

## Finite-sample guarantees for data-driven forward-backward operator methods

Questa è la versione preprint della seguente opera:

*Original*

Finite-sample guarantees for data-driven forward-backward operator methods / Fabiani, Filippo; Franci, Barbara. - In: IEEE TRANSACTIONS ON AUTOMATIC CONTROL. - ISSN 1558-2523. - (2026).  
[10.1109/TAC.2026.3660169]

*Availability:*

This version is available at: 20.500.11771/38178

*Publisher:*

*Published*

DOI:10.1109/TAC.2026.3660169

*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](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)

# Finite-sample guarantees for data-driven forward-backward operator methods

Filippo Fabiani and Barbara Franci

## Abstract

We establish finite sample certificates on the quality of solutions produced by data-based forward-backward (FB) operator splitting schemes. As frequently happens in stochastic regimes, we consider the problem of finding a zero of the sum of two operators, where one is either unavailable in closed form or computationally expensive to evaluate, and shall therefore be approximated using a finite number of noisy oracle samples. Under the lens of algorithmic stability, we then derive probabilistic bounds on the distance between a true zero and the FB output without making specific assumptions about the underlying data distribution. We show that under weaker conditions ensuring the convergence of FB schemes, stability bounds grow proportionally to the number of iterations. Conversely, stronger assumptions yield stability guarantees that are independent of the iteration count. We then specialize our results to a popular FB stochastic Nash equilibrium seeking algorithm and validate our theoretical bounds on a control problem for smart grids, where the energy price uncertainty is approximated by means of historical data.

## Index Terms

Data-driven methods, Robust decision-making, Operator splitting methods, Stochastic optimization.

F. Fabiani is with the IMT School for Advanced Studies Lucca, Piazza San Francesco 19, 55100, Lucca, Italy ([filippo.fabiani@imtlucca.it](mailto: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](mailto:barbara.franci@polito.it)).

## I. INTRODUCTION

Finding a pervasive application in optimization and game theory, operator splitting methods can be generally employed to construct zeros of sums of operators as fixed point iterations [1], [2]. Specifically, they offer an elegant framework and systematic tools for solving problems of the form:

$$\text{Find } \omega^* \in \Omega \text{ s.t. } 0 \in \mathcal{A}(\omega^*) + \mathcal{B}(\omega^*),$$

with decision set  $\Omega \subseteq \mathbb{R}^n$ , and mappings  $\mathcal{A} : \Omega \rightarrow 2^\Omega$ ,  $\mathcal{B} : \Omega \rightarrow 2^\Omega$ . We then say that  $\omega^*$  is a *zero* of  $\mathcal{A} + \mathcal{B}$ , i.e., it belongs to the set  $\text{zer}(\mathcal{A} + \mathcal{B}) := \{\omega \in \mathbb{R}^n : 0 \in \mathcal{A}(\omega) + \mathcal{B}(\omega)\}$ . In this paper, we will specifically investigate the case in which the operator  $\mathcal{A}$  is available and its evaluation, or that of the associated resolvent  $J_{\mathcal{A}} := (\text{Id} + \mathcal{A})^{-1}$  (e.g., a proximity operator), is reasonably cheap, while the operator  $\mathcal{B}$  is either unavailable, or quite challenging to evaluate. The latter shall then be *approximated* in a data-driven fashion by relying on a *finite* number of data.

### A. Data-driven approximation methods and related work

Especially in stochastic regimes, several optimization problems indeed require the evaluation of an expected value mapping, a quantity that can either be hard to compute, or even inaccessible in case the probability distribution of the random variable is unknown. As such, several data-driven approximation schemes have been proposed in the literature [3]–[14]. The latter essentially assume that some batch of random, independent and identically distributed (i.i.d.) samples is available, and then leverage approximation procedures that roughly fall within two large umbrellas: sample average approximation (SAA) and stochastic approximation (SA). In SAA-based approaches, one replaces the expected value formulation with the average over a large number of samples of the random variable from a pool of already existing data [10], [11]. Successively, the approximated, yet deterministic, problem is solved, and convergence to a solution to the original stochastic problem is proved *when the number of samples grows to infinity*. In contrast, popular SA schemes allow one to *sample a realization of the random variable whenever needed*, e.g., at every algorithmic step, thereby originating fully stochastic procedures for computing an optimal solution asymptotically. While SA-based algorithms are more computationally attractive compared

to those leveraging SAA, since they rely on a smaller number of samples at every iteration, they usually require stronger assumptions on the problem data [4], [9]. As a middle ground alternative, algorithms based on variance-reduced SA have been proposed. In particular, there are two available options: consider the average over an increasing number of samples, still drawn at every iteration [5], [7], [13], or sample a large (but finite) batch only every so often while one single sample is drawn in the majority of iterations [12], [14]. With a few exceptions [4], [8] however, available procedures prove convergence to an exact solution asymptotically, and hence in case the number of samples necessary for the approximation grows indefinitely. This amounts to a rather strong condition to be met in practice, unless one is able to generate i.i.d. samples “for free” through Monte Carlo-based simulations.

To overcome this issue, several efforts have been made in trying to characterize the performance of iterative schemes when only finite information is available. This line of research goes in the direction of understanding the generalization properties of iterative methods, also referred to as *algorithmic stability* analysis [15]. Roughly speaking, a stable learner is one for which the learned solution does not change much, with respect to (w.r.t.) some loss function employed for the “training”, with small changes in the sample set. In light of the extensive use of stochastic gradient descent (SGD) for machine learning techniques, several works analyzed the stability properties of SGD in different domains. For instance, [16], which is considered one of the pioneer works on the stability analysis of SGD for optimization, obtained stability bounds depending on the Lipschitz and strong convexity constants characterizing the cost function, also investigating the convex and non-convex cases, with resulting bounds that may depend on the number of iterations and learning rate. Similar results were obtained by [17], extended in [18] and follow-up papers for adversarial training. More recently, instead, stability bounds of the SGD were discussed also for (VIs) [19] and minmax problems [20], [21]. Finally, a vision not strictly related to SGD can be found in [22], which considered distributed learning algorithms for big data.

### *B. Summary of contribution*

In this paper we take a practical viewpoint. Inspired by a recent trend in system identification [23], to approximate the operator  $\mathcal{B}$  we assume to have available a finite number of i.i.d. samples

drawn from an unknown probability distribution with no restrictions. Then, by focusing on the forward-backward (FB) scheme in Algorithm 1, which is arguably the most used operator splitting technique yielding first-order iterative methods [1], [2], we ask ourselves: how far can we get from some  $\omega^* \in \text{zer}(\mathcal{A} + \mathcal{B})$  by running such a scheme with an approximation of  $\mathcal{B}$  exploiting a finite dataset?

Let  $\omega^{K+1}$  be the output of the resulting data-driven FB obtained after  $K$  iterations. Since exact convergence to a zero should not be expected in this limited information setup, our goal is hence to characterize  $\|\omega^{K+1} - \omega^*\|$  with rigorous data-based certificates. By leveraging tools proper of the algorithmic stability framework [15], we make the following contributions:

- (i) We design a tailored loss function in a way that it is representative for our purposes, i.e., to characterize the distance of  $\omega^{K+1}$  from some  $\omega^*$ ;
- (ii) By considering different monotonicity assumptions on the operators involved, we prove uniform stability of the data-driven FB w.r.t. such a loss function. In line with what observed in convex optimization [16], we note that:
  - The weaker conditions granting convergence of the FB scheme lead to a stability bound proportional to  $K$ ;
  - Stronger assumptions make stability independent on the number of iterations performed.
 As far as we know, none of the abovementioned works has applied algorithmic stability to operator splitting methods based on data, i.e., to a more general framework than mere convex optimization. Our analysis substantially departs from prior approaches, relying on monotone operator-theoretic tools that have not been addressed so far.
- (iii) We derive computable expressions for  $\varepsilon \geq 0$  that depend, among the others, on the amount of data available, so that  $\|\omega^{K+1} - \omega^*\| \leq \varepsilon$  holds with arbitrarily high confidence and regardless of the distribution underlying the data;
- (iv) We apply our bounds to a popular stochastic Nash equilibrium (SNE) seeking algorithm based on the FB scheme, thereby establishing certificates on the distance between the output such algorithm produces and an SNE.

Our theoretical results, which align with and generalize existing literature on algorithmic stability-

based approaches, are finally validated through numerical simulations. Specifically, we analyze a stochastic Nash equilibrium problem (SNEP) modeling a control problem for smart grids, where energy price uncertainty is approximated by means of historical data.

### C. Paper organization and notation

The rest of the paper is organized as follows: in §II we formalize the data-driven problem addressed and provide preliminary concepts on algorithmic stability. In §III we discuss uniform stability properties for the FB scheme w.r.t. a predefined loss function under two different sets of assumptions, as well as derive the resulting distribution-free probabilistic certificates. The latter are then specialized to a popular FB stochastic Nash equilibrium seeking algorithm in §IV, while numerical simulations are finally conducted in §V.

In the remainder we will use Standing Assumption to postulate properties that hold throughout the paper, while we refer to a specific Assumption only when needed.

*Standard notation:*  $\mathbb{N}$ ,  $\mathbb{R}$  and  $\mathbb{R}_{\geq 0}$  denote the set of natural, real, and nonnegative real numbers, respectively.  $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$ , while  $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$ .  $\mathbb{S}^n$  is the space of  $n \times n$  symmetric matrices and  $\mathbb{S}_{(\succ) \succ 0}^n$  is the cone of positive (semi-)definite matrices. The transpose of a matrix  $A \in \mathbb{R}^{n \times n}$  is  $A^\top$ , while  $A \succ 0$  ( $\succcurlyeq 0$ ) denotes its positive (semi)definiteness. For a vector  $u \in \mathbb{R}^n$  and a matrix  $A \succ 0$ ,  $\|u\|$  denotes the standard Euclidean norm, while  $\|\cdot\|_A$  the  $A$ -induced norm  $\|u\|_A := \sqrt{u^\top A u}$ .  $I_n$ ,  $\mathbf{1}_n$ , and  $\mathbf{0}_n$  denote the  $n \times n$  identity matrix, the vector of all 1, and 0, respectively (we omit the dimension  $n$  whenever clear). The operator  $\text{col}(\cdot)$  stacks its arguments in column vectors or matrices of compatible dimensions. For example, given vectors  $x_1, \dots, x_N$  with  $x_i \in \mathbb{R}^{n_i}$  and  $\mathcal{I} = \{1, \dots, N\}$ , we denote  $\mathbf{x} := (x_1^\top, \dots, x_N^\top)^\top = \text{col}((x_i)_{i \in \mathcal{I}}) \in \mathbb{R}^n$ ,  $n := \sum_{i \in \mathcal{I}} n_i$ , and  $\mathbf{x}_{-i} := \text{col}((x_j)_{j \in \mathcal{I} \setminus \{i\}})$ . With a slight abuse of notation, we sometimes use also  $\mathbf{x} = (x_i, \mathbf{x}_{-i})$ . The uniform distribution on  $[a, b]$  is denoted by  $\mathcal{U}(a, b)$ , and the normal distribution with mean  $\mu$  and variance  $\sigma^2$  by  $\mathcal{N}(\mu, \sigma^2)$ .

