

---

# Stochastic Gradient Descent-Ascent and Consensus Optimization for Smooth Games: Convergence Analysis under Expected Co-coercivity

---

**Nicolas Loizou\***  
Mila and DIRO,  
Université de Montréal

**Hugo Berard**  
Mila and DIRO,  
Université de Montréal

**Gauthier Gidel†**  
Mila and DIRO,  
Université de Montréal

**Ioannis Mitliagkas†**  
Mila and DIRO,  
Université de Montréal

**Simon Lacoste-Julien†**  
Mila and DIRO,  
Université de Montréal

## Abstract

Two of the most prominent algorithms for solving unconstrained smooth games are the classical stochastic gradient descent-ascent (SGDA) and the recently introduced stochastic consensus optimization (SCO) [Mescheder et al., 2017]. SGDA is known to converge to a stationary point for specific classes of games, but current convergence analyses require a bounded variance assumption. SCO is used successfully for solving large-scale adversarial problems, but its convergence guarantees are limited to its deterministic variant. In this work, we introduce the *expected co-coercivity* condition, explain its benefits, and provide the first last-iterate convergence guarantees of SGDA and SCO under this condition for solving a class of stochastic variational inequality problems that are potentially non-monotone. We prove linear convergence of both methods to a neighborhood of the solution when they use constant step-size, and we propose insightful stepsize-switching rules to guarantee convergence to the exact solution. In addition, our convergence guarantees hold under the arbitrary sampling paradigm, and as such, we give insights into the complexity of minibatching.

## 1 Introduction

Motivated from the recent interest in solving adversarial formulations in machine learning such as generative adversarial networks (GANs) [Goodfellow et al., 2014], we consider in this paper a more abstract formulation of the problem and focus on solving the following *unconstrained* stochastic variational inequality (VI) problem:<sup>2</sup>

$$\text{Find } x^* \in \mathbb{R}^d \text{ such that } \xi(x^*) = \frac{1}{n} \sum_{i=1}^n \xi_i(x^*) = 0, \quad (1)$$

where each  $\xi_i : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is Lipschitz continuous. Further, we assume that the problem (1) has a unique<sup>3</sup> solution  $x^*$  and that the operator  $\xi$  is  $\mu$ -*quasi-strongly monotone*: there is a  $\mu \geq 0$  such that:

$$\langle \xi(x), x - x^* \rangle \geq \mu \|x - x^*\|^2 \quad \forall x \in \mathbb{R}^d \quad (2)$$

---

\*Corresponding author: nicolasloizou1@gmail.com. †Canada CIFAR AI Chair.

<sup>2</sup>While our presentation focuses on this finite-sum structure, most of our convergence results can easily be adapted to the general stochastic setting (see App. D). Also, we do not use the full power of variational inequalities that usually have constraints [Harker and Pang, 1990], but standard algorithms for (1) are coming from this literature [Gidel et al., 2018].

<sup>3</sup>This assumption can be relaxed; but for simplicity of exposition we enforce it.

If  $\mu = 0$ , then we say that  $\xi$  satisfies the variational stability condition:  $\langle \xi(x), x - x^* \rangle \geq 0$  [Hsieh et al., 2020]. In the variational inequality literature, condition (2) is also known as strong stability condition [Mertikopoulos and Zhou, 2019] or as strong Minty variational inequality (MVI) [Diakonikolas et al., 2021, Song et al., 2020].

Problem (1) generalizes the solution of several types of *stochastic smooth games* [Facchinei and Kanzow, 2007, Scutari et al., 2010, Mertikopoulos and Zhou, 2019]. The simplest example is the unconstrained min-max optimization problem (also called a *zero-sum* game):

$$\min_{x_1 \in \mathbb{R}^{d_1}} \max_{x_2 \in \mathbb{R}^{d_2}} \frac{1}{n} \sum_{i=1}^n g_i(x_1, x_2), \quad (3)$$

where each component function  $g_i : \mathbb{R}^{d_1} \times \mathbb{R}^{d_2} \rightarrow \mathbb{R}$  is assumed to be smooth. Here,  $\xi_i$  represents the appropriate concatenation of the block-gradients of  $g_i$ :  $\xi_i(x) := (\nabla_{x_1} g_i(x_1, x_2); -\nabla_{x_2} g_i(x_1, x_2))$ , where  $x := (x_1; x_2)$ . Solving (1) then amounts to finding a stationary point  $x^* = (x_1^*; x_2^*)$  for (3), which under a convex-concavity assumption for  $g_i$  for example, implies that it is a global solution for the min-max problem. More generally, we might seek the pure Nash equilibrium of a  $k$ -players game, where each player  $j$  is simultaneously trying to find the action  $x_j^*$  which minimizes with respect to  $x_j \in \mathbb{R}^{d_j}$  their own cost function  $\frac{1}{n} \sum_{i=1}^n f_i^{(j)}(x_j, x_{-j})$ , while the other players are playing  $x_{-j}$ , which represents  $x = (x_1, \dots, x_k)$  with the component  $j$  removed. Here,  $\xi_i(x)$  is the concatenation over all possible  $j$ 's of  $\nabla_{x_j} f_i^{(j)}(x_j, x_{-j})$ .

Such finite sum formulations appear in several machine learning applications such as generative adversarial networks (GANs) [Goodfellow et al., 2014], robust learning [Wen et al., 2014] or even some formulations of reinforcement learning [Pfau and Vinyals, 2016]. A standard algorithm that has been used to solve (1) is the stochastic version of the classical gradient method [Dem'yanov and Pevnyi, 1972, Nemirovski et al., 2009] or its variance reduced version [Balamurugan and Bach, 2016], that we call stochastic gradient descent-ascent (SGDA) in this paper.<sup>4</sup> More recently, Mescheder et al. [2017] analyzed some limitations of the gradient method in the context of GAN training and proposed an alternative efficient algorithm which could be used to solve (1) that they called consensus optimization (CO), which combines gradient updates with the minimization of  $\|\xi(x)\|^2$ . While the practical version of their algorithm for large  $n$  is stochastic (SCO, that randomly samples  $i$ 's) and displayed good performance [Mescheder et al., 2017], the only global convergence rate guarantees existing in the literature so far is only for the deterministic variant [Azizian et al., 2020, Abernethy et al., 2021].

The classical results from the stochastic VI literature are inappropriate for several reasons. First, a *uniform bound* over  $x$  on the variance  $\mathbb{E} [\|\xi_i(x) - \xi(x)\|^2]$  is typically assumed to get convergence guarantees (see e.g. Nemirovski et al. [2009], Gidel et al. [2018], Mertikopoulos and Zhou [2019], Yang et al. [2020], Lin et al. [2020b]), but this is not compatible with the *unconstrained* aspect of (1). For example, suppose  $g_i$  is a quadratic function, then the variance typically goes to infinity as  $x \rightarrow \infty$ . More appropriate relaxed assumptions have been considered to prove the convergence of other algorithms for (1) such as the stochastic extragradient method [Hsieh et al., 2020, Mishchenko et al., 2020] and its variance-reduced version [Chavdarova et al., 2019], or the stochastic Hamiltonian gradient method [Loizou et al., 2020], but not yet to the best of our knowledge for SGDA nor SCO. Second, the classical analysis for SGDA [Nemirovski et al., 2009] typically considers the convergence of the *average of the iterates* rather than for the *last-iterate*. However, as pointed out among others by Daskalakis et al. [2018] and Chavdarova et al. [2019], getting last-iterate convergence is important to apply the methods on potentially non-monotone problems such as GANs, where averaging is not appropriate. The only non-asymptotic last iterate convergence result for SGDA that we are aware of is Lin et al. [2020b], which focuses on a different class of problems (not assuming quasi-strong monotonicity) but it relies on strong assumptions on  $\mathbb{E} [\|\xi_i(x)\|^2]$  (see Section 4).

In this paper, we address both of these issues. We generalize the recent improved analysis of SGD [Gower et al., 2019] to the case of unconstrained stochastic variational inequality (1), and prove the last-iterate convergence for both SGDA and SCO without requiring any bounded variance assumption. We focus on quasi-strongly monotone VI problems, a class of structured non-monotone operators for which we are able to provide tight convergence guarantees and avoid the standard issues (cycling and divergence of the methods) appearing in the more general non-monotone regime.

<sup>4</sup>We use this suggestive name motivated from the min-max formulation (3), though we also call SGDA the simple update  $x^{k+1} = x^k - \alpha \xi_{i_k}(x^k)$  to solve (1) in the more general non-zero sum game scenario.

**Main Contributions.** The key contributions of this work are summarized as follows:

- We propose the *expected co-coercivity* (EC) assumption, which is the appropriate generalization of the expected smoothness assumption from Gower et al. [2019] to Problem (1). We explain the benefits of EC and show that it is strictly weaker than the bounded variance assumption and “growth conditions” previously used for the analysis of stochastic algorithms for (1).
- Using the EC assumption, we prove the first last-iterate convergence guarantees for stochastic gradient descent-ascent (SGDA) on (1) without any unrealistic noise assumption. We show a linear convergence rate to a neighborhood of  $x^*$  when constant step-size is used, and a  $O(1/k)$  rate to the exact solution when using a decreasing step-size rule. For the latter, we propose a theoretically motivated switching rule from a constant to a decreasing step-size to get faster convergence.
- Using the EC assumption, we provide the first convergence analysis of a stochastic variant (SCO) of the consensus optimization (CO) algorithm proposed by Mescheder et al. [2017] and previously used to train GANs. In particular, we prove last-iterate convergence for SCO, for both constant and decreasing step-sizes. As a corollary of our results, we obtain an improved convergence analysis for the deterministic CO. Furthermore, we explain how the update rule of the stochastic Hamiltonian gradient descent [Loizou et al., 2020] is a special case of the SCO and show that in this scenario, our analysis matches the theoretical guarantees presented in Loizou et al. [2020].
- Inspired by recent results from the optimization literature [Gower et al., 2019], we give the first stochastic reformulation of the variational inequality problem (1) which enables us to provide convergence guarantees of SGDA and SCO under the *arbitrary sampling paradigm* [Richtárik and Takáč, 2016]. This allows us to give insights into the complexity of minibatching.

