

Certifying ϵ -equilibria in stochastic games with limited data availability

Questa è la versione sottoposta a revisione paritaria (postprint) della seguente opera:

Original

Certifying ϵ -equilibria in stochastic games with limited data availability / Fabiani, F., Franci, B.. - In: EUROPEAN JOURNAL OF CONTROL. - ISSN 1435-5671. - (2026). [10.1016/j.ejcon.2026.101558]

Availability:

This version is available at: 20.500.11771/43358

Publisher:

Elsevier

Published

DOI:10.1016/j.ejcon.2026.101558

Terms of use:

This publication is made accessible in accordance with the terms for deposit in the institutional repository, as defined by the IMT School for Advanced Studies Lucca's Open Access Policy.
(https://library.imtlucca.it/sites/default/files/regolamento-policy-open-access-imtlib_0.pdf).

Si prega di consultare le pagine informative dell'editore relative alle politiche di autoarchiviazione.

(Article begins on next page)

Journal Pre-proof

Certifying ε -equilibria in stochastic games with limited data availability

Filippo Fabiani, Barbara Franci

PII: S0947-3580(26)00111-1
DOI: <https://doi.org/10.1016/j.ejcon.2026.101558>
Reference: EJCON 101558

To appear in: *European Journal of Control*

Received date: 27 May 2026
Accepted date: 18 July 2026

Please cite this article as: Filippo Fabiani, Barbara Franci, Certifying ε -equilibria in stochastic games with limited data availability, *European Journal of Control* (2025), doi: <https://doi.org/10.1016/j.ejcon.2026.101558>



This is a PDF of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability. This version will undergo additional copyediting, typesetting and review before it is published in its final form. As such, this version is no longer the Accepted Manuscript, but it is not yet the definitive Version of Record; we are providing this early version to give early visibility of the article. Please note that Elsevier's sharing policy for the Published Journal Article applies to this version, see: <https://www.elsevier.com/about/policies-and-standards/sharing#4-published-journal-article>. Please also note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2026 Published by Elsevier Ltd on behalf of European Control Association.

Certifying ε -equilibria in stochastic games with limited data availability

Filippo Fabiani and Barbara Franci

Abstract—We focus on stochastic generalized Nash equilibrium problems (SGNEPs) to establish probabilistic guarantees on the quality of the solution produced by a data-driven version of the extragradient (EG) algorithm for stochastic generalized Nash equilibrium (SGNE) seeking. Specifically, we consider a setting where the distribution of the random variable influencing the multi-agent process is unknown, and only a finite number of realizations is available for its characterization. While convergence to a SGNE should not be expected with finite data, we leverage a sample average-based approximation of the game mapping to derive distribution-free certificates bounding the distance between the true SGNE and the solution produced by the data-driven EG method. Consistently, the proposed bound shrinks as the number of samples grows to infinite.

I. INTRODUCTION

Several works have recently focused on stochastic generalized Nash equilibrium (SGNE) seeking in stochastic generalized Nash equilibrium problems (SGNEPs) [1]–[4]. Besides its tight connection with (SVIs) [1], this research area has gained popularity due to its broad applicability in real-world scenarios that involve the coordination and control of self-interested agents such as the electricity or gas markets [5], [6] and ride-hailing services [7].

In words, a SGNEP amounts to a collection of mutually coupled optimization problems in which the cost each agent aims at minimizing does not only depend on the decisions of its opponents, but also on a random variable whose distribution is typically *unknown*. Therefore, evaluating the expected cost of each agent is not feasible. To overcome this issue, available approximation schemes provide an estimation of the pseudogradient mapping that can then be implemented within iterative procedures eventually leading to an SGNE.

Among the most popular techniques, *stochastic approximation* relies on samples of the random variable to be drawn at every iteration. Specifically, one can either select a single sample and control the approximation with a vanishing step-size sequence [2], [4], [7], or increase the batch size and average the resulting operators [3], [4], [8]. The underlying approaches can demonstrate asymptotic convergence to an SGNE, however, this inevitably requires an infinite number of uncertainty realizations, which is often unrealistic in practical scenarios where only limited samples are available. Prominent examples can be found in smart grids and electricity markets based on weather forecast data. Some progress have been recently made in cases where the cost functions

consist of finite sums, i.e., with the batch size being finite by definition, although too big to be easily computed [9], [10].

Motivated by these practical considerations, we examine the scenario where only a finite number of samples are given, and use them to approximate the pseudogradient mapping of the SGNEP at hand, which is then incorporated into a popular extragradient (EG) method [8], [11]. Given such a limited data availability, exact convergence to an SGNE should not be expected. Thus, we leverage algorithmic stability arguments [12] to derive a distribution-free certificate bounding the distance between the true SGNE and the solution obtained through the data-driven EG scheme. Consistently, it turns out that our bound shrinks with increasing samples, and hence allows one to recover an SGNE with infinite data.

Besides [13], this is the first work addressing SGNE seeking in a limited data framework while providing rigorous guarantees on the quality of the solution produced that i) are easy to compute, ii) hold with high confidence, and iii) are valid regardless the probability distribution underlying data. For direct comparison, Table I reports the most common algorithms for S(G)NE seeking, together with the amount of samples needed at each iteration for convergence.

A. Algorithmic stability notions and results

Let $\mathcal{D}_s := \{\xi^{(1)}, \dots, \xi^{(s)}\}$ be a dataset consisting of s independent and identically distributed (i.i.d.) samples drawn from an *unknown* probability measure $\mathbb{P}$ attached to Ξ^s , where $\xi^{(i)} \in \Xi$ for all $i = 1, \dots, s$. Then, an *algorithm* amounts to an indexed family of mappings $\{A_s\}_{s \geq 0}$, with $A_s : \Xi^s \rightarrow \mathbb{R}^n$ taking a dataset $\mathcal{D}_s$ and returning a hypothesis $H_s := A_s(\{\xi^{(i)}\}_{i=1}^s) = A_s(\mathcal{D}_s)$ [16]. Given $\mathcal{D}_s$, we denote with $\mathcal{D}_s^i = \{\xi^{(1)}, \dots, \xi^{(i-1)}, \xi', \xi^{(i+1)}, \dots, \xi^{(s)}\}$ the set obtained by replacing the i -th element of $\mathcal{D}_s$ with $\xi' \in \Xi$ drawn according to $\mathbb{P}$, i.i.d. with respect to (w.r.t.) $\mathcal{D}_s \setminus \{\xi^{(i)}\}$. Accordingly, the associated hypothesis is $H_{s,i} := A_s(\mathcal{D}_s^i)$.

