &= \langle \mathbb{F}(\mathbf{x}_1) - \mathbb{F}(\mathbf{x}_2), x_1 - x_2 \rangle + \langle L\lambda_1 - L\lambda_2, \lambda_1 - \lambda_2 \rangle \\ &\geq \beta \|\mathbb{F}(\mathbf{x}_1) - \mathbb{F}(\mathbf{x}_2)\|^2 + \frac{1}{2d^*} \|L\lambda_1 - L\lambda_2\|^2 \\ &\geq \min \left\{ \beta, \frac{1}{2d^*} \right\} (\|\mathbb{F}(\mathbf{x}_1) - \mathbb{F}(\mathbf{x}_2)\|^2 + \|L\lambda_1 - L\lambda_2\|^2) \\ &\geq \theta \|\bar{\mathcal{A}}(\omega_1) - \bar{\mathcal{A}}(\omega_2)\|^2. \end{aligned}$$

The operator $\bar{\mathcal{B}}$ is given by a sum, therefore it is maximally monotone if both the addend are [9, Prop. 20.23]. The first part is maximally monotone by Assumption 1 and Moreau Theorem [9, Thm. 20.25] and the second part is a skew symmetric matrix [9, Cor. 20.28]. Statement (iii) follows from Statement (i) and (iv) follows from (ii) [1, Lem. 7]. ■

APPENDIX II SEQUENCE OF RANDOM VARIABLES

In this appendix we recall some results on sequences of random variables, given the probability space $(\Xi, \mathcal{F}, \mathbb{P})$.

Let us define the filtration $\mathcal{F} = \{\mathcal{F}^k\}$, that is, a family of σ -algebras such that $\mathcal{F}_0 = \sigma(X_0)$, for all $k \geq 1$, $\mathcal{F}_k = \sigma(X_0, \xi_1, \xi_2, \dots, \xi_k)$ and $\mathcal{F}^k \subseteq \mathcal{F}^{k+1}$ for all $k \geq 0$.

The Robbins-Siegmund Lemma is widely used in literature to prove a.s. convergence of sequences of random variables. It first appeared in [30].

Lemma 3 (Robbins-Siegmund Lemma, [30]): Let $\mathcal{F} = (\mathcal{F}^k)_{k \in \mathbb{N}}$ be a filtration. Let $\{\alpha^k\}_{k \in \mathbb{N}}$, $\{\theta^k\}_{k \in \mathbb{N}}$, $\{\eta^k\}_{k \in \mathbb{N}}$ and $\{\chi^k\}_{k \in \mathbb{N}}$ be non negative sequences such that $\sum_k \eta^k < \infty$, $\sum_k \chi^k < \infty$ and let

$$\forall k \in \mathbb{N}, \quad \mathbb{E}[\alpha_{k+1} | \mathcal{F}^k] + \theta^k \leq (1 + \chi^k) \alpha^k + \eta^k \quad a.s.$$

Then $\sum_k \theta^k < \infty$ and $\{\alpha^k\}_{k \in \mathbb{N}}$ converges a.s. to a non negative random variable.

We also need this result for L_p norms, known as Burkholder-Davis-Gundy inequality [31].

Lemma 4 (Burkholder-Davis-Gundy inequality): Let $\{\mathcal{F}^k\}$ be a filtration and $\{U^k\}_{k \geq 0}$ a vector-valued martingale relative to this filtration. Then, for all $p \in [1, \infty)$, there exists a universal constant $c_p > 0$ such that for every $k \geq 1$

$$\mathbb{E} \left[\left(\sup_{0 \leq i \leq N} \|U_i\| \right)^p \right]^{\frac{1}{p}} \leq c_p \mathbb{E} \left[\left(\sum_{i=1}^N \|U_i - U_{i-1}\|^2 \right)^{\frac{p}{2}} \right]^{\frac{1}{p}}. \quad \square$$

We also recall the Minkowski inequality: for given functions $f, g \in L^p(\Xi, \mathcal{F}, \mathbb{P})$, $\mathcal{G} \subseteq \mathcal{F}$ and $p \in [0, \infty]$

$$\mathbb{E} [\|f + g\|^p | \mathcal{G}]^{\frac{1}{p}} \leq \mathbb{E} [\|f\|^p | \mathcal{G}]^{\frac{1}{p}} + \mathbb{E} [\|g\|^p | \mathcal{G}]^{\frac{1}{p}}.$$