## 2 Arbitrary Sampling: Stochastic Reformulation of Problem (1)

In this work, we provide theorems through which we can analyze all minibatch variants of the two algorithms under study, SGDA and SCO. To do this, we construct a so-called “stochastic reformulation” of the variational inequality problem (1). Our approach is inspired by recently proposed stochastic reformulations of standard optimization problems, like the empirical risk minimization in Gower et al. [2019] and linear systems in Richtárik and Takáč [2020], Loizou and Richtárik [2020a,b].

In each step of our algorithms, we assume we are given access to unbiased estimates  $g(x) \in \mathbb{R}^d$  of the operator such that  $\mathbb{E}[g(x)] = \xi(x)$ . For example, we can use a minibatch to form an estimate of the operator such as  $g(x) = \frac{1}{b} \sum_{i \in S} \xi_i(x)$ , where  $S \subset \{1, \dots, n\}$  will be chosen uniformly at random and  $|S| = b$ . To allow for any form of minibatching, we use the *arbitrary sampling* notation  $g(x) = \xi_v(x) := \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x)$ , where  $v \in \mathbb{R}_+^n$  is a random *sampling vector* drawn from some distribution  $\mathcal{D}$  such that  $\mathbb{E}_{\mathcal{D}}[v_i] = 1$ , for  $i = 1, \dots, n$ . Note that the unbiasedness follows immediately from this definition of the sampling vector:  $\mathbb{E}_{\mathcal{D}}[\xi_v(x)] = \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathcal{D}}[v_i] \xi_i(x) = \xi(x)$ .

Thus, with each user-defined distribution  $\mathcal{D}$ , we are able to introduce a stochastic reformulation of problem (1) as follows:

$$\text{Find } x^* \in \mathbb{R}^d \text{ such that } \mathbb{E}_{\mathcal{D}} \left[ \xi_v(x^*) := \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x^*) \right] = 0. \quad (4)$$

Since  $\xi_v(x)$  as an unbiased estimate of the operator  $\xi(x)$ , we can now use stochastic (simultaneous) gradient descent-ascent (SGDA) to solve (4) as follows:

$$x^{k+1} = x^k - \alpha_k \xi_{v^k}(x^k), \quad (5)$$

where  $v^k \sim \mathcal{D}$  is sampled i.i.d at each iteration and  $\alpha_k > 0$  is a stepsize. We highlight that in our analysis, we allow to select *any* distribution  $\mathcal{D}$  that satisfies  $\mathbb{E}_{\mathcal{D}}[v_i] = 1 \forall i$ , and for different selection of  $\mathcal{D}$ , (5) yields different interpretation as an SGDA method for solving the original problem (1).

In this work, we mostly focus on the  $b$ -minibatch sampling, however note that our analysis holds for every form of minibatching and for several choices of sampling vectors  $v$ .

**Definition 2.1** (Minibatch sampling). Let  $b \in [n]$ . We say that  $v \in \mathbb{R}^n$  is a  $b$ -minibatch sampling if for every subset  $S \in [n]$  with  $|S| = b$ , we have that  $\mathbb{P} \left[ v = \frac{n}{b} \sum_{i \in S} e_i \right] = 1 / \binom{n}{b} := \frac{b!(n-b)!}{n!}$

By using a double counting argument, one can show that if  $v$  is a  $b$ -minibatch sampling, it is also a valid sampling vector ( $\mathbb{E}_{\mathcal{D}}[v_i] = 1$ ) [Gower et al., 2019]. See Gower et al. [2019] for other choices of sampling vectors  $v$ .

### 3 Expected Co-coercivity and Connection to Other Assumptions

Before introducing the condition of expected co-coercivity, we first review some details on co-coercivity, an intermediate notion between monotonicity and strong monotonicity [Zhu and Marcotte, 1996], and explain where it belongs as assumption in the literature of variational inequalities and min-max optimization.

**Co-coercive operators.** The co-coercive condition is relatively standard in operator splitting literature [Davis and Yin, 2017, Vũ, 2013] and for variational inequalities [Zhu and Marcotte, 1996]. It was used to analyze the celebrated forward-backward algorithm (a.k.a, proximal gradient) [Lions and Mercier, 1979, Chen and Rockafellar, 1997, Palaniappan and Bach, 2016] that is known not to converge for general monotone operators [Bauschke et al., 2011].

**Definition 3.1** (Co-coercivity / Co-coercive around  $w^*$ ). We say that an operator  $\xi$  is  $\ell$ -co-coercive if there exist  $\ell > 0$  such that,<sup>a</sup>  $\|\xi(x) - \xi(y)\|^2 \leq \ell \langle \xi(x) - \xi(y), x - y \rangle \quad \forall x, y \in \mathbb{R}^d$ . If there exist  $w^* \in \mathbb{R}^d$  and  $\ell > 0$  such that  $\|\xi(x) - \xi(w^*)\|^2 \leq \ell \langle \xi(x) - \xi(w^*), x - w^* \rangle \quad \forall x \in \mathbb{R}^d$ . then we say that the operator  $\xi$  is  $\ell$ -co-coercive around  $w^*$ . Note that in the last definition the point  $w^*$  is not necessarily a point where  $\xi(w^*) = 0$ .

Note that from Cauchy-Swartz’s inequality, one can get that a  $\ell$ -co-coercive operator is  $\ell$ -Lipschitz. In single-objective minization, one can show the converse statement by using convex duality. Thus, a gradient of a function is  $L$ -co-coercive if and only if the function is convex and  $L$ -smooth (i.e.  $L$ -Lipschitz gradients) [Bauschke et al., 2011]. However, in general, a  $L$ -Lipchitz operator is *not*  $L$ -co-coercive. What we can show instead is that a  $L$ -Lipschitz and  $\mu$ -strongly monotone operator is  $\ell$ -co-coercive with  $\ell \in [L, L^2/\mu]$  [Facchinei and Pang, 2007]. Note that both ranges of the spectrum may occur. For instance, Chavdarova et al. [2019] present a sufficient condition in zero-sum games to have  $\ell = O(L)$ . Note that one can easily show that a sum of co-coercive operators is also co-coercive. Let us now provide a proposition summarizing the implications between (strong) monotonicity and co-coercivity.

**Proposition 3.2.** For a  $L$ -Lipschitz operator  $\xi$ , the following implications hold:

$$\begin{array}{ccccc}
 \mu\text{-strongly monotone} & \implies & \frac{L^2}{\mu}\text{-co-coercive} & \implies & \text{monotone} \\
 \downarrow & & \downarrow & & \downarrow \\
 \mu\text{-quasi-strongly monotone} & \implies & \frac{L^2}{\mu}\text{-co-coercive around } x^* & \implies & \text{variational} \\
 & & & & \text{stability condition}
 \end{array}$$

Let us also note that while a  $\ell$ -co-coercive operator is always  $\ell$ -Lipschitz continuous, it is possible for an operator to be  $\ell$ -co-coercive around  $x^*$  and *not* be Lipschitz continuous. This highlights the wider applicability of the  $\ell$ -co-coercivity around  $x^*$  assumption that is all we need for several of our convergence results, in contrast to the Lipschitz continuity of  $\xi$  which is typically assumed in the variational inequality literature. In Appendix A.6, we provide such example of a  $\mu$ -quasi strongly monotone operator that is  $\ell$ -co-coercive around  $x^*$ , which is not monotone nor Lipschitz continuous.

#### 3.1 Expected Co-coercivity (EC)

In our analysis of SGDA and SCO, we rely on a generic and remarkably weak assumption that we call *expected co-coercivity* (EC). In this section, we formally define EC, provide sufficient conditions for it to hold and relate it to the existing gradient assumptions.

**Assumption 3.3** (Expected Cocoercivity). We say that  $\xi$  is  $\ell_\xi$ -co-coercive in expectation with respect to a distribution  $\mathcal{D}$  if there exists  $\ell_\xi > 0$  such that

$$\mathbb{E}_{\mathcal{D}} \left[ \|\xi_v(x) - \xi_v(x^*)\|^2 \right] \leq \ell_\xi \langle \xi(x), x - x^* \rangle \quad \forall x \in \mathbb{R}^d. \quad (\text{EC})$$

For simplicity, we will write  $\xi \in EC(\ell_\xi)$  to say that EC holds and we will refer to  $\ell_\xi$  as the expected co-coercivity constant.

<sup>a</sup>Note that in our definition we consider the inverse of the co-coercive constant from Lions and Mercier [1979] which is the constant  $\ell$  such that  $\langle \xi(x) - \xi(y), x - y \rangle \geq \ell \|\xi(x) - \xi(y)\|^2$ .

The convergence results in this paper will depend on the following operator noise at  $x^*$  that is finite for any reasonable sampling distribution  $\mathcal{D}$  for the sampling vector  $v$ :

$$\sigma^2 := \mathbb{E}_{\mathcal{D}}[\|\xi_v(x^*)\|^2] < \infty. \quad (6)$$

As we discuss below, common assumptions used to prove convergence of stochastic algorithms for solving the VI problem is uniform boundedness of the stochastic operator  $\mathbb{E}\|\xi_i(x)\|^2 \leq c$  or uniform boundedness of the variance  $\mathbb{E}\|\xi_i(x) - \xi(x)\|^2 \leq c$ . However these assumptions either do not hold or are true only for restrictive set of problems. In our work we do not assume such bounds. Instead we use the following direct consequence of Assumption 3.3.

**Lemma 3.4.** If  $\xi \in EC(\ell_\xi)$ , then  $\mathbb{E}\|\xi_v(x)\|^2 \leq 2\ell_\xi \langle \xi(x), x - x^* \rangle + 2\sigma^2$ .

Let us now provide some more familiar sufficient conditions which guarantee that the EC condition holds and give closed form expression for the expected co-coercivity parameter.

**Proposition 3.5.** Let  $\xi_i$  be  $\ell_i$  co-coercive (or  $\ell_i$  co-coercive around  $x^*$ ), then  $\xi \in EC(\ell_\xi)$ . Let  $\ell_{\max} = \max\{\ell_i\}_{i=1}^n$  and  $\ell$  be the co-coercive constant of  $\xi$ , if we let  $v$  to be a  $b$ -minibatch sampling, then  $\ell_\xi = \frac{n}{b} \frac{b-1}{n-1} \ell + \frac{1}{b} \frac{n-b}{n-1} \ell_{\max}$  and  $\sigma^2 = \frac{1}{b} \frac{n-b}{n-1} \sigma_1^2$ , where  $\sigma_1^2 := \frac{1}{n} \sum_{i=1}^n \|\xi_i(x^*)\|^2$ .