The performance of some algorithm, and particularly of the produced hypothesis H w.r.t. an example ξ , are usually measured through a *loss function* $\ell : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. Attached to the latter, one can define the so-called *generalization error*, a random variable (since the hypothesis generated depends on $\mathcal{D}_s$) defined as follows:

$$r(A, s) = \mathbb{E}_{\mathbb{P}} [\ell(H_s, \xi)], \quad (1)$$

where we use A instead of A_s as first argument. Computing (1) then requires $\mathbb{P}$, which might be unavailable in several cases. An estimate of (1) refers to the *empirical error*:

$$\hat{r}(A, s) = \frac{1}{s} \sum_{i=1}^s \ell(H_s, \xi^{(i)}). \quad (2)$$

F. Fabiani is with the IMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100, Lucca, Italy (filippo.fabiani@imtlucca.it). B. Franci is with the Department of Mathematical Sciences, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129, Torino, Italy (barbara.franci@polito.it).

	FBF [14]	Tik [7]	RFB [4]	RFB [4]	VRG [15]	FB [3]	FB [13]	LEG [10]	EG [8]	EG (This paper)
Sample(s)	N_k	1	N_k	1	$1 + NB(p)$	N_k	s	$1 + NB(p)$	N_k	s
Convergence	$k \rightarrow \infty$	$k \rightarrow \infty$	$k \rightarrow \infty$	$k \rightarrow \infty$	$k \rightarrow \infty$	$k \rightarrow \infty$	$k < K$	$k \rightarrow \infty$	$k \rightarrow \infty$	$k < K$
Equilibrium	VI/SNE	SGNE	SGNE	SNE	VI/SNE*	SGNE	ε -SGNE	SGNE*	SGNE/VI	ε -SGNE

TABLE I: The top row reports available S(G)NE algorithms, where we tacitly omit an “S” for “stochastic” in front of each acronym, and use “F” for “forward”, “B” for “backward”, “L” for “loopless”; “Tik” stands for Tikhonov regularization, RFB denotes a relaxed FB, while VRG is a variance reduced gradient iteration. The samples needed at every iteration are indicated in the second row: 1 means stochastic approximation, N_k indicates a time-varying increasing batch size, $0 \leq s < \infty$, while $B(p)$ is a Bernoulli distribution of probability p . The latter to indicate that, at every iteration, the underlying algorithm relies on $1 + N$ samples with probability p , or 1 with probability $1 - p$. The finite-time or asymptotic convergence is reported in the third row ($0 \leq K < \infty$), while the bottom one indicates whether such convergence returns an SNE, SGNE, or one of these by considering cost functions with finite sum (marked with an asterisk).

Definition 1 ([17]): An algorithm $\{A_s\}_{s \geq 0}$ has uniform stability β w.r.t. the loss function ℓ if for all $s \geq 0$, dataset with s samples $\mathcal{D}_s$, and $i \in \{1, \dots, s\}$, $|\ell(H_s, \xi) - \ell(H_{s^i}, \xi)| \leq \beta$ for all $\xi \in \Xi$. $\square$

In words, an algorithm is uniformly stable if changing one sample $\xi^{(i)}$ in its input dataset does not significantly alter the output w.r.t. to some predefined loss ℓ . In general, the smaller the β the better the stability features possessed by the algorithm at hand, measured through ℓ . The following result, based on uniform stability, will be key in the remainder:

Lemma 1 ([17, Th. 3.2]): Let $\{A_s\}_{s \geq 0}$ be an algorithm with uniform stability β w.r.t. a loss function $0 \leq \ell(H_s, \xi) \leq M$, $M \geq 0$, for all $\xi \in \Xi$ and $\mathcal{D}_s$. Then, for any $s \geq 1$ and $\delta \in (0, 1)$, it holds that: $\mathbb{P}^s\{\mathcal{D}_s \in \Xi^s : r(A, s) \leq \hat{r}(A, s) + \beta + (s\beta + M)\sqrt{2 \ln(1/\delta)/s}\} \geq 1 - \delta$. $\square$

Specifically, Lemma 1 provides with an upper bound on the risk associated to hypothesis $H_s = A_s(\mathcal{D}_s)$, which holds with arbitrary confidence $1 - \delta$, such that i) if β scales proportional to $1/s$, vanishes as $s \rightarrow \infty$, and ii) it is exponential in the confidence parameter δ . The former condition is called *consistency* property, and provides a “tight” certificate in the sense that the upper bound in Lemma 1 vanishes as $s \rightarrow \infty$.

B. Notation and operator-theoretic definitions

$\mathbb{R}$ denotes the set of real numbers and $\bar{\mathbb{R}} = \mathbb{R} \cup \{+\infty\}$. The standard inner product is represented by $\langle \cdot, \cdot \rangle : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$, while $\|\cdot\|$ denotes the associated Euclidean norm. $\mathbf{0}_m$ indicates the vector with m entries all equal to 0. Given N vectors $x_1, \dots, x_N$, with each $x_i \in \mathbb{R}^{n_i}$, $\mathbf{x} := \text{col}(x_1, \dots, x_N) = [x_1^\top, \dots, x_N^\top]^\top$, while $\mathbf{x}_{-i} := \text{col}((x_j)_{j \neq i}) = [x_1^\top, \dots, x_{i-1}^\top, x_{i+1}^\top, \dots, x_N^\top]^\top$. With some abuse of notation, we also use $\mathbf{x} = (x_i, \mathbf{x}_{-i})$. The uniform distribution on $[a, b]$ is denoted by $\mathcal{U}(a, b)$, and the normal distribution with mean μ and variance σ^2 by $\mathcal{N}(\mu, \sigma^2)$.