When combined with the Burkholder-Davis-Gundy inequality, it leads to the fact that for all $p \geq 2$, there exists a constant $c_p > 0$ such that, for every $k \geq 1$,

$$\mathbb{E} \left[\left(\sup_{0 \leq i \leq N} \|U_i\| \right)^p \right]^{\frac{1}{p}} \leq c_p \sqrt{\sum_{k=1}^N \mathbb{E} (\|U_i - U_{i-1}\|^2)^{\frac{p}{2}}}.$$

APPENDIX III PROOF OF THEOREM 1

In this section, we prove convergence of Algorithm 1. Since the iterations of Algorithm 1 are obtained by expanding (14), convergence of the sequence $(\mathbf{x}^k, \boldsymbol{\lambda}^k)$ to a v-GNE of the game in (4) follows by the convergence of the FB iterations in (14).

We now prove a preliminary result concerning the sequence generated by Algorithm 1. Due to spacing limitations, we use $\bar{\mathcal{A}}$ instead of $\hat{\mathcal{A}}$.

Lemma 5: Let Assumptions 1-8 hold. Then, the sequence $(\omega^k)_{k \in \mathbb{N}}$ generated by (14) satisfies the inequality

$$\begin{aligned} \mathbb{E} [\|\omega^{k+1} - \omega^*\|_{\Phi}^2 | \mathcal{F}^k] &\leq \|\omega^k - \omega^*\|_{\Phi}^2 + \\ &+ 2\mathbb{E} [\|\Phi^{-1} \varepsilon^k\|_{\Phi}^2 | \mathcal{F}^k] + \frac{1}{2} \left(\frac{1}{\zeta} - 1 \right) \text{res}_{\Phi}(\omega^k)^2. \end{aligned} \quad (21)$$

Proof: We start by using firmly nonexpansiveness of the resolvent [9, Cor. 23.9] and the fact that if ω^* is a solution then $\omega^* = (\text{Id} + \Phi^{-1} \bar{\mathcal{B}})^{-1} (\text{Id} - \Phi^{-1} \bar{\mathcal{A}}) \omega^*$:

$$\begin{aligned} \|\omega^{k+1} - \omega^*\|_{\Phi}^2 &\leq \|\omega^k - \omega^*\|_{\Phi}^2 + 2\langle \omega^k - \omega^*, \Phi^{-1} \varepsilon^k \rangle_{\Phi} \\ &- 2\langle \omega^k - \omega^*, \Phi^{-1} (\bar{\mathcal{A}}(\omega^k) - \bar{\mathcal{A}}(\omega^*)) \rangle_{\Phi} + \\ &- \|\omega^k - \omega^{k+1}\|_{\Phi}^2 + 2\langle \omega^k - \omega^{k+1}, \Phi^{-1} (\hat{\mathcal{A}}(\omega^k) - \bar{\mathcal{A}}(\omega^*)) \rangle_{\Phi} \end{aligned}$$

By Young's inequality with $\zeta > 1$ and such that Assumption 8 is satisfied, we have that

$$\begin{aligned} 2\langle \omega^k - \omega^{k+1}, \Phi^{-1} (\hat{\mathcal{A}}(\omega^k) - \bar{\mathcal{A}}(\omega^*)) \rangle_{\Phi} &\leq \frac{1}{\zeta} \|\omega^k - \omega^{k+1}\|_{\Phi}^2 + \\ &+ \zeta \|\Phi^{-1} (\bar{\mathcal{A}}(\omega^k) - \bar{\mathcal{A}}(\omega^*))\|_{\Phi}^2 + \zeta \|\Phi^{-1} \varepsilon^k\|_{\Phi}^2 + \\ &+ 2\zeta \langle \Phi^{-1} \bar{\mathcal{A}}(\omega^k) - \Phi^{-1} \bar{\mathcal{A}}(\omega^*), \varepsilon^k \rangle_{\Phi} \end{aligned} \quad (22)$$

Then, by using cocoercivity and including (22), we obtain:

$$\begin{aligned} \|\omega^{k+1} - \omega^*\|_{\Phi}^2 &\leq \|\omega^k - \omega^*\|_{\Phi}^2 + \zeta \|\Phi^{-1} \varepsilon^k\|_{\Phi}^2 + \\ &+ \left(\frac{1}{\zeta} - 1 \right) \|\omega^k - \omega^{k+1}\|_{\Phi}^2 + 2\langle \omega^k - \omega^*, \Phi^{-1} \varepsilon^k \rangle_{\Phi} + \\ &+ \left(\frac{\zeta \|\Phi^{-1}\|}{\theta} - 2 \right) \langle \omega^k - \omega^*, \Phi^{-1} \bar{\mathcal{A}}(\omega^k) - \Phi^{-1} \bar{\mathcal{A}}(\omega^*) \rangle_{\Phi} \\ &+ 2\zeta \langle \Phi^{-1} \bar{\mathcal{A}}(\omega^k) - \Phi^{-1} \bar{\mathcal{A}}(\omega^*), \varepsilon^k \rangle_{\Phi} \end{aligned} \quad (23)$$