In the above Proposition 3.5, we show how co-coercivity of  $\xi_i$  implies expected co-coercivity. However, the opposite implication does not necessarily hold. Indeed the expected co-coercivity can hold even when we do not assume that  $\xi_i$  are co-coercive, as we show in the next proposition.

**Proposition 3.6.** Let  $\xi$  be quasi-strongly monotone and let  $\xi_i$  be  $L_i$ -Lipschitz continuous for all  $i \in [n]$ . Then  $\xi \in EC(\ell_\xi)$ .

**Connection to Other Assumptions** In the optimization literature, the standard convergence analysis of stochastic gradient algorithms like SGD relied on bounded gradient ( $\mathbb{E}\|\nabla f_i(x)\|^2 \leq c$ ) or bounded variance assumptions ( $\mathbb{E}\|\nabla f_i(x) - \nabla f(x)\|^2 \leq c$ ) [Recht et al., 2011, Hazan and Kale, 2014, Rakhlin et al., 2012] or growth condition ( $\mathbb{E}\|\nabla f_i(x)\|^2 \leq c_1 \|\nabla f(x)\|^2 + c_2$ ) [Bottou et al., 2018, Schmidt et al., 2017]. However, a recent line of work shows that these assumptions might be restrictive or never be satisfied<sup>5</sup> and proposed alternative conditions [Nguyen et al., 2018, Vaswani et al., 2018, Gower et al., 2019, 2021, Khaled et al., 2020, Khaled and Richtárik, 2020, Assran et al., 2019, Koloskova et al., 2020, Patel and Zhang, 2021, Loizou et al., 2020, 2021]. One of the weakest assumptions used for the convergence analysis of SGD in the smooth setting, is expected smoothness (ES) proposed in Gower et al. [2019] (see last row of Table 1). Our expected co-coercivity condition (EC) can be seen as the generalization of ES in the operator setting.

In the literature of stochastic methods for solving the variational inequality problem and min-max optimization problem, similar assumptions have been made. In particular for the analysis of stochastic algorithms, papers assume either bounded operators [Nemirovski et al., 2009, Abernethy et al., 2021] or bounded variance [Juditsky et al., 2011, Yang et al., 2020, Lin et al., 2020a, Luo et al., 2020, Tran Dinh et al., 2020] and growth condition [Lin et al., 2020b]. In all of the these conditions, the values of parameters  $c$ ,  $c_1$  and  $c_2$  (see Table 1) usually do not have a closed form expression – they are simply assumed to exist. However, to the best of our knowledge, there is no analysis using a concept similar to our expected co-coercivity. All existing analyses of SGDA for (quasi)-strongly monotone and co-coercive operators require the much stronger extra assumptions of “bounded noise” or “bounded variance” to guarantee convergence, while for SCO, there are no known convergence guarantees in the literature. Note that through Lemma 3.4, (EC) implies bounds on the gradient with closed-form problem-depended expressions for these constants. We also mention that other appropriate relaxed assumptions have been considered to prove the convergence of algorithms for (1) [Hsieh et al., 2020, Mishchenko et al., 2020, Chavdarova et al., 2019, Loizou et al., 2020], but not yet for SGDA nor SCO. For a wider literature review in the area, see Appendix E.

In Table 1, we illustrate the correspondence between conditions used in the stochastic optimization literature and the stochastic VI problem. For further connections between these conditions, see Appendix A.5. There, for example, we show why assuming bounded gradients together with strong monotonicity lead to an empty set of operators and explain why ES and EC (see last row of Table 1) are equivalent for convex and smooth single-objective optimization problems.

<sup>5</sup>For example, the bounded gradient assumption and strong convexity contradict each other in the unconstrained setting (see [Nguyen et al., 2018] for more details).

Table 1: Correspondence of Assumptions between Optimization and Variational Inequalities

Assumptions	Stochastic Optimization $\min_x f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$	Stochastic Variational Inequality Find $x^*$ such that $\xi(x^*) = \frac{1}{n} \sum_{i=1}^n \xi_i(x^*) = 0$
Bounded Gradient	$\mathbb{E} \ \nabla f_i(x)\ ^2 \leq c$	$\mathbb{E} \ \xi_i(x)\ ^2 \leq c$
Bounded Variance	$\mathbb{E} \ \nabla f_i(x) - \nabla f(x)\ ^2 \leq c$	$\mathbb{E} \ \xi_i(x) - \xi(x)\ ^2 \leq c$
Growth Condition	$\mathbb{E} \ \nabla f_i(x)\ ^2 \leq c_1 \ \nabla f(x)\ ^2 + c_2$	$\mathbb{E} \ \xi_i(x)\ ^2 \leq c_1 \ \xi(x)\ ^2 + c_2$
Expected Smoothness (ES) / Expected Cocoercivity (EC)	$\mathbb{E} [\ \nabla f_v(x) - \nabla f_v(x^*)\ ^2] \leq 2\mathcal{L}(f(x) - f(x^*))$	$\mathbb{E} [\ \xi_v(x) - \xi_v(x^*)\ ^2] \leq \ell_\xi \langle \xi(x), x - x^* \rangle$

## 4 Stochastic Gradient Descent-Ascent

Having presented the update rule of SGDA (5) for solving the stochastic reformulation (4) of the original unconstrained stochastic variational inequality problem (1), let us now provide theorems for its convergence guarantees. We highlight that our theorems hold for any selection of distributions  $\mathcal{D}$  over the random sampling vectors  $v$  and as such they are able to describe the convergence of an infinite array of variants of SGDA each of which is associated with a specific probability law governing the data selection rule used to form minibatches.

**Theorem 4.1** (Constant Step-size). Assume that  $\xi$  is  $\mu$ -quasi strongly monotone and that  $\xi \in EC(\ell_\xi)$ . Choose  $\alpha_k = \alpha \leq \frac{1}{2\ell_\xi}$  for all  $k$ . Then, the iterates of SGDA, given by (5), satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \alpha\mu)^k \|x^0 - x^*\|^2 + \frac{2\alpha\sigma^2}{\mu}, \quad (7)$$

Note that we do not assume that  $\xi$  or  $\xi_i$  are monotone operators in Theorem 4.1. SGDA converges by only assuming that  $\xi$  is quasi-strongly monotone and that EC holds. Theorem 4.1 states that SGDA converges linearly to a neighborhood of  $x^*$  which is proportional to the step-size  $\alpha$  and the noise at the optimum  $\sigma^2$ . We highlight, that since we control distribution  $\mathcal{D}$  we also control the values of  $\ell_\xi$  and  $\sigma^2$ , and in the case of  $b$ -minibatch sampling these values have a closed-form expressions as shown in Proposition 3.5. To the best of our knowledge, Theorem 4.1 is the first last-iterate non-asymptotic convergence guarantee for SGDA for solving quasi-strongly monotone problems without assuming extra conditions on the noise. It is worth mentioning that Lin et al. [2020b] also prove last-iterate convergence of SGDA for different class of problems (they do not assume quasi-strong monotonicity), but the proposed analysis requires much stronger noise conditions. In particular, a bound on the variance with vanishing constants is needed, which, as far as we know, can only be satisfied by running SGDA with growing mini-batch size [Friedlander and Schmidt, 2012] (see also App. E for a more detailed discussion). To highlight further the generality of Theorem 4.1, we note that for the deterministic GDA,  $\sigma^2 = 0$ . Thus, we can obtain the following corollary.

**Corollary 4.2** (Deterministic GDA). Let all assumptions of Theorem 4.1 be satisfied. Let  $|S| = n$  with probability one (each iteration of SGDA uses a full batch gradient). Then by selecting  $\alpha_k = \alpha \leq \frac{1}{2\ell}$  for all  $k$ , the iterates of deterministic GDA satisfy:  $\|x^k - x^*\|^2 \leq (1 - \alpha\mu)^k \|x^0 - x^*\|^2$ .

Even if Corollary 4.2 looks trivial, to the best of our knowledge, Theorem 4.1 is the first convergence theorem of SGDA that includes the deterministic GDA originally provided by Chen and Rockafellar [1997] as a special case.

**Optimal  $b$ -Minibatch Size:** Using standard computations, the convergence rate presented in Theorem 4.1 can be equivalently expressed as iteration complexity result as follows: If we are given any accuracy  $\epsilon > 0$ , choosing stepsize  $\alpha = \min \left\{ \frac{1}{2\ell_\xi}, \frac{\epsilon\mu}{4\sigma^2} \right\}$  and  $k \geq \max \left\{ \frac{2\ell_\xi}{\mu}, \frac{4\sigma^2}{\epsilon\mu^2} \right\} \log \left( \frac{2\|x^0 - x^*\|^2}{\epsilon} \right)$ , implies  $\mathbb{E} \|x^k - x^*\|^2 \leq \epsilon$ . By combining the lower bound on  $k$  with the expressions of  $\ell_\xi$  and  $\sigma^2$  of Proposition 3.5, we have that the iteration complexity (by ignoring the logarithmic terms) becomes  $k \geq \frac{2}{\mu} \max \left\{ \ell_\xi, \frac{2\sigma^2}{\epsilon\mu} \right\}$  where  $\ell_\xi = \frac{n}{b} \frac{b-1}{n-1} \ell + \frac{1}{b} \frac{n-b}{n-1} \ell_{\max}$  and  $\sigma^2 = \frac{1}{b} \frac{n-b}{n-1} \sigma_1^2$ . Thus, the total complexity of the algorithm as a function of the minibatch size  $b$  is given by  $TC(b) \leq \frac{2}{\mu} \max \{ b\ell_\xi, b \frac{2\sigma^2}{\epsilon\mu} \}$ . By following the same steps with Gower et al. [2019], it can be shown that  $b\ell_\xi$  is linearly increasing term in  $b$  while  $b \frac{2\sigma^2}{\epsilon\mu}$  is a linearly decreasing term in  $b$ . Hence, if we define  $b^*$  to be the minibatch size that minimize the total complexity  $TC(b)$  (optimal  $b$ -Minibatch Size) we have that if  $\sigma_1^2 \leq \ell_{\max}$  then  $b^* = 1$  otherwise  $b^* = n \frac{\ell - \ell_{\max} + \frac{2}{\epsilon\mu} \cdot \sigma_1^2}{n\ell - \ell_{\max} + \frac{2}{\epsilon\mu} \cdot \sigma_1^2}$ .