For all $x, y \in \text{dom}(F)$, a mapping $F : \text{dom } F \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$ is κ -Lipschitz continuous if, for some $\kappa > 0$, $\|F(x) - F(y)\| \leq \kappa \|x - y\|$; pseudo-monotone if $\langle F(y), (x - y) \rangle \geq 0 \Rightarrow \langle F(x), (x - y) \rangle \geq 0$; monotone if $\langle F(x) - F(y), x - y \rangle \geq 0$; μ -strongly monotone if, for some $\mu > 0$, $\langle F(x) - F(y), x - y \rangle \geq \mu \|x - y\|^2$. The following result will be useful to apply our stability-based arguments:

Lemma 2 ([18]): Let $F : \mathbb{R}^n \rightarrow \mathbb{R}^n$ be an μ -strongly monotone and κ -Lipschitz continuous mapping, with $\kappa, \mu > 0$. If $\gamma \in (0, 2\mu/\kappa^2)$, then it holds that $\|x - \gamma F(x) - (y - \gamma F(y))\| \leq \rho \|x - y\|$, with $\rho := \sqrt{1 - 2\gamma\mu + \gamma^2\kappa^2} < 1$. $\square$

II. PROBLEM FORMULATION

We start by introducing SGNEPs and discussing a popular reformulation based on barrier functions for imposing (soft) constraint satisfaction. Then, we formalize the considered SGNE seeking problem with limited data availability.

A. Stochastic generalized Nash equilibrium problems

We consider a set $\mathcal{I} = \{1, \dots, N\}$ indexing N noncooperative agents taking part to a Nash game. In particular, each of them chooses its strategy $x_i \in \mathbb{R}^{n_i}$, aiming at minimizing its local cost function $J_i : \mathbb{R}^{n_i} \times \mathbb{R}^{n-n_i} \rightarrow \mathbb{R}$, $n := \sum_{i \in \mathcal{I}} n_i$, which also depends on the decisions of the i -th agent’s opponents $\mathbf{x}_{-i} \in \mathbb{R}^{n-n_i}$. Instead, $\mathbf{x} \in \mathbb{R}^n$ represents the collective decision strategy of the population of agents $\mathcal{I}$, taking values in the set $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : g(\mathbf{x}) \leq \mathbf{0}_m\}$, where $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$. We indicate the feasible set of agent i with $\mathcal{X}_i(\mathbf{x}_{-i}) \subseteq \mathbb{R}^{n_i}$, which explicitly accounts for the possible dependency on the other participants’ decision variables. Local constraints are, instead, implicitly included in $\mathcal{X}$.

Standing Assumption 1: Let $g(\mathbf{x}) = \text{col}((g_j(\mathbf{x}))_{j=1}^m)$, with each $g_j : \mathbb{R}^n \rightarrow \mathbb{R}$. The following conditions hold:

- (i) Each $\mathbf{x} \mapsto g_j(\mathbf{x})$ is convex and of class $\mathcal{C}^2$;
- (ii) The set $\mathcal{X}$ is nonempty, convex, compact, and satisfies Slater’s constraint qualification. $\square$

These assumptions are standard for SGNEPs [4], [19], [20]. Moreover, we assume each agents’ cost function in the form:

$$J_i(x_i, \mathbf{x}_{-i}) := \mathbb{E}_{\mathbb{P}}[f_i(x_i, \mathbf{x}_{-i}, \xi(\omega))], \quad (3)$$

where $f_i : \mathbb{R}^{n_i} \times \mathbb{R}^{n-n_i} \times \mathbb{R}^d \rightarrow \mathbb{R}$. The variable $\xi : \Xi \rightarrow \mathbb{R}^d$ represents a random quantity affecting each local cost, fully characterized through the probability space $(\Xi, \mathcal{F}, \mathbb{P})$, while $\mathbb{E}_{\mathbb{P}}$ denotes the expectation w.r.t. the distribution $\mathbb{P}$ of ξ , which is assumed here to be *unknown*. We make the blanket assumption that $\mathbb{E}_{\mathbb{P}}[f_i(\mathbf{x}, \xi)]$ is well defined for all $\mathbf{x} \in \mathcal{X}$.

With some abuse of notation, from now on we simply use ξ to indicate the sample $\xi(\omega)$, as common in the literature [8].

Thus, each agent wishes to solve its local optimization problem, thereby originating a so-called SGNEP:

$$\forall i \in \mathcal{I} : \begin{cases} \min_{x_i \in \mathbb{R}^{n_i}} & \mathbb{J}_i(x_i, \mathbf{x}_{-i}) \\ \text{s.t.} & g(x_i, \mathbf{x}_{-i}) \leq \mathbf{0}_m. \end{cases} \quad (4)$$

Next, we formalize a popular solution for SGNEPs [3], [4]:

Definition 2: An SGNE is a collective decision strategy $\mathbf{x}^* \in \mathcal{X}$ such that, for all $i \in \mathcal{I}$,

$$\mathbb{J}_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*) \leq \mathbb{J}_i(y_i, \mathbf{x}_{-i}^*) \text{ for all } y_i \in \mathcal{X}_i(\mathbf{x}_{-i}^*).$$

□

According to [1, Prop. 3.5], we make now few assumptions ensuring that at least an SGNE exists for the SGNEP in (4):

Standing Assumption 2: The following conditions hold:

- (i) Each $x_i \mapsto f_i(x_i, \mathbf{x}_{-i}, \xi)$ is convex, of class $\mathcal{C}^1$, and Lipschitz continuous with integrable Lipschitz constant w.r.t. ξ , while $\xi \mapsto f_i(x_i, \mathbf{x}_{-i}, \xi)$ is measurable;
- (ii) Each $x_i \mapsto \mathbb{J}_i(x_i, \mathbf{x}_{-i})$ is convex and of class $\mathcal{C}^1$. □

Among all possible SGNE, we focus on those corresponding to the solution set of the SVI associated to the SGNEP. To characterize them, let us first define the pseudogradient mapping $\mathbb{F} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$ related to each cost function:

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

Thus, the associated SVI problem is formalized as:

$$\text{Find } \mathbf{x}^* \in \mathcal{X} \text{ such that } \langle \mathbb{F}(\mathbf{x}^*), \mathbf{x} - \mathbf{x}^* \rangle \geq 0 \forall \mathbf{x} \in \mathcal{X}. \quad (6)$$