*Operator-theoretic definitions:* Given a set  $\mathcal{X} \subseteq \mathbb{R}^n$ ,  $\iota_{\mathcal{X}} : \mathbb{R}^n \rightarrow \bar{\mathbb{R}}$  denotes the associated indicator function, i.e.,  $\iota_{\mathcal{X}}(u) = 0$  if  $u \in \mathcal{X}$ ,  $\iota_{\mathcal{X}}(u) = +\infty$  otherwise. If  $\mathcal{X}$  is nonempty and

---

**Algorithm 1:** Forward-backward iterative scheme
 

---

**Initialization:** Set  $\gamma > 0$ ,  $x^0 \in \mathbb{R}^n$

**Iteration**  $k \in \mathbb{N}_0$ :

$$\begin{aligned} y^k &= x^k - \gamma \mathcal{B}(x^k) \\ x^{k+1} &= \mathbf{J}_{\gamma \mathcal{A}}(y^k) \end{aligned} \tag{1}$$


---

convex, the normal cone of  $\mathcal{X}$  evaluated at  $u$  is the set-valued mapping  $\mathbf{N}_{\mathcal{X}} : \mathbb{R}^n \rightarrow 2^{\mathbb{R}^n}$ , defined as  $\mathbf{N}_{\mathcal{X}}(u) := \{d \in \mathbb{R}^n : d^\top(v - u) \leq 0, \text{ for all } v \in \mathcal{X}\}$  if  $u \in \mathcal{X}$ ,  $\mathbf{N}_{\mathcal{X}}(u) := \emptyset$  otherwise. The set of fixed points of an operator  $\mathcal{T} : \mathcal{X} \rightarrow \mathcal{X}$  is denoted by  $\text{fix}(\mathcal{T}) := \{x \in \mathcal{X} : \mathcal{T}(x) = x\}$ . Given some function  $g : \mathbb{R}^n \rightarrow \mathbb{R}$ , the proximity and projection mappings are defined as  $\text{prox}_{\gamma g}(u) := \text{argmin}_{v \in \mathbb{R}^n} \{g(v) + \frac{1}{2\gamma} \|u - v\|^2\}$ ,  $\gamma > 0$ , and  $\text{proj}_{\mathcal{X}}(u) := \text{argmin}_{v \in \mathcal{X}} \frac{1}{2} \|u - v\|^2$ , respectively. A mapping  $F : \text{dom}(F) \subseteq \mathbb{R}^n \rightarrow \mathbb{R}^n$  is  $\ell$ -Lipschitz continuous if, for some  $\ell > 0$ ,  $\|F(x) - F(y)\| \leq \ell \|x - y\|$  for all  $x, y \in \text{dom}(F)$ ;  $\eta$ -strongly monotone if, for some  $\eta > 0$ ,  $\langle F(x) - F(y), x - y \rangle \geq \eta \|x - y\|^2$  for all  $x, y \in \text{dom}(F)$ ;  $\beta$ -cocoercive if, for all  $x, y \in \text{dom}(F)$  and for some  $\beta > 0$ ,  $\langle F(x) - F(y), x - y \rangle \geq \beta \|F(x) - F(y)\|^2$ ; maximally monotone if there exists no monotone operator  $G : C \rightarrow \mathbb{R}^n$  such that the graph of  $G$  contains that of  $F$ .

## II. MATHEMATICAL BACKGROUND

As one of the most used operator splitting techniques for designing iterative methods in Nash equilibrium (NE) seeking [5], [24] and machine learning [12], [14], [16], the FB scheme in Algorithm 1 is defined with  $\omega := \text{col}(y, x) \in \mathbb{R}^{2n}$ , learning rate  $\gamma > 0$ , and operators  $\mathcal{A} : \Omega \rightarrow 2^\Omega$ ,  $\mathcal{B} : \Omega \rightarrow \Omega$ , momentarily assumed monotone,  $\Omega \subseteq \mathbb{R}^n$ . By referring to the scheme in (1), the goal then turns into finding  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$ . We formalize next the data-driven problem of interest, as well as recall key notions and results on algorithmic stability theory.

### A. Problem statement

As introduced in §I, to compute some  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$  we need to approximate the operator  $\mathcal{B}$  through a finite number of data. To this end, we will therefore assume to rely on several queries to some *noisy operator oracle*  $\mathcal{O} : \Omega \times \Xi \rightarrow \Omega$ . The latter formally amounts to a Borel function so that, given some  $x \in \Omega$  and noise input  $\xi \in \Xi \subseteq \mathbb{R}^d$ , which is distributed according to an unknown probability  $\mathbb{P}$ , provides an unbiased estimate of the operator  $\mathcal{B}$  as postulated next [19]:

**Standing Assumption 2.1.** For all  $x \in \Omega$ ,  $\mathbb{E}_{\mathbb{P}}[\mathcal{O}(x, \xi)] = \mathcal{B}(x)$ . □

In our analysis, we will hence leverage some approximation  $\hat{\mathcal{B}}$  that depends on  $s$ -data taken from the set  $\Xi$ . The problem we address here is hence the following: given a finite dataset, how far can we get from a point in  $\text{zer}(\mathcal{A} + \mathcal{B})$  by employing  $\hat{\mathcal{B}}$  in (1) rather than the true operator  $\mathcal{B}$ ? Specifically, we will make use of the following data-based approximation for  $\mathcal{B}$ :

$$\hat{\mathcal{B}}_s(x) = \hat{\mathcal{B}}(x, \mathcal{D}_s) := \frac{1}{s} \sum_{i=1}^s \mathcal{O}(x, \xi^{(i)}), \quad (2)$$

which turns the FB iterative scheme (1) into the set of instructions reported in Algorithm 2. We will then refer to (3) as the data-driven variant of the FB in (1). In the remainder, we will thus implicitly assume that:

- (i) A set  $\mathcal{D}_s := \{\xi^{(i)}\}_{i=1}^s \in \Xi^s$ , consisting of  $s$  i.i.d. samples drawn from an *unknown* probability measure  $\mathbb{P}$  attached to  $\Xi^s$  is available;
- (ii) Akin to [16], [19], [25], the oracle  $\mathcal{O}$  possesses the same properties postulated for the true operator  $\mathcal{B}$ ;
- (iii) The learning rate  $\gamma$  is given to meet the requirements discussed later when introducing suitable assumptions.

**Standing Assumption 2.2.** There exists  $M > 0$  such that, for all  $x \in \mathbb{R}^n$ ,  $\|\mathcal{B}(x)\| \leq M$ . □

While relying on i.i.d. samples may represent the main practical limitation for applying the probabilistic bounds we will develop, especially for mere control applications, we remark that, once the dataset  $\mathcal{D}_s \in \Xi^s$  is given, Algorithm 2 happens to be *deterministic*. Specifically, after

---

**Algorithm 2:** Data-driven FB iterative scheme
 

---

**Initialization:** Samples  $\mathcal{D}_s$ , set  $\gamma > 0$ ,  $x^0 \in \mathbb{R}^n$

**Iteration**  $k \in \mathbb{N}_0$ :

$$y^k = x^k - \gamma \hat{\mathcal{B}}_s(x^k) = x^k - \frac{\gamma}{s} \sum_{i=1}^s \mathcal{O}(x^k, \xi^{(i)}) \quad (3)$$

$$x^{k+1} = \mathbf{J}_{\gamma \mathcal{A}}(y^k)$$


---

running the underlying scheme for  $K \geq 1$  iterations, we obtain a deterministic output since all the calculations performed do not involve any source of randomness. In addition, Algorithm 2 turns out to be naturally *symmetric* w.r.t. any  $\mathcal{D}_s$ , since the output obtained after  $K$  steps does not depend on the order of the elements in  $\mathcal{D}_s$ .

Our goal is then to drive some empirical gap function (defined later) to zero by training on the operator  $\hat{\mathcal{B}}$  for several iterations, and only seek to control the generalization gap that readily gives a measure on how far we can get from a point in  $\text{zer}(\mathcal{A} + \mathcal{B})$ . In particular, we will be interested in determining the radius  $\varepsilon \geq 0$  characterizing the set of  $\varepsilon$ -zeros, defined as:

$$\text{zer}_\varepsilon(\mathcal{A} + \mathcal{B}) := \{\omega \in \Omega : \exists z \in \mathcal{A}(\omega) \text{ s.t. } \|z + \mathcal{B}(\omega)\| \leq \varepsilon\}. \quad (4)$$

Our task will be accomplished by taking an algorithmic stability perspective [15], as introduced next.

### B. Preliminaries on algorithmic stability

An *algorithm* is formally defined as an indexed family of mappings  $\{A_s\}_{s \geq 0}$ , with  $A_s : \Xi^s \rightarrow \mathbb{R}^{2n}$  taking some dataset with  $s$  samples and returning a deterministic hypothesis  $H_s := A_s(\{\xi^{(i)}\}_{i=1}^s) = A_s(\mathcal{D}_s)$  [26]. We will consider later the following sets associated with  $\mathcal{D}_s$ :

- The set obtained by *removing* the  $i$ -th element from  $\mathcal{D}_s$ :

$$\mathcal{D}_{s-1}^{\setminus i} = \{\xi^{(1)}, \dots, \xi^{(i-1)}, \xi^{(i+1)}, \dots, \xi^{(s)}\}.$$

In short, we also denote  $H_{s \setminus i} := A_{s-1}(\mathcal{D}_{s-1}^{\setminus i})$ ;

- The set obtained by *replacing* the  $i$ -th element from  $\mathcal{D}_s$ :

$$\mathcal{D}_s^i = \{\xi^{(1)}, \dots, \xi^{(i-1)}, \xi', \xi^{(i+1)}, \dots, \xi^{(s)}\},$$

where  $\xi' \in \Xi$  is drawn according to  $\mathbb{P}$ , i.i.d. w.r.t.  $\mathcal{D}_s \setminus \{\xi^{(i)}\}$ . Here, we indicate  $H_{s^i} := A_s(\mathcal{D}_s^i)$ .

The performance of an algorithm is generally evaluated through some function  $\ell : \mathbb{R}^{2n} \times \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$  that measures the *loss* associated to an hypothesis  $H$  w.r.t. an example  $\xi$ . Attached to the loss function, one identifies the so-called *risk* or *generalization error*, which coincides with a random variable (the hypothesis generated indeed depends on  $\mathcal{D}_s$ ), defined as:

$$r(A, s) = \mathbb{E}_{\mathbb{P}} [\ell(H_s, \xi)], \quad (5)$$

where we use  $A$  instead of  $A_s$  as first argument. Note that computing (5) would require  $\mathbb{P}$ , which is however unavailable in the considered framework. Nevertheless, the simplest estimator for (5) amounts to the *empirical error*, which reads as:

$$\hat{r}(A, s) = \frac{1}{s} \sum_{i=1}^s \ell(H_s, \xi^{(i)}). \quad (6)$$

**Definition 2.3** ([15, Def. 6] Uniform stability). *An algorithm  $\{A_s\}_{s \geq 0}$  has uniform stability  $\beta = \beta(s)$  w.r.t. the loss function  $\ell$  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  $\xi \in \Xi$ .  $\square$*

A mapping is then uniformly stable if changing one sample  $\xi^{(i)}$  in its input dataset does not significantly alter the output w.r.t. to the metric identified by the loss  $\ell$  itself. Definition 2.3 is also frequently stated as  $|\ell(H_s, \xi) - \ell(H_{s^i}, \xi)| \leq \beta$  for all  $\xi \in \Xi$  [27], [28], i.e., considering sample replacement rather than removal. This is indeed how it will be employed later.