In the next theorem, we provide an insightful stepsize-switching rule that describes when one should switch from a constant to a decreasing step-size regime to guarantee convergence to  $x^*$  and not to a neighborhood, providing the first convergence analysis of SGDA under such a switching rule.

**Theorem 4.3.** Assume  $\xi$  is  $\mu$ -quasi-strongly monotone and that  $\xi \in EC(\ell_\xi)$ . Let  $\mathcal{K} := \ell_\xi / \mu$  and let  $\alpha_k = \frac{1}{2\ell_\xi}$  for  $k \leq 4\lceil \mathcal{K} \rceil$  and  $\alpha_k = \frac{2k+1}{(k+1)^2\mu}$  for  $K > 4\lceil \mathcal{K} \rceil$ . If  $k \geq 4\lceil \mathcal{K} \rceil$ , then iterates of SGDA, given by (5) satisfy:

$$\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma^2}{\mu^2} \frac{8}{k} + \frac{16\lceil \mathcal{K} \rceil^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right). \quad (8)$$

## 5 Stochastic Consensus Optimization

The Consensus Optimization (CO) Algorithm is a computationally-light<sup>6</sup> second order methods which has been introduced in Mescheder et al. [2017] and it was shown to be an effective method for training GANs in a variety of settings. Liang and Stokes [2019] show first that CO converges linearly in the bilinear case. Abernethy et al. [2021] show that CO can be viewed as a perturbation of the deterministic Hamiltonian gradient descent (HGD) and explain how CO converges at the same rate as HGD, while Azizian et al. [2020] prove convergence of CO for  $\mu$ -strongly monotone operators with positive singular values of the Jacobian matrix  $J = \nabla\xi$ .

However, even if CO is used explicitly in the stochastic setting and in practice only minibatch variants are implemented, to the best of our knowledge all existing analysis focus only on the deterministic setting. Thus, our work is the first that provide convergence guarantees for the Stochastic Consensus Optimization (SCO) and due to our framework, our analysis includes the convergence of the deterministic update as special case. Impressively, our analysis provides tighter rates than previous analysis even in the deterministic setting.

For the results of this section, we assume that each  $\xi_i$  in problem (1) is differentiable. That is, we have access to the Jacobian matrices  $\mathbf{J}_i(x) = \nabla\xi_i(x)$ . Following Loizou et al. [2020], in our analysis we will also assume that the Hamiltonian function is quasi-strongly convex and that it satisfies expected smoothness (see last row of Table 1). These are not strong assumptions, and as an example, they are satisfied for smooth bilinear min-max optimization problems Loizou et al. [2020]. In Appendix C, by extending the results of Loizou et al. [2020], we explain how these assumptions can be satisfied for the quadratic min-max problems.

### 5.1 Setting

The consensus optimization (CO) algorithm as presented in Mescheder et al. [2017] has the following update rule:

$$x^{k+1} = x^k - \alpha\xi(x^k) - \gamma\nabla\mathcal{H}(x^k) \quad (9)$$

where  $\mathcal{H}(x) = \frac{1}{2}\|\xi(x)\|^2$  is the Hamiltonian function and  $\alpha, \gamma > 0$  are the step-sizes. From its definition, it is clear that the update rule is essentially a weighted combination of GDA and the Hamiltonian gradient descent (HGD) of Balduzzi et al. [2018]. In practice, implementing CO in a mini-batch setting (stochastic) leads to biased estimates of the gradient of the Hamiltonian function [Mescheder et al., 2017]. This is one of the main reason that existing analysis was not able to capture the behavior of the method in the stochastic setting. However, recently Loizou et al. [2020] proposed a way to obtain unbiased estimators of the gradient of the Hamiltonian function by expressing the Hamiltonian of a stochastic game as a finite-sum problem. In this work, we adopt the finite-sum structure and the unbiased estimators proposed in Loizou et al. [2020] for the Hamiltonian part and we extend the formulation to capture the arbitrary sampling paradigm. That is, the Hamiltonian function can be expressed as,  $\mathcal{H}(x) = \frac{1}{2}\|\xi(x)\|^2 = \frac{1}{n} \sum_{i=1}^n \frac{1}{n} \sum_{j=1}^n \frac{1}{2} \langle \xi_i(x), \xi_j(x) \rangle$ . In addition, by following the stochastic reformulation setting presented in Section 2, let us have two independent random sampling vectors  $u \sim \mathcal{D}$  and  $v \sim \mathcal{D}$  and let us define:  $\mathcal{H}_{u,v}(x) = \frac{1}{n} \sum_{i=1}^n \frac{1}{n} \sum_{j=1}^n \frac{1}{2} \langle u_i \xi_i(x), v_j \xi_j(x) \rangle$ . Since vectors  $u$  and  $v$  are *independent* sampling vectors, it is clear that  $\mathbb{E}_{u,v}[\mathcal{H}_{u,v}(x)] = \mathcal{H}(x)$ , where  $\mathbb{E}_{u,v}$  denotes the expectation with respect to distribution  $\mathcal{D}$  on both vectors  $u$  and  $v$ . Let us also

<sup>6</sup>At each step, CO requires only the computation of a Jacobian-vector product which can be efficiently evaluated in tasks like training neural networks with comparable computation time of a gradient [Pearlmutter, 1994].

use  $\mathbf{J}_i(x) = \nabla \xi_i(x)$  to express the Jacobian matrices for  $i$ , then the gradient of  $\mathcal{H}_{u,v}(x)$  has the following form:

$$\nabla \mathcal{H}_{u,v}(x) = \frac{1}{2} [\mathbf{J}_u^\top(x) \xi_v(x) + \mathbf{J}_v^\top(x) \xi_u(x)] \quad (10)$$

where  $\mathbf{J}_u(x) = \frac{1}{n} \sum_{i=1}^n u_i \mathbf{J}_i(x)$  and  $\mathbf{J}_v(x) = \frac{1}{n} \sum_{i=1}^n v_i \mathbf{J}_i(x)$ . Similar to Loizou et al. [2020], it can be shown that  $\nabla \mathcal{H}_{u,v}(x)$  is an unbiased estimator of  $\nabla \mathcal{H}(x) = \mathbf{J}^\top(x) \xi(x)$ . That is,  $\mathbb{E}_{u,v}[\nabla \mathcal{H}_{u,v}(x)] = \nabla \mathcal{H}(x)$ . Throughout this section, we will denote the gradient noise of the stochastic Hamiltonian function with  $\sigma_{\mathcal{H}}^2 := \mathbb{E}_{u,v}[\|\nabla \mathcal{H}_{u,v}(x^*)\|^2]$ , which is finite for any reasonable sampling distribution  $\mathcal{D}$ .

## 5.2 Stochastic Consensus Optimization and its Special Cases

Having explained the basic setting and background on consensus optimization algorithm and the Hamiltonian function, let us now present as Algorithm 1 the proposed stochastic consensus optimization (SCO) algorithm. Note that in each iteration  $k$ , two random sampling vectors  $u$  and  $v$  are sampled independently from a user-defined distribution  $\mathcal{D}$ . These vectors are used to evaluate  $\xi_{v^k}$  and  $\nabla \mathcal{H}_{v^k, u^k}(x^k)$ , the unbiased estimators of  $\xi(x^k)$  and  $\nabla \mathcal{H}(x^k)$  at point  $x^k$  respectively. Note also that the update rule of SCO

---

### Algorithm 1 Stochastic Consensus Opt. (SCO)

---

**Input:** Starting step-size  $\alpha_0, \gamma_0 > 0$ . Choose initial point  $x^0 \in \mathbb{R}^d$ . Distribution  $\mathcal{D}$  of samples.

**for**  $k = 0, 1, 2, \dots, K$  **do**

Sample independently  $v^k \sim \mathcal{D}$  and  $u^k \sim \mathcal{D}$

Set step-sizes  $\alpha_k$  and  $\gamma_k$  according to a pre-selected step-size rule

Set  $x^{k+1} = x^k - \alpha_k \xi_{v^k}(x^k) - \gamma_k \nabla \mathcal{H}_{v^k, u^k}(x^k)$

**end for**

**Output:** The last iterate  $x^k$

---

is a weighted combination of SGDA (5) and the stochastic Hamiltonian gradient descent (SHGD) of Loizou et al. [2020]. Thus, it is clear, that if one selects  $\alpha^k = 0, \forall k > 0$  then the method is equivalent to SHGD and if  $\gamma^k = 0, \forall k > 0$  then the method becomes equivalent to the SGDA. In addition if we select sampling vectors  $u = v = (1, 1, \dots, 1) \in \mathbb{R}^n$  with probability 1, then from the definition of  $\xi_{v^k}$  and  $\nabla \mathcal{H}_{v^k, u^k}(x^k)$  we obtain the deterministic CO (9) as special case of our update rule.

## 5.3 Convergence Analysis

Let us now present our main theoretical results describing the performance of SCO. Similar to the previous section, we provide two main theorems for two different step-size selection.

**Theorem 5.1 (Constant Step-size).** Assume  $\xi$  is  $\mu$ -quasi-strongly monotone with  $\mu \geq 0$  and that  $\xi \in EC(\ell_\xi)$ . Assume that the Hamiltonian function  $\mathcal{H}$  is  $\mu_{\mathcal{H}}$ -quasi strongly convex and  $\mathcal{L}_{\mathcal{H}}$ -expected smooth. Then, for  $\gamma_k = \gamma \leq \frac{1}{4\mathcal{L}_{\mathcal{H}}}$  and  $\alpha_k = \alpha \leq \frac{1}{4\ell_\xi}$ , the iterates of SCO satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}} - \alpha\mu)^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu_{\mathcal{H}} + \alpha\mu}. \quad (11)$$

If  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}})^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu_{\mathcal{H}}}.$$

Theorem 5.1 is quite informative, as it highlights that both SGDA and SHG parts should coexist in the update rule of the SCO to guarantee faster convergence for  $\mu$ -quasi-strongly monotone operators (i.e. both step-sizes  $\alpha$  and  $\gamma$  should be positive up to specific values). However, if  $\xi$  simply satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$  (when  $\mu = 0$ ), then the convergence rate of SCO, does not depend on the step-size  $\alpha$  (the SGDA part) and the neighborhood of convergence  $\frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu}$  is smaller when  $\alpha = 0$ . Thus, in this case one needs to simply run the SHGD.