While being a solution to the SVI above is necessary and sufficient to be a stochastic Nash equilibrium (SNE) of a game with no coupling constraints, the set of (v-SGNE) of (6) is a subset of all the equilibria of the SGNEP in (4) [21]. We, however, use v-SGNE and SGNE interchangeably in the following. To impose constraint satisfaction while turning the SGNEP in (4) into an unconstrained stochastic Nash equilibrium problem (SNEP), a popular approach relies on barrier functions [19], [20], [22]. We can then adjust each cost function in (4) by adding a barrier term as follows:

$$\forall i \in \mathcal{I} : \min_{x_i \in \mathbb{R}^{n_i}} \mathbb{J}_i(x_i, \mathbf{x}_{-i}) + \phi(\mathbf{x}), \quad (7)$$

where the function ϕ is a log-barrier function, with $\varrho > 0$:

$$\phi(\mathbf{x}) = \begin{cases} -\varrho \sum_{j=1}^m \log(-g_j(\mathbf{x})), & g(\mathbf{x}) < \mathbf{0}_m \\ +\infty, & \text{otherwise.} \end{cases} \quad (8)$$

From Standing Assumption 1, $\mathbf{x} \mapsto \phi(\mathbf{x})$ is convex and of class $\mathcal{C}^2$ on $\text{int}(\mathcal{X})$, and hence its gradient $\nabla \phi(\mathbf{x})$ is monotone and locally Lipschitz with constant ℓ_ϕ [19], [20].

Remark 1: Solving (7) typically provides an approximate solution of the SGNEP in (4) [20, Lemma 4] while for ϱ shrinking to zero, the SNE of (7) converges to an SGNE of (4) [19, Lemma 2]. In practice, however, setting a small enough $\varrho > 0$ is usually sufficient to recover an SGNE of the SGNEP, and it is actually what we will implicitly assume throughout the paper. This is why, when referring

to a solution to the unconstrained SNEP in (7), we will equivalently refer to an SGNE of (4). □

The resulting pseudogradient mapping for the unconstrained SNEP in (7), $\mathbb{B} : \mathbb{R}^n \rightrightarrows \mathbb{R}^n$, formally reads as:

$$\mathbb{B}(\mathbf{x}) = \mathbb{F}(\mathbf{x}) + \nabla \phi(\mathbf{x}). \quad (9)$$

Standing Assumption 3: The following conditions hold:

- (i) $\mathbb{B}$ is μ -strongly monotone;
- (ii) $\mathbb{B}$ is κ -Lipschitz continuous;
- (iii) $\exists M > 0$ such that $\|\mathbb{B}(\mathbf{x})\| \leq M$, for all $\mathbf{x} \in \text{int}(\mathcal{X})$. □

Remark 2: The first requirement holds, for instance, if the pseudogradient mapping $\mathbb{F}$ in (5) is strongly monotone, since $\nabla \phi(\mathbf{x})$ is monotone. Under this assumption on $\mathbb{F}$, the solution to (6) is also unique [23]. Similarly, Standing Assumption 3.(ii) is a consequence of Standing Assumption 1, and possibly the Lipschitz continuity of $\mathbb{F}$ [10], [19], [20]. Standing Assumption 3.(iii), instead, holds for instance in case of polyhedral constraints, i.e., $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^n : A\mathbf{x} \leq b\}$, for $A \in \mathbb{R}^{m \times n}$ and $b \in \mathbb{R}^m$. Collectively, these assumptions are common in the (S)GNE seeking literature [3], [4]. □

B. A framework with limited data availability

According to [23], if $\mathbf{x}^*$ is an SGNE of (7), then $\mathbb{B}(\mathbf{x}^*) = \mathbf{0}$. Thus, we define the set SGNE_ε , which consist of all those points ε -close to an SGNE of (7) (and hence of (4)), as:

$$\text{SGNE}_\varepsilon := \{\mathbf{x} \in \mathcal{X} : \|\mathbb{B}(\mathbf{x})\| \leq \varepsilon\}.$$

As common in SGNEPs, the mapping $\mathbb{F}(\cdot)$ is either unavailable, or quite challenging to calculate, because it involves the expected value $\mathbb{E}_{\mathbb{P}}[\cdot]$, i.e., a multidimensional integration that depends on the distribution $\mathbb{P}$ of ξ which might be unknown. Therefore, by assuming to have available only a few realizations of the uncertain variable we approximate $\mathbb{B}(\cdot)$ in a data-driven fashion by means of an oracle $O : \Omega \times \Xi \rightarrow \mathbb{R}^n$, assumed as a Borel function satisfying:

$$\text{For all } \mathbf{x} \in \Omega, \mathbb{E}_{\mathbb{P}}[O(\mathbf{x}, \xi)] = \mathbb{B}(\mathbf{x}).$$

In the remainder, we tacitly assume that the oracle O possesses the same properties postulated for the pseudogradient operator $\mathbb{B}$ in (9) [24], [25]. Relying on an oracle is a rather common assumption for stochastic algorithms and essentially means that it is possible to compute the pseudogradient given one sample [8], [9], [13]. In our analysis, we hence leverage the following data-based approximation for $\mathbb{B}(\cdot)$:

$$\hat{B}_s(\mathbf{x}) = \hat{B}(\mathbf{x}, \mathcal{D}_s) := \frac{1}{s} \sum_{r=1}^s O(\mathbf{x}, \xi^{(r)}), \quad (10)$$

which depends on s i.i.d. samples made available from the set Ξ , and collected into $\mathcal{D}_s := \{\xi^{(r)}\}_{r=1}^s \in \Xi^s$.

The question we wish to answer is then the following: given a finite dataset, are we able to quantify through an explicit calculation of some $\varepsilon > 0$, how far we can get from an SGNE by employing $\hat{B}_s$ in place of the actual $\mathbb{B}$?