Next, given the residual $\text{res}(\omega^k)$, we have that

$$\begin{aligned} \text{res}_{\Phi}(\omega^k)^2 &= \|\omega^k - (\text{Id} + \Phi^{-1} \bar{\mathcal{B}})^{-1} (\omega^k - \Phi^{-1} \bar{\mathcal{A}}(\omega^k))\|_{\Phi}^2 \\ &\leq 2\|\omega^k - \omega^{k+1}\|_{\Phi}^2 + 2\|(\text{Id} + \Phi^{-1} \bar{\mathcal{B}})^{-1} (\omega^k - \Phi^{-1} \hat{\mathcal{A}}(\omega^k, \xi^k)) \\ &- (\text{Id} + \Phi^{-1} \bar{\mathcal{B}})^{-1} (\omega^k - \Phi^{-1} \bar{\mathcal{A}}(\omega^k))\|_{\Phi}^2 \\ &\leq 2\|\omega^k - \omega^{k+1}\|_{\Phi}^2 + 2\|\Phi^{-1} \varepsilon^k\|_{\Phi}^2 \end{aligned} \quad (24)$$

where the first equality follows by the definition of ω^{k+1} and the last inequality follows from non expansivity. Then,

$$\|\omega^k - \omega^{k+1}\|_{\Phi}^2 \geq \frac{1}{2} \text{res}_{\Phi}(\omega^k)^2 - \|\Phi^{-1} \varepsilon^k\|_{\Phi}^2$$

Finally, equation (23) becomes

$$\begin{aligned} \|\omega^{k+1} - \omega^*\|_{\Phi}^2 &\leq \|\omega^k - \omega^*\|_{\Phi}^2 + \left(\zeta - \frac{1}{\zeta} + 1 \right) \|\Phi^{-1} \varepsilon^k\|_{\Phi}^2 \\ &+ 2\langle \omega^k - \omega^*, \Phi^{-1} \varepsilon^k \rangle_{\Phi} + \frac{1}{2} \left(\frac{1}{\zeta} - 1 \right) \text{res}_{\Phi}(\omega^k)^2 \\ &+ \left(\frac{\zeta \|\Phi^{-1}\|}{\theta} - 2 \right) \langle \omega^k - \omega^*, \Phi^{-1} \bar{\mathcal{A}}(\omega^k) - \Phi^{-1} \bar{\mathcal{A}}(\omega^*) \rangle_{\Phi} \\ &+ 2\zeta \langle \Phi^{-1} \bar{\mathcal{A}}(\omega^k) - \Phi^{-1} \bar{\mathcal{A}}(\omega^*), \varepsilon^k \rangle_{\Phi} \end{aligned}$$

By Assumption 8 and by monotonicity, the second last term is non-positive, hence, by taking the expected value and by using Assumption 7, we obtain (21). ■

Before proving convergence of the algorithm, we prove a preliminary result on the variance of the stochastic error.

Lemma 6: For all $k \geq 0$, if Assumption 10 hold, we have that

$$\mathbb{E} [\|\epsilon^k\|^2 | \mathcal{F}^k] \leq \frac{c\sigma^2}{S^k} \text{ a.s.}$$

Proof: We first prove that $\mathbb{E} [\|\epsilon^k\|^2 | \mathcal{F}^k]^{\frac{1}{2}} \leq \frac{c\sigma}{\sqrt{S^k}} \text{ a.s.}$, then the claim follows immediately. Define the process $\{M_S^S(x)\}_{i=0}^S$ as $M_0(x) = 0$ and for $1 \leq t \leq S$

$$M_t^S(x) = \frac{1}{S} \sum_{l=1}^t \tilde{F}(x, \xi^{(l)}) - \mathbb{F}(x).$$

Let $\mathcal{F}_t = \sigma(\xi^{(1)}, \dots, \xi^{(t)})$. Then $\{M_t^S(x), \mathcal{F}_t\}_{t=1}^S$ is a martingale starting at 0. Let

$$\Delta M_{t-1}^S(x) = M_t^S(x) - M_{t-1}^S(x) = \tilde{F}(x, \xi^{(t)}) - \mathbb{F}(x).$$