The following result will be key to establish our sample complexity bounds on the generalization error, which we will see that, in some cases, it provides a direct upper bound on the distance from some  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$ . An appropriate choice for  $\ell$  will hence be inevitably crucial for our purposes:

**Lemma 2.4** ([15, Th. 12] Exponential bound with uniform stability). *Let  $\{A_s\}_{s \geq 0}$  be an algorithm with uniform stability  $\beta = \beta(s)$  w.r.t. a loss function  $0 \leq \ell(H_s, \xi) \leq \bar{\ell}$ ,  $\bar{\ell} \geq 0$ , for all  $\xi \in \Xi$  and  $\mathcal{D}_s$ . Then, for any  $s \geq 1$  and  $\delta \in (0, 1)$ , the following bound hold true:*

$$\mathbb{P}^s \left\{ \mathcal{D}_s \in \Xi^s : r(A, s) \leq \hat{r}(A, s) + 2\beta + (4s\beta + \bar{\ell}) \sqrt{\frac{\ln(1/\delta)}{2s}} \right\} \geq 1 - \delta.$$

□

Based on the McDiarmid's inequality [29], Lemma 2.4 says that, with arbitrary confidence  $1 - \delta$ , the risk associated to hypothesis  $H_s = A_s(\mathcal{D}_s)$  is upper bounded by quantities that: i) if  $\beta$  scales inversely with  $s$ , vanishes as  $s \rightarrow \infty$  (consistency property), and ii) it is exponential in the confidence parameter  $\delta$ . Referring to i), we will prove later that it is indeed the case for the data-driven version of the FB scheme in Algorithm 2, once identified  $\omega = \text{col}(y, x)$  as the hypothesis of our algorithm. Since we will consider sample replacement rather than removal, after revisiting the proof of Lemma 2.4 we will employ the following slightly different bound [27, Th. 3.2]:

$$\mathbb{P}^s \left\{ \mathcal{D}_s \in \Xi^s : r(A, s) \leq \hat{r}(A, s) + \beta + (s\beta + \bar{\ell}) \sqrt{\frac{2 \ln(1/\delta)}{s}} \right\} \geq 1 - \delta. \quad (7)$$

We finally note that, since the RHS in both bounds depends on  $\hat{r}(A, s)$ , it is unpractical to design loss functions that depend on possibly unknown quantities such as, e.g., some  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$  itself. This is simply because the term  $\sum_{i=1}^s \ell(H_s, \xi^{(i)})$  in the RHS might not be computed directly.

### III. GENERALIZATION BOUNDS FOR THE DATA-DRIVEN FB ALGORITHM

We now establish stability properties for Algorithm 2 w.r.t. a predefined loss function under two different sets of assumptions commonly used to show convergence of Algorithm 1.

Then, by letting  $H = \omega$  as the generic hypothesis of our data-based FB in (8), we consider the following loss function as a measure of the associated performance, for a given  $\gamma > 0$ :

$$\begin{aligned} \ell(H, \xi) &= \|[0_n \quad I_n]H - \gamma\mathcal{O}([0_n \quad I_n]H, \xi) - [I_n \quad 0_n]H\| \\ &= \|x - \gamma\mathcal{O}(x, \xi) - y\|. \end{aligned} \tag{8}$$

According to the requirements of Lemma 2.4, we need to impose the following mild condition on the underlying loss:

**Standing Assumption 3.1.** *For all  $\xi \in \Xi$  and  $\mathcal{D}_s \in \Xi^s$ ,  $0 \leq \ell(H_s, \xi) \leq \bar{\ell}$ , for some  $\bar{\ell} \geq 0$ .  $\square$*

Although not the only possible choice, it will be clear from Theorems 3.4 and 3.8 and related proofs why such a loss function is particularly convenient for our purposes. As we will note indeed, besides capturing the fixed-point residual, such a specific choice will make our bounds particularly easy to compute in practical cases—see the discussion in §III-C.

#### A. Stability bounds independent on the number of iterations

Before proceeding in the analysis, let us introduce some properties of the involved operators.

**Assumption 3.2.** *The operator  $\mathcal{A} : \Omega \rightarrow 2^\Omega$  is maximally monotone, while  $\mathcal{B} : \Omega \rightarrow \Omega$  is  $\mu$ -strongly monotone and  $\kappa$ -Lipschitz continuous, with  $\mu \leq \kappa$ , and  $\gamma \in (0, 2\mu/\kappa^2)$ .  $\square$*

Under Assumption 3.2, the sequence  $\{x^k\}_{k \in \mathbb{N}}$  produced by Algorithm 1 converges linearly to the unique  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$  [2, Prop. 25.9]. According to the discussion in §II-A, throughout this subsection we will assume that, for all  $\xi \in \Xi$ ,  $x \mapsto \mathcal{O}(x, \xi)$  is  $\mu$ -strongly monotone and  $\kappa$ -Lipschitz continuous. We can then preliminary establish what follows:

**Lemma 3.3.** *Let  $\tau := \sqrt{1 - \gamma(2\mu - \gamma\kappa^2)} < 1$ . Under Assumption 3.2, Algorithm 2 possesses  $(2\gamma M(1 + \tau)/(s(1 - \tau)))$ -uniform stability w.r.t. the loss function  $\ell(H, \xi)$  in (8).  $\square$*

*Proof.* Fix some  $s \geq 1$ , and consider two datasets,  $\mathcal{D}_s, \mathcal{D}_s^i \in \Xi^s$ , both consisting of  $s$ -i.i.d. random samples and differing on the  $i$ -th instance only, i.e., some  $\xi'$  replaces  $\xi^{(i)}$  in  $\mathcal{D}_s^i$ . After

iterating Algorithm 2 for  $K \geq 1$  times by leveraging samples in  $\mathcal{D}_s$  and  $\mathcal{D}_s^i$ , and starting from the same  $x^0 \in \mathbb{R}^n$ , we obtain  $H_s = \omega^K$  and  $H_{s^i} = \tilde{\omega}^K$  as the two corresponding hypotheses. By picking an arbitrary  $\xi \in \Xi$  we readily have:

$$\begin{aligned}
|\ell(H_s, \xi) - \ell(H_{s^i}, \xi)| &\stackrel{(a)}{\leq} \|x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi) - \tilde{x}^{K+1} + \gamma\mathcal{O}(\tilde{x}^{K+1}, \xi) - y^K + \tilde{y}^K\| \\
&\stackrel{(b)}{\leq} \tau \|x^{K+1} - \tilde{x}^{K+1}\| + \|y^K - \tilde{y}^K\| \\
&= \tau \|\mathbf{J}_{\gamma\mathcal{A}}(y^K) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K)\| + \|y^K - \tilde{y}^K\| \\
&\stackrel{(c)}{\leq} (1 + \tau) \|y^K - \tilde{y}^K\|,
\end{aligned} \tag{9}$$

where (a) follows in view of the reverse triangle inequality, (b) since standard calculations exploiting the  $\mu$ -strong monotonicity and  $\kappa$ -Lipschitz continuity of the operator  $\mathcal{O}(\cdot, \xi)$  reveals that:

$$\begin{aligned}
&\| (x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi)) - (\tilde{x}^{K+1} - \gamma\mathcal{O}(\tilde{x}^{K+1}, \xi)) \|^2 \\
&= \|x^{K+1} - \tilde{x}^{K+1}\|^2 - 2\gamma \langle x^{K+1} - \tilde{x}^{K+1}, \mathcal{O}(x^{K+1}, \xi) - \mathcal{O}(\tilde{x}^{K+1}, \xi) \rangle \\
&\quad + \gamma^2 \|\mathcal{O}(x^{K+1}, \xi) - \mathcal{O}(\tilde{x}^{K+1}, \xi)\|^2 \\
&\leq \|x^{K+1} - \tilde{x}^{K+1}\|^2 - 2\mu\gamma \|x^{K+1} - \tilde{x}^{K+1}\|^2 + \gamma^2 \kappa^2 \|x^{K+1} - \tilde{x}^{K+1}\|^2 \\
&= (1 - 2\mu\gamma + \gamma^2 \kappa^2) \|x^{K+1} - \tilde{x}^{K+1}\|^2 =: \tau^2 \|x^{K+1} - \tilde{x}^{K+1}\|^2,
\end{aligned}$$

with  $\tau \in [0, 1)$  since  $\gamma \in (0, 2\mu/\kappa^2)$  and  $\mu \leq \kappa$  in view of Assumption 3.2, and (c) since  $\mathbf{J}_{\gamma\mathcal{A}}(\cdot)$  is the resolvent of a maximally monotone operator  $\mathcal{A}$ , and hence firmly nonexpansive [2, Cor. 23.10], i.e., for all  $\gamma > 0$  it satisfies  $\|\mathbf{J}_{\gamma\mathcal{A}}(y^K) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K)\|^2 \leq \|y^K - \tilde{y}^K\|^2 - \|(y^K - \mathbf{J}_{\gamma\mathcal{A}}(y^K)) - (\tilde{y}^K - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K))\|^2 \leq \|y^K - \tilde{y}^K\|^2$ . According to Definition 2.3, it then remains to bound  $\|y^K - \tilde{y}^K\|$ , possibly as a function of  $1/s$ . To this end, we will exploit the generalized growth lemma established in [25, Lemma 2]. Note that  $\|y^K - \tilde{y}^K\| = \|x^K - \frac{\gamma}{s} \sum_{j=1, j \neq i}^s \mathcal{O}(x^K, \xi^{(j)}) - \frac{\gamma}{s} \mathcal{O}(x^K, \xi^{(i)}) - (\tilde{x}^K - \frac{\gamma}{s} \sum_{j=1, j \neq i}^s \mathcal{O}(\tilde{x}^K, \xi^{(j)}) - \frac{\gamma}{s} \mathcal{O}(\tilde{x}^K, \xi^{(i)}))\|$ . Then, for each of the two trajectories yielding  $y^K$  and  $\tilde{y}^K$ , let us define the following update processes:

$$\begin{aligned}
P(x) &:= x - \frac{\gamma}{s} \sum_{\substack{j=1 \\ j \neq i}}^s \mathcal{O}(x, \xi^{(j)}), & \hat{P}_k(x) &:= -\frac{\gamma}{s} \mathcal{O}(x, \xi^{(i)}) \\
P'(x) &:= P(x), & \hat{P}'_k(x) &:= -\frac{\gamma}{s} \mathcal{O}(x, \xi^{(i)}).
\end{aligned}$$

In this way, at the generic iteration  $k$  we have that  $\|y^k - \tilde{y}^k\| = \|x^k - \frac{\gamma}{s} \sum_{j=1, j \neq i}^s \mathcal{O}(x^k, \xi^{(j)}) - \frac{\gamma}{s} \mathcal{O}(x^k, \xi^{(i)}) - \tilde{x}^k + \frac{\gamma}{s} \sum_{j=1, j \neq i}^s \mathcal{O}(\tilde{x}^k, \xi^{(j)}) + \frac{\gamma}{s} \mathcal{O}(\tilde{x}^k, \xi^{(i)})\| = \|P(x^k) + \hat{P}_k(x^k) - P'(\tilde{x}^k) - \hat{P}'_k(\tilde{x}^k)\|$ .