Algorithm 1: Approximated EG method**Initialization:** Each x_i^0 so that $\mathbf{x}^0 \in \text{int}(\mathcal{X})$ **Iteration** $k \in \mathbb{N}_0$: Agent i 1) Receives $\mathbf{x}_{-i}^k$, then updates

$$\mathbf{y}_i^k = \mathbf{x}_i^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}_i^k, \mathbf{x}_{-i}^k, \xi^{(r)})$$

2) Receives $\mathbf{y}_{-i}^k$, then updates

$$\mathbf{x}_i^{k+1} = \mathbf{x}_i^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{y}_i^k, \mathbf{y}_{-i}^k, \xi^{(r)})$$

III. SGNE SEEKING UNDER THE LENS OF ALGORITHMIC STABILITY

First, we introduce the EG scheme applied to SGNE seeking, and the data-driven version adapted for our analysis. Next, we show its stability properties w.r.t. a tailored loss function, which is crucial for proving our main result.

A. The extragradient method

Our analysis focuses on the procedure reported in Algorithm 1, which consists in the distributed and data-driven version of the following EG in compact form [26, §7.2]:

$$\begin{cases} \mathbf{y}^k = \mathbf{x}^k - \gamma \mathbb{B}(\mathbf{x}^k) \\ \mathbf{x}^{k+1} = \mathbf{x}^k - \gamma \mathbb{B}(\mathbf{y}^k). \end{cases} \quad (11)$$

The expected value required to evaluate $\mathbb{B}$ can then be replaced through a data-driven approximation of the pseudogradient mapping as in (10). To simplify the analysis we have adopted the same, constant step-size $\gamma > 0$ and the same dataset $\mathcal{D}_s$ for all the agents, although a generalization to local, possibly time-varying $\{\gamma_i^k\}_{i \in \mathcal{I}}$ and $\mathcal{D}_s$ is doable.

Remark 3: We note that the scheme in (11) converges almost surely to an SGNE when the true operator $\mathbb{B}(\cdot)$ is at least pseudomonotone and Lipschitz continuous, and an increasing batch $s = s^k \rightarrow \infty$ is available at every iteration [8, Th. 3.18], [10, Th. 1]. Thus, if we knew $\mathbb{B}(\cdot)$, the conditions in Standing Assumption 3 would suffice for the convergence of (11). $\square$

In a data-limited context our goal is hence to investigate how far we can get from the (unique) v-SGNE by running Algorithm 1 for an arbitrary number of iterations K . Specifically, we want to characterize the resulting $\mathbf{x}^{K+1}$ as an ε -SNE of the SNEP in (7). To this end, we will make use of the algorithmic stability arguments discussed in §I-A.

B. A computable bound for ε -SGNE

By replacing $\mathbb{B}$ with the approximation $\hat{B}_s$ in (10), the resulting data-driven variant then proposes the distributed updates in Algorithm 1. According to §I-A, $\omega := \text{col}(\mathbf{y}, \mathbf{x}) \in \mathbb{R}^{2n}$ then amounts to the hypothesis of our algorithm. Under

the postulated assumptions we will prove next that Algorithm 1 is uniformly stable w.r.t. the following loss function, which measures the performance of a generic hypothesis H :

$$\begin{aligned} \ell(H, \xi) &= \|[0_n \ I_n]H - \gamma O([0_n \ I_n]H, \xi)\| \\ &= \|\mathbf{x} - \gamma O(\mathbf{x}, \xi)\|. \end{aligned} \quad (12)$$

In particular, we will obtain a stability parameter $\beta \propto (1/s)$, thereby making the probabilistic condition in Lemma 1 tight.

Lemma 3: Let $\gamma \in (0, 2\mu/\kappa^2)$ be chosen so that $\tau = \tau(\gamma) := (1 + \gamma\kappa)\rho + \gamma\kappa < 1$. Then, the data-driven EG in Algorithm 1 possesses uniform stability proportional to $1/s$ w.r.t. the loss function $\ell(H, \xi)$ in (12). $\square$

Proof: Fix some $s \geq 1$, and consider two datasets, $\mathcal{D}_s, \mathcal{D}_s^t \in \Xi^s$, both consisting of s -i.i.d. random samples and differing on the t -th instance only, i.e., some ξ^t replaces $\xi^{(t)}$ in $\mathcal{D}_s^t$. After iterating (a compact version of) Algorithm 1 for $K \geq 1$ times starting from some $\mathbf{x}^0 \in \text{int}(\mathcal{X})$, we obtain $\omega^{K+1} = \text{col}(\mathbf{y}^K, \mathbf{x}^{K+1})$ and $\tilde{\omega}^{K+1} = \text{col}(\tilde{\mathbf{y}}^K, \tilde{\mathbf{x}}^{K+1})$ as the two hypotheses leveraging samples in $\mathcal{D}_s \in \Xi^s$ and $\mathcal{D}_s^t \in \Xi^s$, respectively. Once picked an arbitrary $\xi \in \Xi$, we have:

$$\begin{aligned} &|\ell(H_s, \xi) - \ell(H_{s^t}, \xi)| \\ &\stackrel{(a)}{\leq} \|\mathbf{x}^{K+1} - \gamma O(\mathbf{x}^{K+1}, \xi) - \tilde{\mathbf{x}}^{K+1} + \gamma O(\tilde{\mathbf{x}}^{K+1}, \xi)\| \\ &\stackrel{(b)}{\leq} \rho \|\mathbf{x}^{K+1} - \tilde{\mathbf{x}}^{K+1}\|, \end{aligned} \quad (13)$$

where (a) follows by the reverse triangle inequality, while (b) from Lemma 2, $\rho \in (0, 1)$. It remains to bound $\|\mathbf{x}^{K+1} - \tilde{\mathbf{x}}^{K+1}\|$. To this end, consider Algorithm 1 at the generic k -th iteration, which produces $\mathbf{x}^{k+1}$ and $\tilde{\mathbf{x}}^{k+1}$ such that:

$$\begin{aligned} &\|\mathbf{x}^{k+1} - \tilde{\mathbf{x}}^{k+1}\| \\ &= \|\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}^k, \xi^{(r)}), \xi^{(r)}) \\ &\quad - (\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\tilde{\mathbf{x}}^k, \xi^{(r)}), \xi^{(r)})\| \\ &\leq \|\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}^k, \xi^{(r)}), \xi^{(r)}) \\ &\quad - (\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\tilde{\mathbf{x}}^k, \xi^{(r)}), \xi^{(r)})\| \\ &\quad + \frac{\gamma}{s} \|O(\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}^k, \xi^{(r)}), \xi^{(t)})\| \\ &\quad + \frac{\gamma}{s} \|O(\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\tilde{\mathbf{x}}^k, \xi^{(r)}), \xi^{(t)})\| \\ &\leq \|\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\mathbf{x}^k, \xi^{(r)}), \xi^{(r)}) \\ &\quad - (\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1}^s O(\tilde{\mathbf{x}}^k, \xi^{(r)}), \xi^{(r)})\| \\ &\quad + \frac{2\gamma M}{s}, \end{aligned}$$