To appreciate the generality of Theorem 5.1, let us present some corollaries and compare the rates with existing results in the literature. First, we get a rate for the deterministic CO algorithm.



**Corollary 5.2** (Deterministic CO). Let all assumptions of Theorem 5.1 be satisfied. Then, for  $\gamma_k = \gamma \leq \frac{1}{4L_{\mathcal{H}}}$  and  $\alpha_k = \alpha \leq \frac{1}{4\ell}$ , the iterates of CO satisfy:  $\|x^k - x^*\|^2 \leq (1 - \gamma\mu_{\mathcal{H}} - \alpha\mu)^k \|x^0 - x^*\|^2$ .

The result of Corollary 5.2 should be compared to the convergence guarantees for CO as provided in Abernethy et al. [2021], where the authors viewed CO as a perturbation of the deterministic Hamiltonian gradient descent (HGD). In particular, Abernethy et al. [2021] gives the rate  $1 - \frac{\mu_{\mathcal{H}}}{4L_{\mathcal{H}}}$  for CO under similar assumptions,<sup>7</sup> which is clearly slower than our  $1 - \frac{\mu_{\mathcal{H}}}{4L_{\mathcal{H}}} - \frac{\mu}{4\ell}$  rate. The authors had explicitly noted that treating the GDA part as an adversarial perturbation most likely should be improved upon. With our analysis, we provide a different, more natural analysis of CO and answer this open problem. In addition, note that by setting  $\gamma_k = 0$  and  $\alpha_k = \alpha < 1/2\ell_{\xi}$ , then SCO becomes equivalent to SGDA and Theorem 5.1 matches the convergence guarantees presented in Theorem 4.1. On the other hand, by setting  $\alpha^k = 0$ , SCO yields equivalent updates to SHGD and our result matches the theoretical guarantees of Loizou et al. [2020] as we show in the next corollary.

**Corollary 5.3.** Under the assumptions of Thm. 5.1, set  $\alpha^k = 0$  and  $\gamma_k = \gamma \leq 1/2\mathcal{L}_{\mathcal{H}}$ . Then SCO is equivalent to SHGD and its iterates satisfy:  $\mathbb{E}[\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}})^k \|x^0 - x^*\|^2 + \frac{2\gamma\sigma_{\mathcal{H}}^2}{\mu_{\mathcal{H}}}$ .

All previous results for SCO show convergence to a neighborhood of  $x^*$ . In the next theorem, by selecting decreasing step-sizes (switching strategy) for the values of  $\alpha$  and  $\gamma$ , we are able to guarantee a sublinear convergence to the exact solution for SCO. To the best of our knowledge this is the first result analyzing SCO with decreasing step-sizes.

**Theorem 5.4.** Assume  $\xi$  is  $\mu$ -quasi-strongly monotone and that  $\xi \in EC(\ell_{\xi})$ . Assume that the Hamiltonian function  $\mathcal{H}$  is  $\mu_{\mathcal{H}}$ -quasi strongly convex and  $\mathcal{L}_{\mathcal{H}}$ -expected smooth. Let  $\alpha_k = \gamma_k$ ,  $\psi = \max\{\ell_{\xi}, \mathcal{L}_{\mathcal{H}}\}$  and  $k^* := 8\frac{\psi}{\mu_{\mathcal{H}} + \mu}$ . Let also,  $\gamma_k = \frac{1}{4\psi}$  for  $k \leq \lceil k^* \rceil$  and  $\gamma_k = \frac{2k+1}{(k+1)^2[\mu_{\mathcal{H}} + \mu]}$  for  $k > \lceil k^* \rceil$ . If  $k \geq \lceil k^* \rceil$ , then SCO iterates satisfy:

$$\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma_{\mathcal{H}}^2 + \sigma^2}{[\mu + \mu_{\mathcal{H}}]^2} \frac{16}{k} + \frac{(k^*)^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right) \quad (12)$$

If  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then SCO is still able to converge sublinearly with  $O\left(\frac{1}{k}\right)$ , to  $x^*$ .

## 6 Numerical Evaluation

The purpose of this experimental section is to corroborate our theoretical results, which form the main contributions of this paper. To do so, we focus on strongly-monotone quadratic games of the following form:

$$\min_{x_1 \in \mathbb{R}^d} \max_{x_2 \in \mathbb{R}^p} \frac{1}{n} \sum_i \frac{1}{2} x_1^{\top} \mathbf{A}_i x_1 + x_1^{\top} \mathbf{B}_i x_2 - \frac{1}{2} x_2^{\top} \mathbf{C}_i x_2 + a_i^{\top} x_1 - c_i^{\top} x_2 \quad (13)$$

We show that the Hamiltonian function of such a game is  $\mu_{\mathcal{H}}$ -quasi-strongly convex and  $L_{\mathcal{H}}$ -smooth in App. C.1. For the game to also be strongly-monotone and co-coercive, we sample the matrices such that  $\mu_A \mathbf{I} \preceq \mathbf{A}_i \preceq L_A \mathbf{I}$ ,  $\mu_C \mathbf{I} \preceq \mathbf{C}_i \preceq L_C \mathbf{I}$ , and  $\mu_B^2 \mathbf{I} \preceq \mathbf{B}_i^{\top} \mathbf{B}_i \preceq L_B^2 \mathbf{I}$ , where  $\mathbf{I}$  is the identity matrix; the exact sampling is described in App. C.2. The bias terms  $a_i, c_i$  are sampled from a normal distribution. We pick the step-size for the different methods according to our theoretical findings. That is, for constant step-size, we select  $\alpha = \frac{1}{2\ell_{\xi}}$  for SGDA (Theorem 4.1),  $\alpha = \frac{1}{4\ell_{\xi}}, \gamma = \frac{1}{4\mathcal{L}_{\mathcal{H}}}$  for SCO (Theorem 5.1), and  $\gamma = \frac{1}{2\mathcal{L}_{\mathcal{H}}}$  for SHGD (Corollary 5.3). For the stepsize-switching rule that guarantees convergence to  $x^*$ , we use the step-sizes proposed in Theorem 4.3 for SGDA and Theorem 5.4 for SCO. Let us first look at the qualitative behavior of SGDA, SCO and SHGD on a simple 2d example where  $x_1, x_2 \in \mathbb{R}$ . We show the trajectories of the different algorithms in Fig. 1a. As expected we observe that the behavior of SCO is in between SGDA and SHGD. Recall that the update rule of SCO is a weighted combination of SGDA and the SHGD. The code to reproduce our results can be found at <https://github.com/hugobb/StochasticGamesOpt>.

**Comparison of Algorithms (Constant step-size).** We look at the convergence of the methods for different games where we vary the condition number  $\kappa_G = \frac{\ell_{\xi}}{\mu}$ . In comparing the methods, we

<sup>7</sup>For this convergence rate, Abernethy et al. [2021] assumed that the Hamiltonian function satisfies the Polyak-Lojasiewicz condition but focus on strongly-convex and strongly-concave min-max optimization problems.




Figure 1: Convergence of SGDA, SCO and SHGD on different quadratic games. **(a)** Trajectories of SCO, SHGD and SGDA on a 2d quadratic game. The arrows represent the direction defined by  $\xi(x)$  at a particular point  $x$ . **(b-d)** Distance to optimality  $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$  as a function of the number of iterations. Each plot corresponds to a game with a particular condition number  $\kappa_G = \frac{\ell\xi}{\mu}$ . The solid lines represent the average performance over the 5 runs and the colored area represent the 95% confidence intervals.

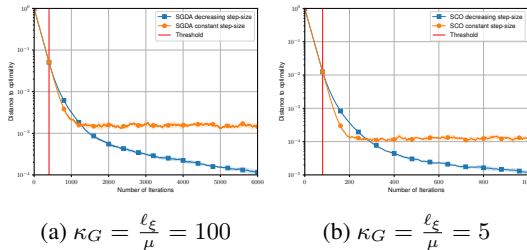


Figure 2: Comparison between constant and decreasing step size regimes of SGDA and SCO. The vertical red lines correspond to the moment we switch from a constant to a decreasing step-size. The solid lines represent the average performance over the 5 runs and the colored area represent the 95% confidence intervals.

use the relative distance to optimality  $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$ . As predicted from our theoretical results, when constant step-size is used, all the methods converge linearly to a neighborhood of the solution (see Figure 1). We observe that the performance of SGDA depends on the condition number: the higher the condition number, the slower the convergence. In contrast, the convergence of both SHGD and SCO is less affected by a larger condition number. We also observe that SGDA is slower than SHGD, but converges to a smaller neighborhood of the solution (see e.g. Fig. 1b). SCO achieves a good trade-off; it converges fast like SHGD and to a small neighborhood like SGDA. An important note is that the size of the neighborhood heavily depends on the selection of the learning rate; we explore this dependence in App. C.3.

**Constant vs Decreasing step-size.** We also compare the performance of SGDA and SCO in the constant and decreasing step-size regimes considered in Theorems 4.1 and 4.3 for SGDA and Theorems 5.1 and 5.4 for SCO. We present our results in Figure 2. As predicted from our theoretical analysis for both methods, the decreasing step-size (switching step-size rule) reaches higher precision compare to the constant step-size. In Figure 2a the vertical red line denotes the value  $4\lceil \ell\xi / \mu \rceil$  predicted in Theorem 4.3 while in Figure 2b the red line denotes the value  $\lceil 8 \frac{\max\{\ell\xi, \mathcal{L}\mathcal{H}\}}{\mu\mathcal{H} + \mu} \rceil$  predicted in Theorem 5.4. Note that for both algorithms the red line is a good approximation of the point where SGDA and SCO need to change their update rules from constant to decreasing step-size.