Thus, by relying on [25, Lemma 2], we obtain that  $\|y^k - \tilde{y}^k\| \leq \theta \|x^k - \tilde{x}^k\| + \|\hat{P}_k(x^k)\| + \|\hat{P}'_k(\tilde{x}^k)\|$ , where  $\theta$  is the parameter of contraction of the update rule  $P(x) + \hat{P}_k(x)$ , which in view of the calculations above is upper bounded by  $\tau$ , smaller than 1 if  $\gamma \in (0, 2\mu/\kappa^2)$  and  $\mu \leq \kappa$ . In view of the boundedness of the operator  $\mathcal{O}(\cdot, \xi)$ , obtained from Standing Assumption 2.2, both  $\|\hat{P}_k(x^k)\|$  and  $\|\hat{P}'_k(\tilde{x}^k)\|$  are then upper bounded by  $\gamma M/s$ , thereby yielding:

$$\begin{aligned} \|y^k - \tilde{y}^k\| &\leq \tau \|x^k - \tilde{x}^k\| + \frac{2\gamma M}{s} = \tau \|\mathbf{J}_{\gamma\mathcal{A}}(y^{k-1}) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^{k-1})\| + \frac{2\gamma M}{s} \\ &\leq \tau \|y^{k-1} - \tilde{y}^{k-1}\| + \frac{2\gamma M}{s}. \end{aligned}$$

Recurring over the  $K$ -steps by considering that i) the resolvent operator  $\mathbf{J}_{\gamma\mathcal{A}}(\cdot)$  is firmly nonexpansive, and hence  $\|\mathbf{J}_{\gamma\mathcal{A}}(y^{k-1}) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^{k-1})\| \leq \|y^{k-1} - \tilde{y}^{k-1}\|$  at every  $k$ , ii) both trajectories originate from the same initial condition  $x^0 \in \mathbb{R}^n$ , and iii) exploiting the relation  $\sum_{j=1}^K a^{K-j} = (1 - a^K)/(1 - a)$ , which holds true for any  $a \in \mathbb{R}$ , we obtain:

$$\|y^K - \tilde{y}^K\| \leq \frac{2\gamma M}{s} \sum_{j=1}^K \tau^{K-j} = \frac{2\gamma M(1 - \tau^K)}{s(1 - \tau)} \leq \frac{2\gamma M}{s(1 - \tau)},$$

Putting everything together with (9), we finally arrive at the following relation:

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

which proves uniform stability with bound inverse to  $s$ . ■

We can then prove one of the main results of this section:

**Theorem 3.4.** *Fix  $s \geq 1$  and  $\delta \in (0, 1)$ . Given any dataset  $\mathcal{D}_s \in \Xi^s$ , under Assumption 3.2 there exists  $\varepsilon = \varepsilon(s, \delta) > 0$  such that  $x^{K+1} \in \text{zer}_\varepsilon(\mathcal{A} + \mathcal{B})$  with probability at least  $1 - \delta$ , where  $x^{K+1}$  is obtained by iterating Algorithm 2  $K$ -times. □*

*Proof.* From [2, Prop. 25.1(iv)], we have that  $\text{zer}(\mathcal{A} + \mathcal{B}) = \text{fix}(\mathbf{J}_{\gamma\mathcal{A}} \circ (\text{Id} - \gamma\mathcal{B}))$ . Therefore, a possible metric to evaluate the distance between a generic  $x^{k+1}$  and  $x^*$  is  $\|\mathbf{J}_{\gamma\mathcal{A}}(x^{k+1} - \gamma\mathcal{B}(x^{k+1})) -$

$x^{k+1}$ ], which in turn measures the gap between consecutive iterations on variable  $x$ . After running Algorithm 2  $K$ -times, consider the following relations:

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

where we have used in (a) the fact that the oracle operator is unbiased in view of Standing Assumption 2.1, together with the second relation in (3), in (b) the firm nonexpansiveness of the resolvent of  $\mathcal{A}$ , for any  $\gamma > 0$ , in (c) the fact that the data-driven FB in (3) yields a deterministic output  $\omega^K$  once fixed the dataset  $\mathcal{D}_s$ , as well as the linearity of the expected value, and in (d) the Jensen's inequality [30]. Thus, upper bounding the risk  $\mathbb{E}_{\mathbb{P}}[\ell(H_s, \xi)]$  with loss function as in (8) provides an equivalent information on the distance of  $x^{K+1}$  from a fixed point of the dynamics produced by iterating the operator  $J_{\gamma\mathcal{A}} \circ (\text{Id} - \gamma\mathcal{B})$  with perfect knowledge on  $\mathcal{B}$ . Therefore, with probability at least  $1 - \delta$ , applying the bound in (7) leads to  $\|J_{\gamma\mathcal{A}}(x^{K+1} - \gamma\mathcal{B}(x^{K+1})) - x^{K+1}\| \leq \epsilon$  with

$$\epsilon := \frac{1}{s} \sum_{i=1}^s \|x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi^{(i)}) - y^K\| + \frac{2\gamma M(1+\tau)}{s(1-\tau)} + \left( \frac{2\gamma M(1+\tau)}{1-\tau} + \bar{\ell} \right) \sqrt{\frac{2 \ln(1/\delta)}{s}}.$$

This allows us to conclude that  $x^{K+1}$  is an  $\epsilon$ -zero of the sum of  $\mathcal{A}$  and  $\mathcal{B}$ , according to (4), with  $\varepsilon := \epsilon/\gamma$ . In fact, since  $\|J_{\gamma\mathcal{A}}(x^{K+1} - \gamma\mathcal{B}(x^{K+1})) - x^{K+1}\| \leq \epsilon$ , we directly have:

$$\begin{aligned}
\epsilon &\geq \|(\text{Id} + \gamma\mathcal{A})^{-1} \circ (\text{Id} - \gamma\mathcal{B})(x^{K+1}) - x^{K+1}\| \\
&= \|(\text{Id} - \gamma\mathcal{B})(x^{K+1}) - (\text{Id} + \gamma\mathcal{A})(x^{K+1})\| \\
&= \|- \gamma\mathcal{B}(x^{K+1}) - \gamma z\| = \gamma \|z + \mathcal{B}(x^{K+1})\|,
\end{aligned}$$

in view of the single-valuedness of the resolvent  $J_{\gamma\mathcal{A}}(\cdot)$  and its properties, where  $z \in \mathcal{A}(x^{K+1})$ , concluding the proof.  $\blacksquare$

### B. Weaker assumptions yield $K$ -dependent stability bounds

Compared to §III-A, we now weaken the assumptions on the operator  $\mathcal{B}$  and, as a consequence, on the oracle  $\mathcal{O}$ .

**Assumption 3.5.** *The operator  $\mathcal{A} : \Omega \rightarrow 2^\Omega$  is maximally monotone, while  $\mathcal{B} : \Omega \rightarrow \Omega$  is  $\theta$ -cocoercive,  $\theta > 0$ , and  $\gamma \in (0, 2\theta)$ .  $\square$*

According to [2, Th. 25.8], under Assumption 3.5 the sequence  $\{x^k\}_{k \in \mathbb{N}}$  produced by Algorithm 1 converges to some  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$  (following the notation of [2, Th. 25.8],  $\lambda^k = 1$  is indeed a valid choice since  $\delta$  there is strictly greater than 1). In view of the discussion in §II-A, the noisy oracle inherits the same conditions made on  $\mathcal{B}$ . In this case,  $\mathcal{O}(\cdot, \xi)$  is hence  $\theta$ -cocoercive for a given  $\xi \in \Xi$ . We have the following result:

**Lemma 3.6.** *For fixed  $s > 0$ , let  $\{\xi^{(1)}, \dots, \xi^{(s)}\}$  be a given set of samples with  $\xi^{(i)} \in \Xi$ , for all  $i = 1, \dots, s$ . Then, the operator  $\sum_{i=1}^s \mathcal{O}(\cdot, \xi^{(i)})$  is  $(\theta/s)$ -cocoercive.  $\square$*

*Proof.* Given any  $x, y \in \mathbb{R}^n$ , the following relations hold true:

$$\begin{aligned} \left\| \sum_{i=1}^s \mathcal{O}(x, \xi^{(i)}) - \sum_{i=1}^s \mathcal{O}(y, \xi^{(i)}) \right\|^2 &\leq s \sum_{i=1}^s \left\| \mathcal{O}(x, \xi^{(i)}) - \mathcal{O}(y, \xi^{(i)}) \right\|^2 \\ &\leq \frac{s}{\theta} \sum_{i=1}^s \langle x - y, \mathcal{O}(x, \xi^{(i)}) - \mathcal{O}(y, \xi^{(i)}) \rangle \\ &= \frac{s}{\theta} \langle x - y, \sum_{i=1}^s (\mathcal{O}(x, \xi^{(i)}) - \mathcal{O}(y, \xi^{(i)})) \rangle \\ &= \frac{s}{\theta} \langle x - y, \sum_{i=1}^s \mathcal{O}(x, \xi^{(i)}) - \sum_{i=1}^s \mathcal{O}(y, \xi^{(i)}) \rangle, \end{aligned}$$

where in the second inequality we have used the  $\theta$ -cocoercivity of  $\mathcal{O}(\cdot, \xi)$  for fixed  $\xi \in \Xi$ .  $\blacksquare$

We then have uniform stability of Algorithm 2 depending on the number of iterations:

**Lemma 3.7.** *Under Assumption 3.5, Algorithm 2 possesses  $(4\gamma MK/s)$ -uniform stability w.r.t. the loss function  $\ell(H, \xi)$  in (8).  $\square$*

*Proof.* Fix some  $s \geq 1$ , and consider two datasets,  $\mathcal{D}_s, \mathcal{D}_s^i \in \Xi^s$ , both consisting of  $s$ -i.i.d. random samples and differing on the  $i$ -th instance only, i.e., some  $\xi'$  replaces  $\xi^{(i)}$  in  $\mathcal{D}_s^i$ . After iterating Algorithm 2 for  $K \geq 1$  times by leveraging samples in  $\mathcal{D}_s$  and  $\mathcal{D}_s^i$ , and starting from the same  $x^0 \in \mathbb{R}^n$ , we obtain  $H_s = \omega^K$  and  $H_{s^i} = \tilde{\omega}^K$  as the two corresponding hypotheses. By picking an arbitrary  $\xi \in \Xi$  we immediately obtain:

$$\begin{aligned} |\ell(H_s, \xi) - \ell(H_{s^i}, \xi)| &\stackrel{(a)}{\leq} \|x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi) - \tilde{x}^{K+1} + \gamma\mathcal{O}(\tilde{x}^{K+1}, \xi) - y^K + \tilde{y}^K\| \\ &\stackrel{(b)}{\leq} \|x^{K+1} - \tilde{x}^{K+1}\| + \|y^K - \tilde{y}^K\| \\ &= \|\mathbf{J}_{\gamma\mathcal{A}}(y^K) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K)\| + \|y^K - \tilde{y}^K\| \\ &\stackrel{(c)}{\leq} 2\|y^K - \tilde{y}^K\|, \end{aligned}$$

where (a) follows in view of the reverse triangle inequality, (b) since in view of Assumption 3.5 the operator  $\mathcal{O}(\cdot, \xi)$  is  $\theta$ -cocoercive and  $\gamma \in (0, 2\theta)$ , and hence we know from [2, Prop. 4.33] that  $\text{Id}(\cdot) - \gamma\mathcal{O}(\cdot, \xi)$  is  $(\gamma/2\theta)$ -averaged, yielding  $\|x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi) - (\tilde{x}^{K+1} - \gamma\mathcal{O}(\tilde{x}^{K+1}, \xi))\|^2 \leq \|x^{K+1} - \tilde{x}^{K+1}\|^2 - \frac{2\theta-\gamma}{\gamma} \|\gamma\mathcal{O}(x^{K+1}, \xi) - \gamma\mathcal{O}(\tilde{x}^{K+1}, \xi)\|^2 \leq \|x^{K+1} - \tilde{x}^{K+1}\|^2$ , and (c) since  $\mathbf{J}_{\gamma\mathcal{A}}(\cdot)$  is the resolvent of the maximally monotone operator  $\mathcal{A}$ , and hence firmly nonexpansive [2, Cor. 23.10], i.e., for all  $\gamma > 0$  it satisfies  $\|\mathbf{J}_{\gamma\mathcal{A}}(y^K) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K)\|^2 \leq \|y^K - \tilde{y}^K\|^2 - \|(y^K - \mathbf{J}_{\gamma\mathcal{A}}(y^K)) - (\tilde{y}^K - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^K))\|^2 \leq \|y^K - \tilde{y}^K\|^2$ . According to Definition 2.3, also in this case it remains to bound  $\|y^K - \tilde{y}^K\|$ , possibly as a function of  $1/s$ .