where the last inequality follows by virtue of Standing Assumption 3.(iii). Note that the arguments of both outer operators $O(\cdot)$ coincide with the gradient descent update evaluated from $\mathbf{x}^k$ and $\tilde{\mathbf{x}}^k$ with the same pool of samples, denoted as $\mathbf{z}^{k+1} = \mathbf{x}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\mathbf{x}^k, \xi^{(r)})$ and $\tilde{\mathbf{z}}^{k+1} = \tilde{\mathbf{x}}^k - \frac{\gamma}{s} \sum_{r=1, r \neq t}^s O(\tilde{\mathbf{x}}^k, \xi^{(r)})$, respectively. Thus, after summing and subtracting both $\mathbf{z}^{k+1}$ and $\tilde{\mathbf{z}}^{k+1}$ inside the norm above, standard manipulations yield:

$$\begin{aligned} \|\mathbf{x}^{k+1} - \tilde{\mathbf{x}}^{k+1}\| &\leq \|\mathbf{z}^{k+1} - \tilde{\mathbf{z}}^{k+1}\| + \frac{\gamma}{s} \sum_{r=1}^s \|O(\mathbf{z}^{k+1}, \xi^{(r)}) \\ &\quad - O(\tilde{\mathbf{z}}^{k+1}, \xi^{(r)}) + O(\mathbf{x}^k, \xi^{(r)}) - O(\tilde{\mathbf{x}}^k, \xi^{(r)})\| + \frac{2\gamma M}{s}, \\ &\leq (1 + \gamma\kappa) \|\mathbf{z}^{k+1} - \tilde{\mathbf{z}}^{k+1}\| + \gamma\kappa \|\mathbf{x}^k - \tilde{\mathbf{x}}^k\| + \frac{2\gamma M}{s}. \end{aligned}$$

By relying on the growth lemma [25, Lemma 2], we can then upper bound the last relation, particularly the term $\|\mathbf{z}^{k+1} - \tilde{\mathbf{z}}^{k+1}\|$, as follows: $\|\mathbf{x}^{k+1} - \tilde{\mathbf{x}}^{k+1}\| \leq ((1 + \gamma\kappa)\rho + \gamma\kappa)\|\mathbf{x}^k - \tilde{\mathbf{x}}^k\| + \frac{2\gamma M}{s}$, where $\rho < 1$ is the contraction factor characterizing the gradient descent step in Lemma 2 for $\gamma \in (0, 2\mu/\kappa^2)$. Then, recursing over the K -steps, and considering that both trajectories leading to $\mathbf{x}^{K+1}$ and $\tilde{\mathbf{x}}^{K+1}$ originate from the same initial conditions leaves us with:

$$\begin{aligned} \|\mathbf{x}^{K+1} - \tilde{\mathbf{x}}^{K+1}\| &\leq \frac{2\gamma M}{s} \sum_{k=1}^K ((1 + \gamma\kappa)\rho + \gamma\kappa)^{K-k} \\ &\leq \frac{2\gamma M}{s(1 - \gamma\kappa - (1 + \gamma\kappa)\rho)}, \end{aligned}$$

where we have exploited the relation $\sum_{k=1}^K \alpha^{K-k} = \frac{1 - \alpha^K}{1 - \alpha} < \frac{1}{1 - \alpha}$ which holds true for any $\alpha \in [0, 1)$, i.e., $(1 + \gamma\kappa)\rho + \gamma\kappa = \tau < 1$. Finally, combining the latter inequality with (13) allows us to conclude that:

$$|\ell(H_s, \xi) - \ell(H_{s^t}, \xi)| \leq \frac{2\gamma\rho M}{s(1 - \tau)},$$

thereby proving uniform stability proportional to $1/s$. ■

Remark 4: Lemma 3, and later Theorem 1, assume the existence of some γ so that $(1 + \gamma\kappa)\rho + \gamma\kappa < 1$. The latter condition ensures a sort of global contraction for the data-driven EG scheme. While establishing explicit conditions on γ guaranteeing $\tau < 1$ is hard, one can numerically observe that meeting such a condition is possible, making the resulting assumption reasonable—see Appendix A. □

Then, equipped with this uniform stability result, we are now ready to state and prove what follows:

Theorem 1: Fix $s \geq 1$, $\delta \in (0, 1)$, and let $\gamma \in (0, 2\mu/\kappa^2)$ be chosen so that $\tau(\gamma) < 1$. Given any dataset $\mathcal{D}_s \in \Xi^s$, there exists a computable $\varepsilon = \varepsilon(s, \delta) > 0$ such that $\mathbf{x}^{K+1} \in \text{SGNE}_\varepsilon$ with probability at least $1 - \delta$, where $\mathbf{x}^{K+1}$ is a subvector of ω^{K+1} obtained by iterating the data-driven EG in Algorithm 1 K -times starting from $\mathbf{x}^0 \in \text{int}(\mathcal{X})$. □