Then, by Equation (17), we have

$$\mathbb{E} [\|\Delta M_{t-1}^S\|^2]^{\frac{1}{2}} = \frac{1}{S} \mathbb{E} [\|\tilde{F}(x, \xi^{(t)}) - \mathbb{F}(x)\|^2]^{\frac{1}{2}} \leq \frac{\sigma}{S}.$$

By applying Lemma 4, we have

$$\begin{aligned} \mathbb{E} [\|M_S^S(x)\|^2]^{\frac{1}{2}} &\leq c_2 \sqrt{\sum_{i=1}^N \mathbb{E} \left[\left\| \frac{\tilde{F}(x, \xi^{(t)}) - \mathbb{F}(x)}{S} \right\|^2 \right]} \\ &\leq c_2 \sqrt{\frac{1}{S^2} \sum_{i=1}^N \mathbb{E} [\|\tilde{F}(x, \xi^{(t)}) - \mathbb{F}(x)\|^2]} \leq \frac{c_2 \sigma}{\sqrt{S}}. \end{aligned}$$

We note that from (11) $\hat{F}(x, \xi) = \frac{1}{S^k} \sum_{t=1}^{S^k} \tilde{F}(x, \xi^{(t)})$ and $M_{S^k}^{S^k}(x^k) = \epsilon^k$, hence by taking the square, the claim follows. ■

We note that if Lemma 6 holds, then it follows that

$$\mathbb{E} [\|\Phi^{-1} \epsilon^k\|_{\Phi}^2 | \mathcal{F}^k] \leq \frac{c\sigma^2 \|\Phi^{-1}\|}{S^k}.$$

We are now ready to prove the main convergence result.

Proof: [Proof of Theorem 1] By Lemmas 5 and 6 we have that

$$\begin{aligned} \mathbb{E} [\|\omega^{k+1} - \omega^*\|_{\Phi}^2 | \mathcal{F}^k] &\leq \|\omega^k - \omega^*\|_{\Phi}^2 + 2 \frac{c\sigma^2 \|\Phi^{-1}\|}{S^k} \\ &\quad + \frac{1}{2} \left(\frac{1}{\zeta} - 1 \right) \text{res}(\omega^k)^2. \end{aligned}$$

By Lemma 3 we conclude that the sequence $(\omega^k)_{k \in \mathbb{N}}$ is bounded and has a cluster point $\bar{\omega}$. Since $\sum \theta^k < \infty$, it follows that $\lim_{k \rightarrow \infty} \text{res}(\omega^k) = 0$ as $k \rightarrow \infty$ and $\text{res}(\bar{\omega}) = 0$. ■