To this end, we note that  $\|y^K - \tilde{y}^K\| = \|x^K - \frac{\gamma}{s} \sum_{j=1}^s \mathcal{O}(x^K, \xi^{(j)}) - (\tilde{x}^K - \frac{\gamma}{s} \sum_{j=1, j \neq i}^s \mathcal{O}(\tilde{x}^K, \xi^{(j)}) - \frac{\gamma}{s} \mathcal{O}(\tilde{x}^K, \xi'))\|$ . More generally, at the  $k$ -th iteration we have:

$$\begin{aligned} \|y^k - \tilde{y}^k\| &= \left\| x^k - \frac{\gamma}{s} \sum_{j=1}^s \mathcal{O}(x^k, \xi^{(j)}) - \tilde{x}^k + \frac{\gamma}{s} \sum_{j=1}^s \mathcal{O}(\tilde{x}^k, \xi^{(j)}) \right\| \\ &\leq \left\| x^k - \frac{\gamma}{s} \sum_{\substack{j=1 \\ j \neq i}}^s \mathcal{O}(x^k, \xi^{(j)}) - \left( \tilde{x}^k - \frac{\gamma}{s} \sum_{\substack{j=1 \\ j \neq i}}^s \mathcal{O}(\tilde{x}^k, \xi^{(j)}) \right) \right\| \\ &\quad + \left\| \frac{\gamma}{s} \mathcal{O}(x^k, \xi^{(i)}) \right\| + \left\| \frac{\gamma}{s} \mathcal{O}(\tilde{x}^k, \xi') \right\| \\ &\leq \|x^k - \tilde{x}^k\| + \frac{2\gamma M}{s} = \|\mathbf{J}_{\gamma\mathcal{A}}(y^{k-1}) - \mathbf{J}_{\gamma\mathcal{A}}(\tilde{y}^{k-1})\| + \frac{2\gamma M}{s} \\ &\leq \|y^{k-1} - \tilde{y}^{k-1}\| + \frac{2\gamma M}{s}, \end{aligned}$$

where the second to last inequality follows by combining Lemma 3.6 with the boundedness of the operator  $\mathcal{O}(\cdot, \xi)$  in Standing Assumption 2.2. From [2, Prop. 4.33], the main consequence of Lemma 3.6 is indeed the  $(\gamma/2\theta)$ -averagedness of the operator  $\text{Id}(\cdot) - \frac{\gamma}{s} \sum_{\substack{j=1 \\ j \neq i}}^s \mathcal{O}(\cdot, \xi^{(j)})$  when  $\gamma < 2\theta$ , which is imposed through Assumption 3.5. The last inequality then follows again due to the firm nonexpansiveness of  $J_{\gamma\mathcal{A}}(\cdot)$ .

Thus, recursing over the  $K$ -steps by considering that i) the resolvent operator  $J_{\gamma\mathcal{A}}(\cdot)$  is (firmly) nonexpansive, and ii) both trajectories leading to  $y^K$  and  $\tilde{y}^K$  originate from the same initial condition  $x^0 \in \mathbb{R}^n$ , yields  $\|y^K - \tilde{y}^K\| \leq 2\gamma MK/s$ . Putting everything together, we finally obtain what follows:

$$|\ell(H_s, \xi) - \ell(H_{s^i}, \xi)| \leq \frac{4\gamma MK}{s},$$

which proves uniform stability with bound inverse to  $s$  and linear in  $K$ , concluding the proof.  $\blacksquare$

We can finally prove the following sample complexity bound:

**Theorem 3.8.** *Fix  $s \geq 1$  and  $\delta \in (0, 1)$ . Given any dataset  $\mathcal{D}_s \in \Xi^s$ , under Assumption 3.5 there exists  $\varepsilon = \varepsilon(s, \delta, K) > 0$  such that  $x^{K+1} \in \text{zer}_\varepsilon(\mathcal{A} + \mathcal{B})$  with probability at least  $1 - \delta$ , where  $x^{K+1}$  is obtained by iterating Algorithm 2  $K$ -times.  $\square$*

*Proof.* As in the proof of Theorem 3.4, upper bounding the risk  $\mathbb{E}_{\mathbb{P}}[\ell(H_s, \xi)]$  provides an equivalent information on the distance of  $x^{K+1}$  from a fixed point of the dynamics produced by iterating the operator  $J_{\gamma\mathcal{A}} \circ (\text{Id} - \gamma\mathcal{B})$ . Therefore, with probability at least  $1 - \delta$ , applying the bound in (7) leads to  $\|J_{\gamma\mathcal{A}}(x^{K+1} - \gamma\mathcal{B}(x^{K+1})) - x^{K+1}\| \leq \varepsilon$  with

$$\varepsilon := \frac{1}{s} \sum_{i=1}^s \|x^{K+1} - \gamma\mathcal{O}(x^{K+1}, \xi^{(i)}) - y^K\| + \frac{4\gamma MK}{s} + (4\gamma MK + \bar{\ell}) \sqrt{\frac{2 \ln(1/\delta)}{s}}.$$

Akin to the conclusion of the proof of Theorem 3.4,  $x^{K+1}$  is then an  $\varepsilon$ -zero in the spirit of (4), with  $\varepsilon := \varepsilon/\gamma$ .  $\blacksquare$

**Remark 3.9.** *Our results are not only consistent with, but also generalize existing ones on SGD in convex optimization, since our operator-theoretic framework embraces a wider class of problems. When the cost function at hand is simply convex, SGD is characterized by a stability bound proportional to the number of iterations performed [16, Th. 3.8]. Consistently, in §III-B*

we show that Algorithm 2 exhibits a stability bound proportional to  $K$  when the operator  $\mathcal{B}$  is merely cocoercive. Conversely, if the cost is strongly convex, there is no dependency on the number of iterations [16, Th. 3.9], and we obtain an analogous result in §III-A with strongly monotone  $\mathcal{B}$ .  $\square$

### C. Discussion on our results

The generalization properties established for Algorithm 2 hence provide computable bounds on the distance between its output, after a finite number of steps, and a zero of the operator  $\mathcal{A} + \mathcal{B}$ . As discussed next, this is particularly relevant for real-world applications where data collection through repeated experiments is either prohibitively expensive or infeasible, thus calling for the efficient use of the available, limited samples:

- **Reliability of the solution:** Offering high confidence-type of bounds makes the obtained  $x^{K+1}$  reliable w.r.t. any possible dataset drawn from  $\Xi^s$ . No matter what  $\mathcal{D}_s$  is employed to run Algorithm 2, the solution produced will then always be an  $\varepsilon$ -zero of  $\mathcal{A} + \mathcal{B}$ , with  $\varepsilon$  vanishing as  $O(s^{-1/2})$ , up to a set of arbitrarily small measure  $\delta$ ;
- **Easy-to-compute bounds:** Although *a-posteriori*, i.e., once  $x^{K+1}$  is revealed, the specific choice of the loss function in (8), and hence the empirical risk  $\hat{r}(A, s)$  associated, makes our bounds particularly easy to compute. Notice that for our purposes, i.e., to characterize  $x^{K+1}$  as an  $\varepsilon$ -zero of  $\mathcal{A} + \mathcal{B}$ , we could even have chosen  $\ell(H, \xi) = \|x^* - [0_n \ I_n]H\| = \|x^* - \mathbf{J}_{\gamma\mathcal{A}}(x - \gamma\mathcal{O}(x, \xi))\|$  for some  $x^* \in \text{zer}(\mathcal{A} + \mathcal{B})$ , thereby obtaining similar stability bounds as in Lemmas 3.3 and 3.7. However, computing  $\varepsilon$  in this case is prohibitive unless one knows in advance such an  $x^*$ , making the zero-finding problem meaningless;
- **Distribution-free:** Along the line of results based on non-parametric statistics and learning theory, such as probably approximately correct (PAC) learning [26] or scenario approach theory [31], we further emphasize that we only require the data to be i.i.d.. Our bounds are thus derived without imposing any specific condition on the unknown probability distribution  $\mathbb{P}$  underlying available samples. This provides robust guarantees regardless of how skewed or heavy-tailed the true distribution might be.

#### IV. FINITE-SAMPLE GUARANTEES FOR EQUILIBRIUM SEEKING IN STOCHASTIC NASH GAMES

As a suitable motivating application, we now illustrate how a standard NE seeking algorithm in a stochastic regime fits the framework considered in this paper, and hence benefit the sample complexity bounds developed in §III.

##### A. A stochastic Nash equilibrium problem

We consider a set of  $N$  self-interested agents, indexed by  $\mathcal{I} := \{1, \dots, N\}$ , where each of them controls a strategy  $x_i \in \Omega_i \subseteq \mathbb{R}^{n_i}$ . The goal is then to simultaneously solve a collection of mutually coupled stochastic optimization problems:

$$\forall i \in \mathcal{I} : \min_{x_i \in \mathbb{R}^{n_i}} \mathbb{E}_{\mathbb{P}} [J_i(x_i, \mathbf{x}_{-i}, \xi)] + g_i(x_i), \quad (10)$$

where each  $J_i : \mathbb{R}^n \times \mathbb{R}^d \rightarrow \mathbb{R}$  is some measurable function so that, together with  $g_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}$ , constitutes the cost of the  $i$ -th agent, where  $n := \sum_{i \in \mathcal{I}} n_i$  and  $\mathbf{x}_{-i} := \text{col}((x_j)_{j \in \mathcal{I} \setminus \{i\}}) \in \mathbb{R}^{n-n_i}$ . The vector  $\xi \in \mathbb{R}^d$  denotes, instead, some uncertainty affecting each  $J_i$ , and corresponds to a random variable taking values in some  $\Xi \subseteq \mathbb{R}^d$  according to a possibly unknown probability distribution  $\mathbb{P}$ . Each function  $g_i(\cdot)$  typically represents a nonsmooth term modeling local constraints via an additive indicator function  $\iota_{\Omega_i}(\cdot)$ . Thus, 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  $\xi$ .

Let  $\mathbb{J}_i(x_i, \mathbf{x}_{-i}) := \mathbb{E}_{\mathbb{P}} [J_i(x_i, \mathbf{x}_{-i}, \xi)] + g_i(x_i)$  and  $\Omega := \prod_{i \in \mathcal{I}} \Omega_i$ . Aiming at solving the collection of optimization problems in (10) simultaneously, thereby resulting into a SNEP [32], the solution concept of our interest is formalized next:

**Definition 4.1** ([33, Definition 2.1] Stochastic Nash equilibrium). *A collective decision vector  $\mathbf{x}^* \in \mathbb{R}^n$  is a SNE if, for all  $i \in \mathcal{I}$ ,*

$$\mathbb{J}_i(x_i^*, \mathbf{x}_{-i}^*) \leq \inf_{y_i \in \mathbb{R}^{n_i}} \mathbb{J}_i(y_i, \mathbf{x}_{-i}^*). \quad \square$$

In words, an SNE is a set of strategies where no agent can decrease its cost function by unilaterally deviating from its decision. Moreover, we will call  $\varepsilon$ -SNE any collective strategy

profile  $\mathbf{x}^* \in \Omega$  that satisfies, for all  $i \in \mathcal{I}$ ,  $\mathbb{J}_i(\mathbf{x}_i^*, \mathbf{x}_{-i}^*) \leq \inf_{y_i \in \mathbb{R}^{n_i}} \mathbb{J}_i(y_i, \mathbf{x}_{-i}^*) + \varepsilon$ , for some  $\varepsilon > 0$ .