## 7 Conclusion and Future Directions of Research

We provided the first last-iterate convergence analysis of SGDA and SCO without requiring any strong bounded noise assumption, by introducing the much weaker expected co-coercivity assumption. We proved last-iterate convergence for both methods for a class of unconstrained variational inequality problems that are potentially non-monotone (quasi-strongly monotone problems), with both constant and decreasing step-sizes. Future work includes extending our results beyond the  $\mu$ -quasi strongly monotone assumption for SGDA (assuming only co-coercivity), where we expect to obtain a slower sublinear rate. We also believe the proposal of expected co-coercivity to be of independent interest; it could be used to provide an efficient analysis of other algorithms to solve (1) under the arbitrary sampling paradigm, and it would also be interesting to generalize the analysis to the constrained formulation of variational inequalities.

## Acknowledgments and Disclosure of Funding

**Funding.** Nicolas Loizou acknowledges support by the IVADO Postdoctoral Funding Program. This research was partially supported by the Canada CIFAR AI Chair Program and by a Google Focused Research award. Simon Lacoste-Julien is a CIFAR Associate Fellow in the Learning in Machines & Brains program.

**Competing interests.** Simon Lacoste-Julien additionally works part time as the head of the SAIT AI Lab, Montreal from Samsung. Hugo Berard was working part time as a research intern at Facebook, Montreal.

## References

- J. Abernethy, K. A. Lai, and A. Wibisono. Last-iterate convergence rates for min-max optimization: Convergence of Hamiltonian gradient descent and consensus optimization. In *ALT*, 2021.
- M. Assran, N. Loizou, N. Ballas, and M. Rabbat. Stochastic gradient push for distributed deep learning. *ICML*, 2019.
- W. Azizian, I. Mitliagkas, S. Lacoste-Julien, and G. Gidel. A tight and unified analysis of gradient-based methods for a whole spectrum of differentiable games. In *AISTATS*, 2020.
- P. Balamurugan and F. Bach. Stochastic variance reduction methods for saddle-point problems. In *NeurIPS*, 2016.
- D. Balduzzi, S. Racaniere, J. Martens, J. Foerster, K. Tuyls, and T. Graepel. The mechanics of n-player differentiable games. In *ICML*, 2018.
- H. H. Bauschke, P. L. Combettes, et al. *Convex analysis and monotone operator theory in Hilbert spaces*, volume 408. Springer, 2011.
- L. Bottou, F. E. Curtis, and J. Nocedal. Optimization methods for large-scale machine learning. *SIAM Review*, 60(2):223–311, 2018.
- L. Brighi and R. John. Characterizations of pseudomonotone maps and economic equilibrium. *Journal of Statistics and Management Systems*, 5(1-3):253–273, 2002.
- T. Chavdarova, G. Gidel, F. Fleuret, and S. Lacoste-Julien. Reducing noise in GAN training with variance reduced extragradient. In *NeurIPS*, 2019.
- G. H. Chen and R. T. Rockafellar. Convergence rates in forward–backward splitting. *SIAM Journal on Optimization*, 7(2):421–444, 1997.
- S. C. Choi, W. S. DeSarbo, and P. T. Harker. Product positioning under price competition. *Management Science*, 36(2):175–199, 1990.
- C. Daskalakis, A. Ilyas, V. Syrgkanis, and H. Zeng. Training GANs with optimism. In *ICLR*, 2018.
- C. Daskalakis, S. Skoulakis, and M. Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1466–1478, 2021.
- D. Davis and W. Yin. A three-operator splitting scheme and its optimization applications. *Set-valued and variational analysis*, 25(4):829–858, 2017.
- V. F. Dem’yanov and A. B. Pevnyi. Numerical methods for finding saddle points. *USSR Computational Mathematics and Mathematical Physics*, 12(5):11–52, 1972.
- J. Diakonikolas, C. Daskalakis, and M. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *AISTATS*, 2021.
- A. M. Elizarov and A. Kalimullina. Maximization of the lift/drag ratio of airfoils with a turbulent boundary layer: Sharp estimates, approximation, and numerical solutions. *Computational Mathematics and Mathematical Physics*, 49(3):559–572, 2009.
- F. Facchinei and C. Kanzow. Generalized Nash equilibrium problems. *4OR*, 5(3):173–210, 2007.
- F. Facchinei and J.-S. Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer Science & Business Media, 2007.

- M. Friedlander and M. Schmidt. Hybrid deterministic-stochastic methods for data fitting. *SIAM Journal on Scientific Computing*, 34(3):A1380–A1405, 2012.
- G. Gidel, H. Berard, G. Vignoud, P. Vincent, and S. Lacoste-Julien. A variational inequality perspective on generative adversarial networks. In *ICLR*, 2018.
- N. Golowich, S. Pattathil, and C. Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. In *NeurIPS*, 2020a.
- N. Golowich, S. Pattathil, C. Daskalakis, and A. Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *COLT*, 2020b.
- I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial nets. In *NeurIPS*, 2014.
- R. Gower, O. Sebbouh, and N. Loizou. SGD for structured nonconvex functions: Learning rates, minibatching and interpolation. In *AISTATS*, 2021.
- R. M. Gower, N. Loizou, X. Qian, A. Sailanbayev, E. Shulgin, and P. Richtárik. SGD: General analysis and improved rates. In *ICML*, 2019.
- P. T. Harker and J.-S. Pang. Finite-dimensional variational inequality and nonlinear complementarity problems: a survey of theory, algorithms and applications. *Mathematical programming*, 48(1):161–220, 1990.
- E. Hazan and S. Kale. Beyond the regret minimization barrier: optimal algorithms for stochastic strongly-convex optimization. *The Journal of Machine Learning Research*, 15(1):2489–2512, 2014.
- Y.-G. Hsieh, F. Iutzeler, J. Malick, and P. Mertikopoulos. Explore aggressively, update conservatively: Stochastic extragradient methods with variable stepsize scaling. In *NeurIPS*, 2020.
- A. Juditsky, A. Nemirovski, and C. Tauvel. Solving variational inequalities with stochastic mirror-prox algorithm. *Stochastic Systems*, 1(1):17–58, 2011.
- A. Kannan and U. V. Shanbhag. Optimal stochastic extragradient schemes for pseudomonotone stochastic variational inequality problems and their variants. *Computational Optimization and Applications*, 74(3):779–820, 2019.
- H. Karimi, J. Nutini, and M. Schmidt. Linear convergence of gradient and proximal-gradient methods under the Polyak-Łojasiewicz condition. In *ECML-PKDD*, 2016.
- A. Khaled and P. Richtárik. Better theory for SGD in the nonconvex world. *arXiv preprint arXiv:2002.03329*, 2020.
- A. Khaled, O. Sebbouh, N. Loizou, R. M. Gower, and P. Richtárik. Unified analysis of stochastic gradient methods for composite convex and smooth optimization. *arXiv preprint arXiv:2006.11573*, 2020.
- A. Koloskova, N. Loizou, S. Boreiri, M. Jaggi, and S. U. Stich. A unified theory of decentralized SGD with changing topology and local updates. *ICML*, 2020.
- Y. Li and Y. Yuan. Convergence analysis of two-layer neural networks with relu activation. In *NeurIPS*, 2017.
- T. Liang and J. Stokes. Interaction matters: A note on non-asymptotic local convergence of generative adversarial networks. In *AISTATS*, 2019.
- T. Lin, C. Jin, and M. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *ICML*, 2020a.
- T. Lin, Z. Zhou, P. Mertikopoulos, and M. Jordan. Finite-time last-iterate convergence for multi-agent learning in games. In *ICML*, 2020b.
- P.-L. Lions and B. Mercier. Splitting algorithms for the sum of two nonlinear operators. *SIAM Journal on Numerical Analysis*, 16(6):964–979, 1979.
- N. Loizou and P. Richtárik. Convergence analysis of inexact randomized iterative methods. *SIAM Journal on Scientific Computing*, 42(6):A3979–A4016, 2020a.
- N. Loizou and P. Richtárik. Momentum and stochastic momentum for stochastic gradient, newton, proximal point and subspace descent methods. *Computational Optimization and Applications*, 77(3):653–710, 2020b.