Proof: After iterating the instructions reported in Algorithm 1 K -times, consider the following relations:

$$\begin{aligned} \|\mathbf{x}^* - \gamma\mathbb{B}(\mathbf{x}^*)\| &= 0 \leq \|\mathbf{x}^{K+1} - \gamma\mathbb{B}(\mathbf{x}^{K+1})\| \\ &\stackrel{(a)}{=} \|\mathbf{x}^{K+1} - \gamma\mathbb{E}_{\mathbb{P}}[O(\mathbf{x}^{K+1}, \xi)]\| \\ &\stackrel{(b)}{=} \|\mathbb{E}_{\mathbb{P}}[\mathbf{x}^{K+1} - \gamma O(\mathbf{x}^{K+1}, \xi)]\| \\ &\stackrel{(c)}{\leq} \mathbb{E}_{\mathbb{P}}[\|\mathbf{x}^{K+1} - \gamma O(\mathbf{x}^{K+1}, \xi)\|] \\ &= \mathbb{E}_{\mathbb{P}}[\ell(H_s, \xi)] = r(\mathcal{A}, s), \end{aligned}$$

where we have used in (a) the fact that the oracle operator is unbiased, in (b) the fact that the data-driven EG in Algorithm 1 yields a deterministic output ω^{K+1} once fixed the dataset $\mathcal{D}_s$, as well as the linearity of the expected value, and in (c) the standard Jensen's inequality. Thus, upper bounding the risk $\mathbb{E}_{\mathbb{P}}[\ell(H_s, \xi)]$ provides an equivalent information on the distance of $\mathbf{x}^{K+1}$ from a fixed point of the dynamics produced by iterating the operator $\text{Id} - \gamma\mathbb{B} \circ (\text{Id} - \gamma\mathbb{B})$ with perfect knowledge on $\mathbb{B}$. Therefore, with probability at least $1 - \delta$, the direct application of the




Fig. 1: Box plots capturing the relative approximation error $\|\mathbf{x}^{K+1} - \mathbf{x}^*\|$ produced by different run K of Algorithm 1 over 20 trials with fixed number of samples $s = 1000$. The blue down-pointing triangles denote the resulting averaged bounds ε offered by Theorem 1 with $\delta = 0.05$.

bound in Lemma 3 leads to $\|\mathbf{x}^{K+1} - \gamma\mathbb{B}(\mathbf{x}^{K+1})\| \leq \varepsilon$ with $\varepsilon := \sum_{r=1}^s \|\mathbf{x}^{K+1} - \gamma O(\mathbf{x}^{K+1}, \xi^{(r)})\| / s + \beta + (s\beta + M)\sqrt{2\ln(1/\delta)/s}$ and $\beta = 2\gamma\rho M/s(1 - \tau)$. ■

Thus, Theorem 1 provides a computable bound, which holds with arbitrarily high confidence $1 - \delta$ w.r.t. any possible dataset drawn from Ξ^s and probability distribution $\mathbb{P}$ underlying these data, on the distance between the output of Algorithm 1 after K steps, and an SGNE of the SGNEP in (4). Note that the specific choice of the loss function in (12) makes the calculation of ε possible a-posteriori (i.e., after K iterations). In fact, all of the terms $\|\mathbf{x}^{K+1} - \gamma O(\mathbf{x}^{K+1}, \xi^{(r)})\|$ do not explicitly depend on unknown quantities, e.g., the true SGNE $\mathbf{x}^*$, while the other terms depend on design parameters (δ, γ), problem data (s, M, μ, κ), or both (ρ, τ).

IV. ILLUSTRATIVE EXAMPLE

We now test our bounds on an academic example involving $N = 20$ agents controlling an unconstrained scalar decision variable (i.e., $n_i = 1$) to minimize a quadratic cost function $f_i(x_i, \mathbf{x}_{-i}, \xi) = a_i x_i^2 + (b_i + \xi)x_i + (\sum_{j \in \mathcal{I} \setminus \{i\}} c_{i,j} x_j)x_i$. Specifically, each parameter in the cost is chosen randomly, i.e., $a_i \sim \mathcal{U}(0.2, 0.8)$, $b_i \sim \mathcal{U}(0.02, 0.075)$, $c_{i,j} \sim \mathcal{N}(0, 0.01)$, while the stochastic variable ξ is distributed according to $\mathcal{N}(-1, 0.2)$. Note that the results from the previous section do not directly involve quantities related to the collective strategy profile $\mathbf{x}$, e.g., $n = \sum_{i \in \mathcal{I}} n_i$, nor $\nabla\phi(\cdot)$, and hence the constraints – see the discussion in Remark 2. This explains the unconstrained nature of the considered game with scalar strategies, which has the sole purpose of computing the proposed bounds numerically. In all the instances, the true SNE $\mathbf{x}^*$ is computed via [3, Alg. 1].

We start by investigating the behaviour of Algorithm 1 by fixing the number of samples s and varying K . Specif-




Fig. 2: Box plots capturing the relative approximation error $\|x^{101} - x^*\|$ produced by $K = 100$ iterations of Algorithm 1 over 20 trials considering different dataset size. The blue down-pointing triangles denote the resulting averaged bounds ε offered by Theorem 1 with $\delta = 0.05$.

ically, we perform 20 trials for each value of $K \in \{60, 80, \dots, 1000\}$, and populate the dataset $\mathcal{D}_s$ with $s = 1000$ samples. In Fig. 1 we illustrate the numerical results obtained with confidence level $\delta = 0.05$ and step-size γ chosen to meet the condition in Theorem 1. While, in practice, Algorithm 1 well approximates the unique SNE when run for $K \geq 300$, the theoretical bound established in Theorem 1 reveals an intrinsic dependency on K itself. In fact, the averaged bound obtained appears dominated by the empirical error term, which progressively reduces its impact with larger K , settling around 0.35. By running Algorithm 1 for $K = 1000$ iterations with $s = 1000$ samples, for instance, one can then claim that the data-driven EG produces an error in the approximation of x^* of at most 35% with probability of at least 95%. Note that this holds regardless the distribution of ξ and the dataset $\mathcal{D}_{1001}$ at hand.

Successively, we compare the theoretical bound ε provided by Theorem 1 with the actual distance $\|x^{1001} - x^*\|$ (both scaled by $\|x^*\|$) by considering a small number of iterations $K = 101$ and several values for s . Also in this case, we perform 20 trials for each $s \in \{500, 600, \dots, 3000\}$. In Fig. 2, we note that the relative error produced by the data-driven EG in Algorithm 1, measured by the box plots, is quite significant (around 30%), yet almost equivalent for the different s considered. Consistently, the theoretical bound ε decreases as the number of samples adopted increases.

As a concluding remark, we note that, in general, the proposed certificates are not expected to be tight. The resulting numerical gap between the theoretical bound and the empirical evaluation—approximately one order of magnitude in the worst case, as seen in Fig. 1—is quite common for these type of probabilistic guarantees holding with high-confidence. Similar results have been observed in other domains, e.g.,

robust control design [27], or robustness evaluation [28].