We now make some standard assumptions on the data of the SNEP at hand [5]:

**Standing Assumption 4.2** (Local cost function). *The following conditions hold true for all  $i \in \mathcal{I}$ :*

- (i) *The function  $x_i \mapsto g_i(x_i)$  is lower semicontinuous and convex with nonempty, compact, and convex domain,  $\text{dom}(g_i) \subseteq \mathbb{R}^{n_i}$ ;*
- (ii) *For all  $\mathbf{x}_{-i} \in \mathbb{R}^{n-n_i}$  and  $\xi \in \Xi$ , the function  $x_i \mapsto J_i(x_i, \mathbf{x}_{-i}, \xi)$  is convex, Lipschitz continuous, and continuously differentiable. For all  $x_i \in \mathbb{R}^{n_i}$  and  $\mathbf{x}_{-i} \in \mathbb{R}^{n-n_i}$ , the function  $\xi \mapsto J_i(x_i, \mathbf{x}_{-i}, \xi)$  is measurable with integrable Lipschitz constant  $\kappa_i(\mathbf{x}_{-i}, \xi) > 0$  w.r.t.  $\xi$ ;*
- (iii) *For all  $\mathbf{x}_{-i} \in \mathbb{R}^{n-n_i}$ ,  $x_i \mapsto \mathbb{E}_{\mathbb{P}} [J_i(x_i, \mathbf{x}_{-i}, \xi)]$  is convex and continuously differentiable.  $\square$*

The conditions postulated in Standing Assumption 4.2 guarantee the existence of an SNE, while they are not sufficient to establish uniqueness [33]. Thus, the derivation of some iterative SNE seeking algorithms able to compute an equilibrium is typically achieved by recasting the SNEP in (10) as a zero-finding problem, i.e., by looking for some  $\mathbf{x}^*$  satisfying:

$$\text{For all } i \in \mathcal{I} : \mathbf{0}_{n_i} \in \mathbb{E}_{\mathbb{P}} [\nabla_{x_i} J_i(x_i^*, \mathbf{x}_{-i}^*, \xi)] + \partial g_i(x_i^*), \quad (11)$$

where, in view of Standing Assumption 4.2.(ii), we are entitled to exchange the expected value and the gradient. We recall that the inclusions above can be derived by imposing the Karush–Kuhn–Tucker (KKT) conditions to each problem (10). By introducing  $\mathbb{F}(\mathbf{x}) := \mathbb{E}_{\mathbb{P}} [F(\mathbf{x}, \xi)]$  with  $F(\mathbf{x}, \xi) := \text{col}((F_i(x_i, \mathbf{x}_{-i}, \xi))_{i \in \mathcal{I}})$  and  $F_i(x_i, \mathbf{x}_{-i}, \xi) := \nabla_{x_i} J_i(x_i, \mathbf{x}_{-i}, \xi)$ , and  $\partial g(\mathbf{x}) := \text{col}((\partial g_i(x_i))_{i \in \mathcal{I}})$ , conditions in (11) can be represented in compact form by operators:

$$\mathcal{A} : \mathbf{x} \mapsto \partial g(\mathbf{x}), \quad \mathcal{B} : \mathbf{x} \mapsto \mathbb{F}(\mathbf{x}).$$

Note that the conditions in (11) require the evaluation of the expected value  $\mathbb{E}_{\mathbb{P}} [\cdot]$ , a quantity that is impossible to compute in case the distribution  $\mathbb{P}$  of the random variable  $\xi$  is not available, as in our framework. Therefore, several data-driven approximation schemes for the evaluation of the pseudogradient mapping, i.e., the operator  $\mathcal{B}$  above, have been proposed in the literature

---

**Algorithm 3:** Data-driven proximal gradient method
 

---

**Initialization:** Samples  $\mathcal{D}_s$ ,  $x_i^0 \in \mathbb{R}^{n_i}$  for all  $i \in \mathcal{I}$

**Iteration**  $k \in \mathbb{N}_0$ : Agent  $i$  receives  $\mathbf{x}_{-i}^k$ , then updates

$$x_i^{k+1} = \text{prox}_{\gamma g_i} \left( x_i^k - \frac{\gamma}{s} \sum_{j=1}^s F_i(x_i^k, \mathbf{x}_{-i}^k, \xi^{(j)}) \right)$$


---

[4]–[8]. However, we take a practical perspective by assuming that the agents have collectively available a finite set of i.i.d. realizations for  $\xi$  coming from, e.g., historical data, gathered into  $\mathcal{D}_s$ . This is particularly relevant for instance in autonomous driving [34], traffic coordination [35], as well as smart grids [36] and demand response management [37]. A key example from the latter domain will be analyzed in §V-A.

We now report conditions on  $\mathbb{F}(\cdot)$  ensuring, among the other, the uniqueness of the SNE associated to the SNEP in (10):

**Standing Assumption 4.3.** *For all  $\xi \in \Xi$ , the mapping  $\mathbf{x} \mapsto F(\mathbf{x}, \xi)$  is  $\mu_F$ -strongly monotone and  $\kappa_F$ -Lipschitz continuous, with  $\mu_F \leq \kappa_F$ , and  $\gamma \in (0, 2\mu_F/\kappa_F^2)$ . Given any  $\xi \in \Xi$ , there exists  $M_F > 0$  so that  $\|F(\mathbf{x}, \xi)\| \leq M_F$  for all  $\mathbf{x} \in \mathbb{R}^n$ .  $\square$*

Our analysis focuses on the distributed, proximal-like procedure reported in Algorithm 3, which can be derived by starting from the agent-wise fixed-point FB iteration, i.e., for all  $i \in \mathcal{I}$ :

$$x_i^* = (I_{n_i} - \gamma \partial g_i)^{-1} (x_i^* + \gamma \mathbb{E}_{\mathbb{P}} [\nabla_{x_i} J_i(x_i^*, \mathbf{x}_{-i}^*, \xi)]).$$

From Standing Assumption 4.2 and [2, Prop. 16.44], we have  $(I_{n_i} - \gamma \partial g_i)^{-1} = \text{prox}_{\gamma g_i}$ , leading to the exact version of Algorithm 3. Similar to (2), the expected value can then be replaced through a data-driven approximation of the local operator  $\mathbb{E}_{\mathbb{P}} [F_i(x_i, \mathbf{x}_{-i}, \xi)]$  by averaging  $F_i(\cdot, \cdot, \xi^{(j)})$  at every iteration over the available  $s$ -samples. Also in this case we tacitly assume each  $F_i(\cdot, \cdot, \xi)$  being an unbiased operator, i.e., for all  $\mathbf{x} \in \mathbb{R}^n$ ,  $\mathbb{E}_{\mathbb{P}} [F_i(\mathbf{x}, \xi)] = \mathbb{F}_i(\mathbf{x})$ , for all  $i \in \mathcal{I}$ . Moreover, to simplify the analysis, we have adopted the same, constant learning rate  $\gamma > 0$  for all the agents, albeit a generalization to local, possibly time-varying  $\{\gamma_{k_i}\}_{i \in \mathcal{I}}$  is doable. A further generalization consists in considering local dataset  $\mathcal{D}_{s^i}$ .

Under the conditions postulated in Standing Assumptions 4.2 and 4.3, which are widely employed in the algorithmic game theory literature for stochastic setting [7], [33], we note that Algorithm 3 exhibits almost sure convergence to the SNE in case of an increasing batch size  $s = s_k \rightarrow \infty$ , along with few other assumptions [5, Th. 1]. In a data-limited context, however, our goal is instead to investigate how far we can get from the unique SNE by running Algorithm 3 for  $K$  iterations. Specifically, we want to characterize, in a probabilistic sense, the resulting  $\mathbf{x}^{K+1}$  as an  $\varepsilon$ -SNE of the SNEP in (10).

### B. Data-driven certificates for $\varepsilon$ -SNE

First, we need to recast Algorithm 3 into the framework of algorithmic stability by starting from its compact reformulation:

$$\mathbf{x}^{k+1} = \mathbf{prox}_{\gamma g} \left( \mathbf{x}^k - \frac{\gamma}{s} \sum_{j=1}^s F(\mathbf{x}^k, \xi^{(j)}) \right), \quad (12)$$

where, with some abuse of notation,  $\mathbf{prox}_{\gamma g}(\cdot)$  is meant to be applied agent-wise with  $g_i$  in place of  $g$ . For the sake of our analysis, (12) can be equivalently rewritten by means of an additional variable  $\mathbf{y} \in \mathbb{R}^n$  as follows:

$$\begin{cases} \mathbf{y}^k = \mathbf{x}^k - \frac{\gamma}{s} \sum_{j=1}^s F(\mathbf{x}^k, \xi^{(j)}) \\ \mathbf{x}^{k+1} = \mathbf{prox}_{\gamma g}(\mathbf{y}^k). \end{cases} \quad (13)$$

In fact, due to the properties of the proximity operator we have:

$$\begin{cases} \mathbf{y}^* = \mathbf{x}^* - \gamma \mathbb{F}(\mathbf{x}^*) \\ \mathbf{x}^* = \mathbf{prox}_{\gamma g}(\mathbf{y}^*) = \mathbf{prox}_{\gamma g}(\mathbf{x}^* - \gamma \mathbb{F}(\mathbf{x}^*)), \end{cases}$$

i.e., a fixed point of the dynamics in (13), in case  $\mathbb{F}(\cdot)$  was available and in place of the average  $\frac{1}{s} \sum_{j=1}^s F(\mathbf{x}^k, \xi^{(j)})$ , equivalently produces an SNE  $\mathbf{x}^*$  as subvector, thereby linking the augmented dynamics (13) with the conditions in (11). By making a parallelism with §II-B, starting from  $\mathbf{x}^0 \in \mathbb{R}^n$  we will let coincide with our hypothesis  $H_s$  the output provided by iterating (13)  $K$ -times, i.e.,  $\boldsymbol{\omega}^K := \text{col}(\mathbf{y}^K, \mathbf{x}^{K+1}) \in \mathbb{R}^{2n}$ .

The considerations above suggest us to choose the following function as a candidate loss associated to a hypothesis  $H$ :

$$\begin{aligned}\ell(H, \xi) &= \|[0_n \quad I_n]H - \gamma F([0_n \quad I_n]H, \xi) - [I_n \quad 0_n]H\| \\ &= \|\mathbf{x} - \gamma F(\mathbf{x}, \xi) - \mathbf{y}\|.\end{aligned}\tag{14}$$

We can then establish what follows:

**Lemma 4.4.** *With  $\tau_F := \sqrt{1 - \gamma(2\mu_F - \gamma\kappa_F^2)} < 1$ , Algorithm 3 possesses  $(2\gamma M_F(1 + \tau_F))/(s(1 - \tau_F))$ -uniform stability w.r.t. the loss function  $\ell(H, \xi)$  in (14).  $\square$*

*Proof.* The proof follows exactly the same steps as per the proof of Lemma 3.3, once noted that Standing Assumptions 4.2 and 4.3 together allow (13) to satisfy the conditions postulated in Assumption 3.2.  $\blacksquare$

As a main consequence of the uniform stability exhibited by Algorithm 3, we obtain the following key result:

**Theorem 4.5.** *Fix  $s \geq 1$  and  $\delta \in (0, 1)$ . Given any dataset  $\mathcal{D}_s \in \Xi^s$ , there exists  $\varepsilon = \varepsilon(s, \delta) > 0$  such that  $\mathbf{x}^{K+1}$  is an  $\varepsilon$ -SNE of the SNEP in (10) with probability at least  $1 - \delta$ , where  $\mathbf{x}^{K+1}$  is obtained by iterating Algorithm 3  $K$ -times.  $\square$*

*Proof.* Once observed that measuring the gap between consecutive iterations on the decision variable  $\mathbf{x}$  through  $\|\mathbf{prox}_{\gamma g}(\mathbf{x}^{k+1} - \gamma \mathbb{F}(\mathbf{x}^{k+1})) - \mathbf{x}^{k+1}\|$  provides an equivalent information on the distance to a fixed point of the dynamics  $\mathbf{x}^{k+1} = \mathbf{prox}_{\gamma g}(\mathbf{x}^k - \gamma \mathbb{F}(\mathbf{x}^k))$ , the rest of the proof mimics the steps performed for deriving Theorem 3.4, thus ending up to upper bounding  $\|\mathbf{prox}_{\gamma g}(\mathbf{x}^{K+1} - \gamma \mathbb{F}(\mathbf{x}^{K+1})) - \mathbf{x}^{K+1}\|$  with the risk associated to (14).  $\blacksquare$

Theorem 4.5 then offers a finite-sample probabilistic certificate, in the form of a distance  $\varepsilon$  from the unique SNE, characterizing the output of Algorithm 3 obtained after arbitrary  $K$  iterations. This is crucial in real-world applications where available dataset may even be large, yet limited, and whose samples are only required to be i.i.d., with no restrictions on the probability  $\mathbb{P}$  according to which they are distributed.