- N. Loizou, H. Berard, A. Jolicoeur-Martineau, P. Vincent, S. Lacoste-Julien, and I. Mitliagkas. Stochastic Hamiltonian gradient methods for smooth games. In *ICML*, 2020.
- N. Loizou, S. Vaswani, I. Laradji, and S. Lacoste-Julien. Stochastic polyak step-size for SGD: An adaptive learning rate for fast convergence. *AISTATS*, 2021.
- L. Luo, H. Ye, Z. Huang, and T. Zhang. Stochastic recursive gradient descent ascent for stochastic nonconvex-strongly-concave minimax problems. *NeurIPS*, 2020.
- P. Mertikopoulos and Z. Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1):465–507, 2019.
- L. Mescheder, S. Nowozin, and A. Geiger. The numerics of GAN. In *NeurIPS*, 2017.
- K. Mishchenko, D. Kovalev, E. Shulgin, P. Richtárik, and Y. Malitsky. Revisiting stochastic extragradient. In *AISTATS*, 2020.
- A. Mokhtari, A. Ozdaglar, and S. Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *AISTATS*, 2020.
- I. Necoara, Y. Nesterov, and F. Glineur. Linear convergence of first order methods for non-strongly convex optimization. *Math. Program.*, pages 1–39, 2018.
- A. Nemirovski, A. Juditsky, G. Lan, and A. Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on Optimization*, 19(4):1574–1609, 2009.
- Y. Nesterov. *Introductory Lectures on Convex Optimization: A Basic Course*, volume 87. Springer Science & Business Media, 2013.
- L. Nguyen, P. H. Nguyen, M. van Dijk, P. Richtárik, K. Scheinberg, and M. Takáč. SGD and Hogwild! Convergence without the bounded gradients assumption. In *ICML*, 2018.
- B. Palaniappan and F. Bach. Stochastic variance reduction methods for saddle-point problems. In *NeurIPS*, 2016.
- V. Patel and S. Zhang. Stochastic gradient descent on nonconvex functions with general noise models. *arXiv preprint arXiv:2104.00423*, 2021.
- B. A. Pearlmutter. Fast exact multiplication by the Hessian. *Neural computation*, 6(1):147–160, 1994.
- D. Pfau and O. Vinyals. Connecting generative adversarial networks and actor-critic methods. *arXiv preprint arXiv:1610.01945*, 2016.
- A. Rakhlin, O. Shamir, and K. Sridharan. Making gradient descent optimal for strongly convex stochastic optimization. In *ICML*, 2012.
- B. Recht, C. Re, S. Wright, and F. Niu. Hogwild: A lock-free approach to parallelizing stochastic gradient descent. In *NeurIPS*, 2011.
- P. Richtárik and M. Takáč. On optimal probabilities in stochastic coordinate descent methods. *Optimization Letters*, 10(6):1233–1243, 2016.
- P. Richtárik and M. Takáč. Stochastic reformulations of linear systems: algorithms and convergence theory. *SIAM Journal on Matrix Analysis and Applications*, 41(2):487–524, 2020.
- L. Rosasco, S. Villa, and B. C. Vũ. A stochastic forward-backward splitting method for solving monotone inclusions in hilbert spaces. *arXiv preprint arXiv:1403.7999*, 2014.
- A. Rousseau, P. Sharer, S. Pagerit, and S. Das. Trade-off between fuel economy and cost for advanced vehicle configurations. In *20th International Electric Vehicle Symposium (EVS20), Monaco*, volume 5, 2005.
- M. Schmidt, N. Le Roux, and F. Bach. Minimizing finite sums with the stochastic average gradient. *Math. Program.*, 162(1-2):83–112, 2017.
- G. Scutari, D. P. Palomar, F. Facchinei, and J.-S. Pang. Convex optimization, game theory, and variational inequality theory. *IEEE Signal Processing Magazine*, 27(3):35–49, 2010.
- C. Song, Z. Zhou, Y. Zhou, Y. Jiang, and Y. Ma. Optimistic dual extrapolation for coherent non-monotone variational inequalities. *NeurIPS*, 2020.

- Q. Tran Dinh, D. Liu, and L. Nguyen. Hybrid variance-reduced SGD algorithms for minimax problems with nonconvex-linear function. *NeurIPS*, 2020.
- P. Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.
- S. Vaswani, F. Bach, and M. Schmidt. Fast and faster convergence of SGD for over-parameterized models and an accelerated perceptron. *arXiv preprint arXiv:1810.07288*, 2018.
- B. C. Vũ. A splitting algorithm for dual monotone inclusions involving cocoercive operators. *Advances in Computational Mathematics*, 38(3):667–681, 2013.
- J. Wen, C.-N. Yu, and R. Greiner. Robust learning under uncertain test distributions: Relating covariate shift to model misspecification. In *ICML*, 2014.
- J. Yang, N. Kiyavash, and N. He. Global convergence and variance-reduced optimization for a class of nonconvex-nonconcave minimax problems. *NeurIPS*, 2020.
- Z. Zhou, P. Mertikopoulos, N. Bambos, S. Boyd, and P. W. Glynn. Stochastic mirror descent in variationally coherent optimization problems. *NeurIPS*, 2017.
- Z. Zhou, P. Mertikopoulos, A. L. Moustakas, N. Bambos, and P. Glynn. Robust power management via learning and game design. *Operations Research*, 69(1):331–345, 2021.
- D. L. Zhu and P. Marcotte. Co-coercivity and its role in the convergence of iterative schemes for solving variational inequalities. *SIAM Journal on Optimization*, 6(3):714–726, 1996.

## Checklist

1. For all authors...
  - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
  - (b) Did you describe the limitations of your work? [Yes]
  - (c) Did you discuss any potential negative societal impacts of your work? [N/A] Our main results are of theoretical nature (proposal of new convergence analysis). Therefore, we do not see any immediate societal impact of our results.
  - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
  - (a) Did you state the full set of assumptions of all theoretical results? [Yes]
  - (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] The code is available at <https://github.com/hugobb/StochasticGames0pt>.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
  - (a) If your work uses existing assets, did you cite the creators? [N/A]
  - (b) Did you mention the license of the assets? [N/A]
  - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
  - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
  - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

# Supplementary Material

The supplementary material is organized as follows: In Section [A](#), we give some basic definitions and provide the proofs of the propositions, lemmas and theorems related to the expected co-coercivity condition as presented in Section [3](#) of the main paper. In Section [B](#) we present the proofs of the main theorems and corollaries for the convergence of SGDA and SCO. In Section [C](#) we present the experimental details and provide additional experiments. Finally in Section [D](#) we explain how our convergence results can be easily adapted to the general stochastic setting and in Section [E](#) we provide further related work.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Arbitrary Sampling: Stochastic Reformulation of Problem (1)</b>	<b>3</b>
<b>3</b>	<b>Expected Co-coercivity and Connection to Other Assumptions</b>	<b>4</b>
3.1	Expected Co-coercivity (EC) . . . . .	4
<b>4</b>	<b>Stochastic Gradient Descent-Ascent</b>	<b>6</b>
<b>5</b>	<b>Stochastic Consensus Optimization</b>	<b>7</b>
5.1	Setting . . . . .	7
5.2	Stochastic Consensus Optimization and its Special Cases . . . . .	8
5.3	Convergence Analysis . . . . .	8
<b>6</b>	<b>Numerical Evaluation</b>	<b>9</b>
<b>7</b>	<b>Conclusion and Future Directions of Research</b>	<b>10</b>
<b>A</b>	<b>Proofs of Results on Co-coercivity and Expected Co-coercivity</b>	<b>17</b>
A.1	Proof of Proposition 3.2 . . . . .	17
A.2	Proof of Lemma 3.4 . . . . .	18
A.3	Proof of Proposition 3.5 . . . . .	18
A.4	Proof of Proposition 3.6 . . . . .	21
A.5	Connections of EC to other Assumptions . . . . .	22
A.6	Example: Quasi-strongly Monotone Operator that is not Monotone nor Lipschitz . . . . .	23
A.7	Example: Co-coercivity for Quadratic Games . . . . .	24
<b>B</b>	<b>Proofs of Main Convergence Analysis Results</b>	<b>25</b>
B.1	Proof of Theorem 4.1 . . . . .	25
B.2	Proof of Theorem 4.3 . . . . .	26
B.3	Proof of Theorem 5.1 . . . . .	27
B.4	Proof of Theorem 5.4 . . . . .	30
<b>C</b>	<b>On Experiments</b>	<b>33</b>



C.1	Properties of Hamiltonian Function for Quadratic Games . . . . .	33
C.2	Experimental Details . . . . .	33
C.3	Additional Experiment: Influence of the Step-size on Convergence . . . . .	34
<b>D</b>	<b>Beyond Finite-sum Structure</b>	<b>35</b>
<b>E</b>	<b>More Related Work</b>	<b>35</b>

## A Proofs of Results on Co-coercivity and Expected Co-coercivity

Let us start by re-stating the main definitions of the classes of operators under study.

**Definition A.1** (Lipschitz continuous). An operator  $\xi : \mathbb{R}^d \rightarrow \mathbb{R}^d$  is  $L$ -Lipschitz continuous if there is  $L > 0$  such that:

$$\|\xi(x) - \xi(y)\| \leq L\|x - y\|, \quad \forall x, y \in \mathbb{R}^d \quad (14)$$

**Definition A.2** (Co-coercivity). We say that an operator  $\xi$  is  $\ell$ -co-coercive if there exist  $\ell > 0$  such that:

$$\|\xi(x) - \xi(y)\|^2 \leq \ell \langle \xi(x) - \xi(y), x - y \rangle \quad \forall x, y \in \mathbb{R}^d.$$

**Definition A.3** (Co-coercive around  $w^*$ ). We say that an operator  $\xi$  is  $\ell$ -co-coercive around  $w^*$  if there exist  $w^* \in \mathbb{R}^d$  and  $\ell > 0$  such that

$$\|\xi(x) - \xi(w^*)\|^2 \leq \ell \langle \xi(x) - \xi(w^*), x - w^* \rangle \quad \forall x \in \mathbb{R}^d.$$

Note that in this definition, the point  $w^*$  is not necessarily a point where  $\xi(w^*) = 0$ .

**Definition A.4** (Strongly monotone / monotone). We say that an operator  $\xi$  is  $\mu$ -strongly monotone if there exist  $\mu > 0$  such that

$$\langle \xi(x) - \xi(y), x - y \rangle \geq \mu \|x - y\|^2 \quad \forall x, y \in \mathbb{R}^d.$$

If  $\mu = 0$ , that is

$$\langle \xi(x) - \xi(y), x - y \rangle \geq 0 \quad \forall x, y \in \mathbb{R}^d,$$

then we say that the operator is monotone.

**Definition A.5** (Quasi-Strongly Monotone / Variational Stability Condition). We say that an operator  $\xi$  is  $\mu$ -quasi-strongly monotone if there exist  $\mu > 0$  such that

$$\langle \xi(x), x - x^* \rangle \geq \mu \|x - x^*\|^2 \quad \forall x \in \mathbb{R}^d.$$

Here  $x^*$  is the solution of the stochastic variational inequality problem (1). If  $\mu = 0$ , that is

$$\langle \xi(x), x - x^* \rangle \geq 0 \quad (15)$$

then we say that  $\xi$  satisfies the variational stability condition.

### A.1 Proof of Proposition 3.2

Before stating the proof of Proposition 3.2, we clarify that the assumption of  $L$ -Lipschitzness of  $\xi$  is only used for the implications where  $L$  appear; it is not needed for the other implications. In particular, while a  $\ell$ -co-coercive operator is always  $\ell$ -Lipschitz continuous by using Cauchy-Schwartz,<sup>8</sup> it is possible for an operator to be  $\ell$ -co-coercive around  $x^*$  and *not* be Lipschitz continuous (see such an example in Section A.6). This highlights the wider applicability of the  $\ell$ -co-coercivity around  $x^*$  assumption that is all we need for several of our convergence results, in contrast to the Lipschitz continuity of  $\xi$  which is typically assumed in the variational inequality literature.