REFERENCES

- [1] P. Yi and L. Pavel, "An operator splitting approach for distributed generalized Nash equilibria computation," *Automatica*, vol. 102, pp. 111–121, 2019.
- [2] A. A. Kulkarni and U. V. Shanbhag, "On the variational equilibrium as a refinement of the generalized Nash equilibrium," *Automatica*, vol. 48, no. 1, pp. 45–55, 2012.
- [3] L. Pavel, "Distributed gne seeking under partial-decision information over networks via a doubly-augmented operator splitting approach," *IEEE Transactions on Automatic Control*, 2019.
- [4] 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.
- [5] L. Pavel, "An extension of duality to a game-theoretic framework," *Automatica*, vol. 43, no. 2, pp. 226–237, 2007.
- [6] F. Facchinei and C. Kanzow, "Generalized Nash equilibrium problems," *Annals of Operations Research*, vol. 175, no. 1, pp. 177–211, 2010.
- [7] G. Belgioioso and S. Grammatico, "Projected-gradient algorithms for generalized equilibrium seeking in aggregative games are preconditioned forward-backward methods," in *2018 European Control Conference (ECC)*. IEEE, 2018, pp. 2188–2193.
- [8] —, "Semi-decentralized Nash equilibrium seeking in aggregative games with separable coupling constraints and non-differentiable cost functions," *IEEE control systems letters*, vol. 1, no. 2, pp. 400–405, 2017.
- [9] H. H. Bauschke, P. L. Combettes *et al.*, *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011, vol. 408.
- [10] B. Franci and S. Grammatico, "A damped forward-backward algorithm for stochastic generalized nash equilibrium seeking," in *2020 European Control Conference (ECC)*. IEEE, 2020, pp. 1117–1122.
- [11] H. Xu and D. Zhang, "Stochastic Nash equilibrium problems: sample average approximation and applications," *Computational Optimization and Applications*, vol. 55, no. 3, pp. 597–645, 2013.
- [12] D. Watling, "User equilibrium traffic network assignment with stochastic travel times and late arrival penalty," *European Journal of Operational Research*, vol. 175, no. 3, pp. 1539–1556, 2006.
- [13] R. Henrion and W. Römisch, "On m-stationary points for a stochastic equilibrium problem under equilibrium constraints in electricity spot market modeling," *Applications of Mathematics*, vol. 52, no. 6, pp. 473–494, 2007.
- [14] V. DeMiguel and H. Xu, "A stochastic multiple-leader Stackelberg model: analysis, computation, and application," *Operations Research*, vol. 57, no. 5, pp. 1220–1235, 2009.
- [15] I. Abada, S. Gabriel, V. Briat, and O. Massol, "A generalized Nash-Cournot model for the Northwestern European natural gas markets with a fuel substitution demand function: The gammes model," *Networks and Spatial Economics*, vol. 13, no. 1, pp. 1–42, 2013.
- [16] U. Ravat and U. V. Shanbhag, "On the characterization of solution sets of smooth and nonsmooth convex stochastic Nash games," *SIAM Journal on Optimization*, vol. 21, no. 3, pp. 1168–1199, 2011.
- [17] J. Koshal, A. Nedic, and U. V. Shanbhag, "Regularized iterative stochastic approximation methods for stochastic variational inequality problems," *IEEE Transactions on Automatic Control*, vol. 58, no. 3, pp. 594–609, 2013.
- [18] C.-K. Yu, M. Van Der Schaar, and A. H. Sayed, "Distributed learning for stochastic generalized Nash equilibrium problems," *IEEE Transactions on Signal Processing*, vol. 65, no. 15, pp. 3893–3908, 2017.
- [19] F. Yousefian, A. Nedić, and U. V. Shanbhag, "On smoothing, regularization, and averaging in stochastic approximation methods for stochastic variational inequality problems," *Mathematical Programming*, vol. 165, no. 1, pp. 391–431, 2017.
- [20] A. Iusem, A. Jofré, R. I. Oliveira, and P. Thompson, "Extragradient method with variance reduction for stochastic variational inequalities," *SIAM Journal on Optimization*, vol. 27, no. 2, pp. 686–724, 2017.
- [21] H. Jiang and H. Xu, "Stochastic approximation approaches to the stochastic variational inequality problem," *IEEE Transactions on Automatic Control*, vol. 53, no. 6, pp. 1462–1475, 2008.
- [22] L. Rosasco, S. Villa, and B. C. Vũ, "Stochastic forward-backward splitting for monotone inclusions," *Journal of Optimization Theory and Applications*, vol. 169, no. 2, pp. 388–406, 2016.
- [23] R. I. Bot, P. Mertikopoulos, M. Staudigl, and P. T. Vuong, "Forward-backward-forward methods with variance reduction for stochastic variational inequalities," *arXiv preprint arXiv:1902.03355*, 2019.
- [24] H. Robbins and S. Monro, "A stochastic approximation method," *The annals of mathematical statistics*, pp. 400–407, 1951.
- [25] F. Facchinei and J. S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.
- [26] D. P. Palomar and Y. C. Eldar, *Convex optimization in signal processing and communications*. Cambridge university press, 2010.
- [27] F. Facchinei, A. Fischer, and V. Piccialli, "On generalized Nash games and variational inequalities," *Operations Research Letters*, vol. 35, no. 2, pp. 159–164, 2007.
- [28] A. Auslender and M. Teboulle, "Lagrangian duality and related multiplier methods for variational inequality problems," *SIAM Journal on Optimization*, vol. 10, no. 4, pp. 1097–1115, 2000.
- [29] A. N. Iusem, A. Jofré, R. I. Oliveira, and P. Thompson, "Variance-based extragradient methods with line search for stochastic variational inequalities," *SIAM Journal on Optimization*, vol. 29, no. 1, pp. 175–206, 2019.
- [30] H. Robbins and D. Siegmund, "A convergence theorem for non negative almost supermartingales and some applications," in *Optimizing methods in statistics*. Elsevier, 1971, pp. 233–257.
- [31] D. W. Stroock, *Probability theory: an analytic view*. Cambridge university press, 2010.