**Remark 4.6.** *The uniform stability established in Lemma 3.7 offers an opportunity to extend our bounds for  $\varepsilon$ -SNE to stochastic generalized Nash equilibrium problems (SGNEPs). Here, the agents aim at solving mutually coupled optimization problems with shared constraints  $h : \mathbb{R}^n \rightarrow \mathbb{R}^m$ , i.e.,*

$$\forall i \in \mathcal{I} : \min_{x_i \in \Omega_i} \mathbb{J}_i(x_i, \mathbf{x}_{-i}) \text{ s.t. } h(x_i, \mathbf{x}_{-i}) \leq 0.$$

Also in this case, an equilibrium can be seen as a zero of the sum of the following two operators [7], [24], [38]:

$$\mathcal{A} : \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} \mapsto \begin{bmatrix} \partial g(\mathbf{x}) \\ \mathbb{N}_{\mathbb{R}_{\geq 0}^m}(\lambda) \end{bmatrix} + \begin{bmatrix} \nabla h(\mathbf{x})^\top \lambda \\ -h(\mathbf{x}) \end{bmatrix}, \quad \mathcal{B} : \begin{bmatrix} \mathbf{x} \\ \lambda \end{bmatrix} \mapsto \begin{bmatrix} \mathbb{F}(\mathbf{x}) \\ 0 \end{bmatrix},$$

where, typically, the (extended) mapping  $\mathcal{B}$  is cocoercive [38, Lemma 1]. Following similar arguments as in Lemma 3.7, we can then prove  $K$ -dependent uniform stability of a tailored data-driven FB for SGNEPs and establish a bound analogous to the one in Theorem 3.8.  $\square$

## V. NUMERICAL EXPERIMENTS

We now verify our theoretical bounds on several numerical examples. All simulations are run in MATLAB on a laptop with an Apple M2 chip featuring an 8-core CPU and 16 GB RAM.

### A. Plug-in electric vehicles charging coordination

To test the bound established in Theorem 4.5, a particular instance of that in Theorem 3.4, we adopt a stochastic version of a classic charging coordination problems for plug-in electric vehicles (PEVs), sketched next for completeness. Specifically, we consider a set of  $N$  PEVs populating a distribution grid, indexed by  $\mathcal{I} = \{1, \dots, N\}$ , where each agent aims to determine a day-ahead charging schedule subject to a stochastic cost of electricity [39] and few other fees. To this end, each PEV directly controls variable  $x_i \in \Omega_i \subseteq \mathbb{R}^T$  corresponding to the energy injection over a discrete interval  $\mathcal{T} = \{1, \dots, T\}$ . Each user then aims at minimizing a private cost in the form:

$$J_i(x_i, \sigma(\mathbf{x})) = \|x_i\|_{Q_i}^2 + c_i^\top x_i + \xi^\top x_i + \|\sigma(\mathbf{x}) - \bar{\sigma}\|_P^2, \quad (15)$$

TABLE I: Simulation parameters – §V-A

Param.	Description	Value
$T$	Time interval	14
$N$	Number of PEVs	20
$Q_i$	Weight matrix - Battery degradation	$\sim \mathcal{U}(0.002, 0.008) \cdot I_T$
$c_i$	Weight vector - Battery degradation	$\sim \mathcal{U}(0.02, 0.075) \cdot \mathbf{1}_T$
$P$	Weight matrix – Deviation from $\bar{\sigma}$	$I_{168}$
$\bar{\sigma}$	Aggregate reference signal	$\mathbf{1}_T$
$\bar{x}_i$	Power injection cap	2.5
$\zeta_i$	Minimum charging amount	$\sim \mathcal{U}(12, 18)$
$\mu_F$	Strong monotonicity constant	0.0127
$\kappa_F$	Lipschitz constant	0.1159
$M_F$	Upper bound – $\ F(\mathbf{x}, \xi)\ $	39.2192

where  $\|x_i\|_{Q_i}^2 + c_i^\top x_i$  models the  $i$ -th battery degradation cost, with  $0 \prec Q_i \in \mathbb{R}^{T \times T}$  and  $c_i \in \mathbb{R}^T$ , while  $\xi \in \mathbb{R}^T$  denotes the stochastic day-ahead price of energy, supported over some  $\Xi \subset \mathbb{R}^T$  with unknown probability distribution  $\mathbb{P}$ . Moreover,  $\sigma(\mathbf{x})$  represents the aggregate demand of the overall population of PEVs, usually defined as  $\sigma(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N x_i \in \Omega$ ,  $\Omega = \prod_{i \in \mathcal{I}} \Omega_i$ , whose deviation from some reference signal  $\bar{\sigma} \in \mathbb{R}^T$  is penalized through  $0 \prec P \in \mathbb{R}^{T \times T}$ . To complete the definition of our SNEP, we let  $\Omega_i = \{x_i \in [0, \bar{x}_i]^T : \mathbf{1}_T^\top x_i \geq \zeta_i\}$ . While each  $\bar{x}_i \geq 0$  limits the nonnegative power injection at every interval,  $\zeta_i \geq 0$  forces a minimum level of charging amount over the time horizon  $\mathcal{T}$  for the  $i$ -th user satisfaction.

Our numerical analysis is conducted with the values reported in Tab. I, where  $\mu_F$  and  $\kappa_F$  have been obtained analytically in view of the quadratic structure of the SNEP at hand, while  $M_F$  has been computed numerically. For the data-driven representation of the variable  $\xi$  affecting each cost in (15), we have collected from [40] ten years of day-ahead energy prices (in €/kWh) with granularity of one hour. In particular, the price data refer to the period January 5, 2015–December 31, 2024, i.e., since the platform [40] was launched, for a total of 3649 samples. As common in these type of coordination problems in which the flexible PEVs demand is usually analyzed during off-peak periods, we have focused on the interval 0:00am–1:00pm.

For comparison purposes only, we compute the unique SNE  $\mathbf{x}^*$  through a simplified version




Fig. 1: Left y-axis: Box plots capturing the relative approximation error  $\|\mathbf{x}^{1001} - \mathbf{x}^*\|$  produced by  $K = 1000$  iterations of Algorithm 3 over 50 trials considering different dataset size. Right y-axis: The resulting averaged bounds  $\varepsilon$  (green down-pointing triangles) offered by Theorem 4.5 with  $\delta = 0.05$ .

of [5, Alg. 1]. Besides turning each proximal operator  $\text{prox}_{\gamma g_i}(\cdot)$  into a projection mapping  $\text{proj}_{\Omega_i}(\cdot)$ —the same happens to Algorithm 3 indeed—at every iteration, [5, Alg. 1] relies on  $10k$  i.i.d. samples drawn from a surrogate representation of  $\Xi$  obtained by simply taking the minimum and maximum value over  $\mathcal{T}$  from the pool of 3649 available realizations. Note that, in the practical setting considered, the need of leveraging a limited dataset clashes with the theoretical conditions for exact convergence to the SNE required by [5, Alg. 1], which under our choice of selecting 10 i.i.d. samples per step would run for at most 364 iterations, thereby motivating our quest for finite sample guarantees.

First, we fix the number of iterations  $K = 1000$  and consider  $s \in \{100, 500, 1000, 2000, 3000\}$  to compare the theoretical bound  $\varepsilon$  provided by Theorem 4.5 with the actual distance  $\|\mathbf{x}^{1001} - \mathbf{x}^*\|$ , both terms scaled by  $\|\mathbf{x}^*\|$ . Specifically, we perform 50 trials for each of the  $s$  samples drawn from the pool of available 3649. In Fig. 1, where the box plots refer to the left y-axis, while the green down-pointing triangles to the right one, we report the numerical results obtained with confidence level  $\delta = 0.05$  and learning rate  $\gamma = 0.02$ , which produces  $\bar{\ell} = 24.3852$ . Although the performance in terms of SNE approximation is comparable, running Algorithm 3 with a small




Fig. 2: Top: Aggregate demand  $\sigma(\mathbf{x}^*)$  (blue line) and the approximated one related to  $\mathbf{x}^{1001}$  obtained by running Algorithm 3 with  $s = 3000$ , averaged over 50 trials (red line). Bottom: Mean value (yellow line) and standard deviation (yellow shaded area) of the energy price  $\xi$  over  $\mathcal{T}$ .

number of samples produces widely spread results with several outliers. However, this effect diminishes as the dataset size increases. This trend is consistently reflected in the theoretical bound  $\varepsilon$ . Specifically, while the average value of  $\varepsilon(100, 0.05)$  over 50 trials appears loose, increasing the dataset size results in tighter bounds. In particular, we can claim that the solution produced by Algorithm 3, run for  $K = 1000$  iterations with 3000 samples, produces an error in the approximation of  $\mathbf{x}^*$  of at most 30% with probability of at least 95%. The top plot in Fig. 2 quantifies such a statement by comparing the aggregate PEVs demand at the SNE and its approximation, averaged over the 50 trials, while the bottom plot reports the mean and standard deviation of the stochastic energy price.

Successively, we investigate the behaviour of Algorithm 3 for a fixed number of samples  $s = 3000$ ,  $\delta = 0.05$ ,  $\gamma = 0.02$ , and varying  $K \in \{100, 500, 1000, 5000, 10000\}$ . Quite interestingly, from Fig. 3 we observe that the theoretical bound established in Theorem 4.5 reveals an intrinsic dependency on  $K$  itself. In fact, considering a limited number of iterations produces an averaged bound that is substantially “dominated” by the empirical error




Fig. 3: Box plots capturing the relative approximation error  $\|\mathbf{x}^{K+1} - \mathbf{x}^*\|$  produced by different run  $K$  of Algorithm 3 over 50 trials with fixed number of samples. The green down-pointing triangles denote the resulting averaged bounds  $\varepsilon$  offered by Theorem 4.5 with  $\delta = 0.05$ .




Fig. 4: Aggregate demand  $\sigma(\mathbf{x}^*)$  (blue line) and its approximation through  $\mathbf{x}^{5001}$ , obtained by running Algorithm 3 with  $s = 3000$ . The solid red line denotes the mean of  $\sigma(\mathbf{x}^{5001})$ , while the shaded red area the associated standard deviation.

$\sum_{j=1}^s \|\mathbf{x}^{K+1} - \gamma F(\mathbf{x}^{K+1}, \xi^{(j)}) - \mathbf{y}^K\|$ , which progressively reduces its impact with larger values of  $K$ . In addition, note that with  $K \in \{5000, 10000\}$  one obtains a slight dispersion related to the relative approximation error computed—the resulting aggregate behaviour is compared in Fig. 4. This numerical evidence aligns with the necessity of imposing conditions on the approximation error variance for exact SNE computation—see, for instance, [4, Ass. 2(f), Ass. 3(d)], and [7,

Ass. 9, Prop. 1, Ass. 19].

**Remark 5.1.** *Our certificates are not expected to be tight. Probabilistic bounds based on stability arguments, indeed, work under weak assumptions but tend to be conservative due to their worst-case nature. Note that the gap observed in Fig. 1 and 3 is comparable to the one obtained with similar probabilistic bounds [41]–[43]. Nevertheless, sharpening strategies based on data-dependent bounds [44], [45] or sensitivity refinements [46] can be likewise adopted.  $\square$*

### B. An academic example

We now test the bound obtained in §III-B, explicit function of the iteration index  $K$ , on a generic optimization problem which also represent a typical example for VIs [47], thus naturally fitting our operator splitting setup. Let us then consider the following quadratic optimization problem:

$$\min_{x \in \mathcal{X}} q^\top x + \frac{1}{2} x^\top P x,$$

with  $\mathcal{X} \subseteq \mathbb{R}^n$ . The data vectors/matrices are  $q \in \mathbb{R}^n$  and  $P \in \mathbb{R}^{n \times n}$ ,  $P \succcurlyeq 0$ . Similar to (11), the problem of finding an optimizer  $x^*$  for the above can be rewritten as:

$$\mathbf{0}_n \in \underbrace{q + \bar{P}x}_{=: \mathcal{B}(x)} + \underbrace{N_{\mathcal{X}}(x)}_{=: \mathcal{A}(x)},$$

where  $\bar{P}$  denotes the symmetric part of  $P$ , i.e.,  $\bar{P} := \frac{1}{2}(P + P^\top) \in \mathbb{S}_{\neq 0}^n$ . In view of [2, Cor. 18.16], note that the operator  $\mathcal{B}$  is  $(1/\lambda_{\max}(\bar{P}))$ -cocoercive, and hence  $\gamma \in (0, 2/\lambda_{\max}(\bar{P}))$ .