V. CONCLUSION

We have derived rigorous probabilistic guarantees on the quality of the solution produced by a data-driven version of an EG scheme applied to SGNE seeking with limited data availability. Specifically, we have leveraged a finite number of uncertainty realizations to construct a sample average-based approximation for the game mapping, and exploited algorithmic stability arguments to bound the distance between the true SGNE and the solution obtained through the EG method. A distinct feature of our bounds is that they i) are easy to compute, ii) hold with high confidence, and iii) remain valid regardless of the underlying data distribution.

Future work will explore alternative techniques for solving SGNEPs under the lens of algorithmic stability, potentially incorporating nonlinear operators to better handle both local and coupling constraints. Besides extending results for the EG considered here, prominent examples include algorithms based on standard operator splitting methods, such as forward-backward-forward or Douglas-Rachford.

DECLARATION OF INTERESTS

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

REFERENCES

- [1] 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.
- [2] J. Koshal, A. Nedić, 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.
- [3] B. Franci and S. Grammatico, “A distributed forward-backward algorithm for stochastic generalized Nash equilibrium seeking,” *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5467–5473, 2020.
- [4] —, “Stochastic generalized Nash equilibrium-seeking in merely monotone games,” *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 3905–3919, 2021.
- [5] 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.
- [6] 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, pp. 1–42, 2013.
- [7] 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)*. IEEE, 2022, pp. 4455–4460.
- [8] A. N. 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.
- [9] A. Alacaoglu and Y. Malitsky, “Stochastic variance reduction for variational inequality methods,” in *Conference on Learning Theory*. PMLR, 2022, pp. 778–816.
- [10] B. Franci, “On variance-reduced extragradient methods for stochastic generalized Nash equilibrium problems,” *IEEE Control Systems Letters*, 2025.
- [11] G. M. Korpelevich, “The extragradient method for finding saddle points and other problems,” *Matecon*, vol. 12, pp. 747–756, 1976.
- [12] O. Bousquet and A. Elisseeff, “Stability and generalization,” *The Journal of Machine Learning Research*, vol. 2, pp. 499–526, 2002.

- [13] F. Fabiani and B. Franci, "Finite-sample guarantees for data-driven forward-backward operator methods," *IEEE Transactions on Automatic Control*, pp. 1–12, 2026.
- [14] R. I. Boş, P. Mertikopoulos, M. Staudigl, and P. T. Vuong, "Minibatch forward-backward-forward methods for solving stochastic variational inequalities," *Stochastic Systems*, vol. 11, no. 2, pp. 112–139, 2021.
- [15] R. M. Gower, M. Schmidt, F. Bach, and P. Richtárik, "Variance-reduced methods for machine learning," *Proceedings of the IEEE*, vol. 108, no. 11, pp. 1968–1983, 2020.
- [16] M. Vidyasagar, *Learning and generalisation: With applications to neural networks*. Springer Science & Business Media, 2013.
- [17] S. Kutin and P. Niyogi, "Almost-everywhere algorithmic stability and generalization error," in *Proceedings of the 18th conference on Uncertainty in Artificial Intelligence*, 2002, pp. 275–282.
- [18] J. Lei and U. V. Shanbhag, "Distributed variable sample-size gradient-response and best-response schemes for stochastic Nash equilibrium problems," *SIAM Journal on Optimization*, vol. 32, no. 2, pp. 573–603, 2022.
- [19] A. R. Romano and L. Pavel, "Exact convergence to GNE using penalty methods," in *2022 European Control Conference (ECC)*. IEEE, 2022, pp. 2285–2290.
- [20] —, "An inexact-penalty method for GNE seeking in games with dynamic agents," *IEEE Transactions on Control of Network Systems*, vol. 10, no. 4, pp. 2072–2084, 2023.
- [21] 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.
- [22] F. Fabiani and A. Caiti, "Nash equilibrium seeking in potential games with double-integrator agents," in *2019 18th European Control Conference (ECC)*. IEEE, 2019, pp. 548–553.
- [23] F. Facchinei and J.-S. Pang, *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.
- [24] M. Hardt, B. Recht, and Y. Singer, "Train faster, generalize better: Stability of stochastic gradient descent," in *International Conference on Machine Learning*. PMLR, 2016, pp. 1225–1234.
- [25] F. Farnia and A. Ozdaglar, "Train simultaneously, generalize better: Stability of gradient-based minimax learners," in *International Conference on Machine Learning*. PMLR, 2021, pp. 3174–3185.
- [26] E. K. Ryu and S. Boyd, "Primer on monotone operator methods," *Applied and Computational Mathematics*, vol. 15, no. 1, pp. 3–43, 2016.
- [27] G. Schildbach, L. Fagiano, C. Frei, and M. Morari, "The scenario approach for stochastic model predictive control with bounds on closed-loop constraint violations," *Automatica*, vol. 50, no. 12, pp. 3009–3018, 2014.
- [28] F. Fabiani, K. Margellos, and P. J. Goulart, "Probabilistic feasibility guarantees for solution sets to uncertain variational inequalities," *Automatica*, vol. 137, p. 110120, 2022.

APPENDIX

A. On the condition ensuring contraction

In Fig. 3 we report the binary values associated to the feasibility of the nonlinear optimization problem for the design of an optimal γ satisfying the condition for Theorem 1, i.e.,

$$\begin{aligned} \min_{\gamma \in (0, 2\mu/\kappa^2)} \quad & -\gamma \\ \text{s.t.} \quad & (1 + \gamma\kappa)\sqrt{1 - 2\gamma\mu + \gamma^2\kappa^2} + \gamma\kappa < 1, \end{aligned} \quad (14)$$

for 100 randomly sampled values of $\kappa \sim \mathcal{U}(10^{-6}, 50)$ and $\mu \sim \mathcal{U}(10^{-6}, 50)$. From Fig. 3, the feasible area covers around the 80.6% of the considered cases. Therefore, assuming the existence of some γ ensuring $\tau(\gamma) < 1$ appears a reasonable condition at the basis of our analysis.




Fig. 3: Map showing the feasibility (blue area), infeasibility (green area), or reached solver limits (yellow area) of (14) for 10^4 random combinations of $(\kappa, \mu) \sim \mathcal{U}(10^{-6}, 50)^2$.