To exemplify the case in which the entries of the matrix  $M$  are uncertain, we generate perturbed random matrices  $\mathcal{D}_s = \{\bar{P}^{(1)}, \dots, \bar{P}^{(s)}\}$ ,  $\bar{P}^{(i)} := \bar{P} + \text{diag}(\xi^{(i)}) \in \mathbb{S}^n$  for all  $i \in \{1, \dots, s\}$ , run Algorithm 2 and verify the result established in Theorem 3.8 numerically. Specifically, we consider  $n = 10$ , and generate  $\bar{P}$  as  $\bar{P} = Q\Lambda Q^\top$ , where  $Q \in \mathbb{R}^{10 \times 10}$  denotes the matrix of the QR decomposition of a random matrix with normally distributed entries, while the diagonal  $\Lambda \in \mathbb{R}^{10 \times 10}$  has elements linearly spaced in  $[0, 1]$ . This yields a cocoercivity constant  $1/\lambda_{\max}(Q\Lambda Q^\top) = 1$ . Note that, although symmetric with each element of the random parameter  $\xi$  sampled as  $\mathcal{N}(0, 0.5)$ , not all the produced samples  $\bar{P}^{(i)}$  are guaranteed to be positive semidefinite. In addition, each element of  $q$  is chosen at random according to  $\mathcal{N}(0, 0.5)$ , along with the constraint

TABLE II: Comparison results–Varying  $K$ 

$K$	100	500	1000	5000	10000
$\text{avg}(\Delta x_K) \times 10^{-3}$	443.8	18.2	1.1	1.3	1.4
$\text{avg}(\varepsilon)$	0.49	0.85	1.29	2.37	3.67

set  $\mathcal{X} = [0, a]^{10}$ , with  $a \sim \mathcal{U}(0, 2)$ . By considering  $\delta = 0.05$ , all the other parameters involved in Theorem 3.8 and not explicitly mentioned have been estimated numerically.

In this case we are interested in investigating how the dependency on  $K$  affects the underlying bound. To this end, we fix the number of samples  $s = 10000$  and vary  $K \in \{100, 500, 1000, 5000, 10000\}$ . The numerical results, averaged over 50 different trials with learning rate  $\gamma = 0.01$ , are reported in Tab. II, where  $\Delta x_K := \|x^{K+1} - x^*\|$  and some reference  $x^*$  is computed by using Gurobi [48]. As expected, the dependence on  $K$  in the stability bound from Lemma 3.7 results in rather loose certificates when  $K$  is large. Therefore, when the operator  $\mathcal{B}$  is merely cocoercive, it seems preferable to run Algorithm 2 for only a few iterations, e.g.,  $K \in [500, 1000]$ . Based on our numerical experience, this approach yields a solution  $x^{K+1}$  that is empirically close to some  $x^*$ , and it is accompanied by a tighter probabilistic certificate. This observation is consistent to what already pointed out in [16].

## VI. CONCLUSION AND OUTLOOK

By focusing on the problem of finding a zero of the sum of two operators, where one is either unavailable in closed form or computationally expensive to evaluate, we have derived rigorous certificates on the quality of solutions produced by data-based FB operator splitting schemes. As frequently happens in stochastic optimization, we have indeed approximated such an expensive operator by means of a finite number of samples. Since exact convergence to a zero should not be expected in this limited information setting, we have established probabilistic bounds on the distance between a true zero and the FB output, without making specific assumptions about the underlying data distribution. This has been made possible through a careful design of a tailored loss function representative for our purposes. We have then proved uniform stability of

the data-driven FB w.r.t. such a loss function under different monotonicity assumptions on the operators involved. Once derived computable expressions for an  $\varepsilon$ -zero that hold true with high confidence, we have finally applied our results to a popular SNE seeking algorithm based on the FB scheme.

Future work will analyze the stability properties for other operator splitting methods, e.g., Douglas-Rachford [2, Ch. 25], as well as possible randomized variants [16] of the resulting data-driven schemes. Aiming at improving our probabilistic bounds, one can also investigate the problem considered in this paper under different lenses, leveraging tools from PAC-value iteration or oracle-based complexity measures. Finally, along the line of [49], [50], possible relaxations of the i.i.d. requirement on the dataset at hand could also be explored.

## REFERENCES

- [1] E. K. Ryu and S. Boyd, “Primer on monotone operator methods,” *Applied and Computational Mathematics*, vol. 15, no. 1, pp. 3–43, 2016.
- [2] H. H. Bauschke and P. L. Combettes, *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2011.
- [3] H. Robbins and S. Monro, “A stochastic approximation method,” *The Annals of Mathematical Statistics*, pp. 400–407, 1951.
- [4] 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.
- [5] 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, 2021.
- [6] —, “Distributed projected-reflected-gradient algorithms for stochastic generalized Nash equilibrium problems,” in *2021 European Control Conference (ECC)*. IEEE, 2021, pp. 369–374.
- [7] —, “Stochastic generalized Nash equilibrium-seeking in merely monotone games,” *IEEE Transactions on Automatic Control*, vol. 67, no. 8, pp. 3905–3919, 2021.
- [8] 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.
- [9] 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, pp. 391–431, 2017.
- [10] A. Shapiro, “Monte Carlo sampling methods,” *Handbooks in Operations Research and Management Science*, vol. 10, pp. 353–425, 2003.
- [11] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on stochastic programming: Modeling and theory*. SIAM, 2021.
- [12] A. Alacaoglu and Y. Malitsky, “Stochastic variance reduction for variational inequality methods,” in *Conference on Learning Theory*. PMLR, 2022, pp. 778–816.
- [13] 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.

- [14] 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.
- [15] O. Bousquet and A. Elisseeff, “Stability and generalization,” *The Journal of Machine Learning Research*, vol. 2, pp. 499–526, 2002.
- [16] 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.
- [17] R. Bassily, V. Feldman, C. Guzmán, and K. Talwar, “Stability of stochastic gradient descent on nonsmooth convex losses,” *Advances in Neural Information Processing Systems*, vol. 33, pp. 4381–4391, 2020.
- [18] Y. Xing, Q. Song, and G. Cheng, “On the algorithmic stability of adversarial training,” *Advances in Neural Information Processing Systems*, vol. 34, pp. 26 523–26 535, 2021.
- [19] E. Zhao, T. Chavdarova, and M. Jordan, “Learning variational inequalities from data: Fast generalization rates under strong monotonicity,” *arXiv preprint arXiv:2410.20649*, 2024.
- [20] Y. Lei, Z. Yang, T. Yang, and Y. Ying, “Stability and generalization of stochastic gradient methods for minimax problems,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 6175–6186.
- [21] A. Ozdaglar, S. Pattathil, J. Zhang, and K. Zhang, “What is a good metric to study generalization of minimax learners?” *Advances in Neural Information Processing Systems*, vol. 35, pp. 38 190–38 203, 2022.
- [22] X. Wu, J. Zhang, and F.-Y. Wang, “Stability-based generalization analysis of distributed learning algorithms for big data,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 31, no. 3, pp. 801–812, 2019.
- [23] A. Tsiamis, I. Ziemann, N. Matni, and G. J. Pappas, “Statistical learning theory for control: A finite-sample perspective,” *IEEE Control Systems Magazine*, vol. 43, no. 6, pp. 67–97, 2023.
- [24] P. Yi and L. Pavel, “An operator splitting approach for distributed generalized Nash equilibria computation,” *Automatica*, vol. 102, pp. 111–121, 2019.
- [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] M. Vidyasagar, *Learning and generalisation: With applications to neural networks*. Springer Science & Business Media, 2013.
- [27] 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.
- [28] O. Bousquet, Y. Klochkov, and N. Zhivotovskiy, “Sharper bounds for uniformly stable algorithms,” in *Conference on Learning Theory*. PMLR, 2020, pp. 610–626.
- [29] C. McDiarmid, “On the method of bounded differences,” *Surveys in combinatorics*, vol. 141, no. 1, pp. 148–188, 1989.
- [30] M. D. Perlman, “Jensen’s inequality for a convex vector-valued function on an infinite-dimensional space,” *Journal of Multivariate Analysis*, vol. 4, no. 1, pp. 52–65, 1974.
- [31] M. C. Campi and S. Garatti, *Introduction to the scenario approach*. SIAM, 2018.
- [32] J. Lei and U. V. Shanbhag, “Stochastic Nash equilibrium problems: Models, analysis, and algorithms,” *IEEE Control Systems Magazine*, vol. 42, no. 4, pp. 103–124, 2022.
- [33] 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.
- [34] P. Palanisamy, “Multi-agent connected autonomous driving using deep reinforcement learning,” in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–7.
- [35] T. Liu, L. Cui, B. Pang, and Z.-P. Jiang, “A unified framework for data-driven optimal control of connected vehicles in mixed traffic,” *IEEE Transactions on Intelligent Vehicles*, vol. 8, no. 8, pp. 4131–4145, 2023.

- [36] X. Xu, Y. Jia, Y. Xu, Z. Xu, S. Chai, and C. S. Lai, “A multi-agent reinforcement learning-based data-driven method for home energy management,” *IEEE Transactions on Smart Grid*, vol. 11, no. 4, pp. 3201–3211, 2020.
- [37] M. Motalleb, A. Annaswamy, and R. Ghorbani, “A real-time demand response market through a repeated incomplete-information game,” *Energy*, vol. 143, pp. 424–438, 2018.
- [38] 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.
- [39] G. Liu, Y. Xu, and K. Tomsovic, “Bidding strategy for microgrid in day-ahead market based on hybrid stochastic/robust optimization,” *IEEE Transactions on Smart Grid*, vol. 7, no. 1, pp. 227–237, 2015.
- [40] E. Transparency Platform, “IT-CNORTH Day-ahead Prices,” <https://newtransparency.entsoe.eu/market/energyPrices>.
- [41] 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.
- [42] F. Fabiani, K. Margellos, and P. J. Goulart, “On the robustness of equilibria in generalized aggregative games,” in *2020 59th IEEE Conference on Decision and Control (CDC)*. IEEE, 2020, pp. 3725–3730.
- [43] —, “Probabilistic feasibility guarantees for solution sets to uncertain variational inequalities,” *Automatica*, vol. 137, p. 110120, 2022.
- [44] S. Shalev-Shwartz, O. Shamir, N. Srebro, and K. Sridharan, “Learnability, stability and uniform convergence,” *The Journal of Machine Learning Research*, vol. 11, pp. 2635–2670, 2010.
- [45] I. Kuzborskij and C. Lampert, “Data-dependent stability of stochastic gradient descent,” in *International Conference on Machine Learning*. PMLR, 2018, pp. 2815–2824.
- [46] Z. Deng, H. He, and W. Su, “Toward better generalization bounds with locally elastic stability,” in *International Conference on Machine Learning*. PMLR, 2021, pp. 2590–2600.
- [47] F. Facchinei and J. Pang, *Finite-Dimensional Variational Inequalities and Complementarity Problems*, ser. Springer Series in Operations Research and Financial Engineering. Springer New York, 2013.
- [48] Gurobi Optimization, LLC, “Gurobi Optimizer Reference Manual,” 2024. [Online]. Available: <https://www.gurobi.com>
- [49] M. Mohri and A. Rostamizadeh, “Stability bounds for stationary  $\varphi$ -mixing and  $\beta$ -mixing processes,” *Journal of Machine Learning Research*, vol. 11, no. 2, 2010.
- [50] R. R. Zhang, X. Liu, Y. Wang, and L. Wang, “McDiarmid-type inequalities for graph-dependent variables and stability bounds,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.