<sup>8</sup> $\|\xi(x) - \xi(x')\|^2 \leq \ell \langle \xi(x) - \xi(x'), x - x' \rangle \leq \ell \|\xi(x) - \xi(x')\| \|x - x'\| \implies \|\xi(x) - \xi(x')\| \leq \ell \|x - x'\|.$

*Proof.* Most of these implications can be found in [Facchinei and Pang \[2007\]](#).

$\mu$ -strongly monotone  $\implies \frac{L^2}{\mu}$ -co-coercive: The proof of this result is a direct application of strong monotonicity and Lipschitzness properties:

$$\|\xi(x) - \xi(x')\|^2 \leq L^2 \|x - x'\|^2 \leq \frac{L^2}{\mu} \langle \xi(x) - \xi(x'), x - x' \rangle, \quad \forall x, x' \in \mathbb{R}^d.$$

$\ell$ -co-coercive  $\implies$  monotone: It comes from the fact that a norm is non-negative.

$$0 \leq \|\xi(x) - \xi(x')\|^2 \leq \ell \langle \xi(x) - \xi(x'), x - x' \rangle, \quad \forall x, x' \in \mathbb{R}^d.$$

monotone  $\implies$  variational stability condition: It comes from the fact that monotonicity applied to  $x' = x^*$  is variational stability condition.

Quasi  $\mu$ -strongly monotone  $\implies \frac{L^2}{\mu}$ -co-coercive relatively to  $x^*$ : (It is the only implication that is not proven in [Facchinei and Pang \[2007\]](#)) The proof of this result is a direct application of quasi-strong monotonicity and Lipschitzness properties:

$$\|\xi(x) - \xi(x^*)\|^2 \leq L^2 \|x - x^*\|^2 \leq \frac{L^2}{\mu} \langle \xi(x), x - x^* \rangle, \quad \forall x \in \mathbb{R}^d.$$

$\frac{L^2}{\mu}$ -co-coercive relatively to  $x^*$   $\implies$  variational stability condition: It comes from the fact that a norm is non-negative.

$$0 \leq \|\xi(x) - \xi(x^*)\|^2 \leq \ell \langle \xi(x), x - x^* \rangle, \quad \forall x \in \mathbb{R}^d.$$

□

## A.2 Proof of Lemma 3.4

*Proof.*

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} \|\xi_v(x)\|^2 &= \mathbb{E}_{\mathcal{D}} \|\xi_v(x) - \xi_v(x^*) + \xi_v(x^*)\|^2 \\ &\leq 2\mathbb{E}_{\mathcal{D}} \|\xi_v(x) - \xi_v(x^*)\|^2 + 2\mathbb{E}_{\mathcal{D}} \|\xi_v(x^*)\|^2 \\ &\stackrel{EC}{\leq} 2\ell_{\xi} \langle \xi(x), x - x^* \rangle + 2\mathbb{E}_{\mathcal{D}} \|\xi_v(x^*)\|^2 \\ &\stackrel{(6)}{\leq} 2\ell_{\xi} \langle \xi(x), x - x^* \rangle + 2\sigma^2. \end{aligned} \tag{16}$$

The first inequality follows from the estimate  $\|a + b\|^2 \leq 2\|a\|^2 + 2\|b\|^2$ . □

## A.3 Proof of Proposition 3.5

Before we formally present the proof of Proposition 3.5, let us first establish some random set terminology.

Let  $C \subseteq [n]$  and let  $e_C := \sum_{i \in C} e_i$ , where  $\{e_1, \dots, e_n\}$  are the standard basis vectors in  $\mathbb{R}^n$ . These subsets will be selected using a random set valued map  $S$ , in the literature referred to by the name *sampling*. A sampling is uniquely characterized by choosing subset probabilities  $p_C \geq 0$  for all subsets  $C$  of  $[n]$ :

$$\mathbb{P}[S = C] = p_C, \quad \forall C \subset [n], \tag{17}$$

where  $\sum_{C \subseteq [n]} p_C = 1$ . In this work, following the terminology of [Gower et al. \[2019, 2021\]](#), our results hold for *proper* samplings.

**Definition A.6.** A sampling  $S$  is called proper if  $p_i \stackrel{\text{def}}{=} \mathbb{P}[i \in S] = \sum_{C: i \in C} p_C$  is positive for all  $i$ .

As we mentioned in the main paper, in this work we focus on  $b$ -minibatch sampling (see Definition 2.1) however we highlight again that our results hold for the larger class of sampling vectors  $v \in \mathbb{R}^n$  that satisfy  $\mathbb{E}_{\mathcal{D}}[v_i] = 1$ , for  $i = 1, \dots, n$ .

For example, the random vector  $v = v(S)$  given by  $v = \sum_{i \in S} \frac{1}{p_i} e_i$  is a sampling vector. This can be easily proved, by noticing that  $v_i = \mathbf{1}_{(i \in S)} / p_i$ , where  $\mathbf{1}_{(i \in S)}$  is the indicator function of the event

$i \in S$ . Then, It follows that  $\mathbb{E}[v_i] = \mathbb{E}[\mathbf{1}_{\{i \in S\}}]/p_i = 1$ . Commonly used samplings that captured by our theory are the *independent sampling*, *partition sampling*, *single-element sampling* and *importance sampling*. For more details on these different samplings check [Gower et al. \[2019\]](#).

By definition 2.1 of  $b$ -minibatch sampling, it holds that  $\mathbb{P}[i \in S] = p_i = \frac{b}{n}$ , and  $\mathbb{P}[i, j \in S] = \frac{b}{n} \frac{b-1}{n-1}$ .

All the sampling schemes presented in [Gower et al. \[2019\]](#) had the following additional property: there exists a constant  $z > 0$  such that

$$\frac{\mathbb{P}[i, j \in S]}{\mathbb{P}[i \in S] \mathbb{P}[j \in S]} = z, \quad \forall i, j \in \{1, \dots, n\}, i \neq j. \quad (18)$$

For  $b$ -minibatch sampling,  $z = \frac{n}{b} \frac{b-1}{n-1}$  [[Gower et al., 2021](#)].

Let us now present the proof of Proposition 3.5.

*Proof.* Since  $\xi$  is  $\ell_i$ -co-coercive around  $x^*$  then we have that  $\xi$  is  $\ell$ -co-coercive around  $x^*$ . That is  $\forall x \in \mathbb{R}^d$  it holds,

$$\|\xi_i(x) - \xi_i(x^*)\|^2 \leq \ell_i \langle \xi_i(x) - \xi_i(x^*), x - x^* \rangle \quad (19)$$

$$\|\xi(x) - \xi(x^*)\|^2 \leq \ell \langle \xi(x) - \xi(x^*), x - x^* \rangle = \ell \langle \xi(x), x - x^* \rangle. \quad (20)$$

Noticing that

$$\begin{aligned} \|\xi_v(x) - \xi_v(x^*)\|^2 &= \frac{1}{n^2} \left\| \sum_{i \in S} \frac{1}{p_i} (\xi_i(x) - \xi_i(x^*)) \right\|^2 \\ &= \sum_{i, j \in S} \left\langle \frac{1}{np_i} (\xi_i(x) - \xi_i(x^*)), \frac{1}{np_j} (\xi_j(x) - \xi_j(x^*)) \right\rangle, \end{aligned}$$

we have

$$\begin{aligned} \mathbb{E}[\|\xi_v(x) - \xi_v(x^*)\|^2] &= \sum_C p_C \sum_{i, j \in C} \left\langle \frac{1}{np_i} (\xi_i(x) - \xi_i(x^*)), \frac{1}{np_j} (\xi_j(x) - \xi_j(x^*)) \right\rangle \\ &= \sum_{i, j=1}^n \sum_{C: i, j \in C} p_C \left\langle \frac{1}{np_i} (\xi_i(x) - \xi_i(x^*)), \frac{1}{np_j} (\xi_j(x) - \xi_j(x^*)) \right\rangle \\ &= \sum_{i, j=1}^n \frac{\mathbb{P}[i, j \in S]}{p_i p_j} \left\langle \frac{1}{n} (\xi_i(x) - \xi_i(x^*)), \frac{1}{n} (\xi_j(x) - \xi_j(x^*)) \right\rangle, \end{aligned}$$

where we used a double counting argument in the 2nd equality. Now since  $\mathbb{P}[i, j \in S]/(p_i p_j) = z$  for  $i \neq j$  (18) and  $\mathbb{P}[i, i \in S] = p_i$  we have from the above that

$$\begin{aligned} \mathbb{E}[\|\xi_v(x) - \xi_v(x^*)\|^2] &= \sum_{i \neq j} z \left\langle \frac{1}{n} (\xi_i(x) - \xi_i(x^*)), \frac{1}{n} (\xi_j(x) - \xi_j(x^*)) \right\rangle \\ &\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{1}{p_i} \|\xi_i(x) - \xi_i(x^*)\|^2 \\ &= \sum_{i, j=1}^n z \left\langle \frac{1}{n} (\xi_i(x) - \xi_i(x^*)), \frac{1}{n} (\xi_j(x) - \xi_j(x^*)) \right\rangle \\ &\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{1}{p_i} (1 - p_i z) \|\xi_i(x) - \xi_i(x^*)\|^2 \quad (21) \end{aligned}$$

$$\begin{aligned} &\stackrel{(19)}{\leq} z \|\xi(x) - \xi(x^*)\|^2 \\ &\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{\ell_i}{p_i} (1 - p_i z) \langle \xi_i(x) - \xi_i(x^*), x - x^* \rangle \\ &\stackrel{(20)}{\leq} \left( z\ell + \max_{i=1, \dots, n} \frac{\ell_i}{np_i} (1 - p_i z) \right) \langle \xi(x), x - x^* \rangle. \quad (22) \end{aligned}$$

Comparing the above to the definition of expected co-coercivity (EC) we have that

$$\ell_\xi = z\ell + \max_{i=1,\dots,n} \frac{\ell_i}{np_i} (1 - p_i z). \quad (23)$$

Using that for  $b$ -minibatch sampling it holds  $\mathbb{P}[i \in S] = p_i = \frac{b}{n}$  and  $z = \frac{n}{b} \frac{b-1}{n-1}$  we obtain

$$\ell_\xi = \frac{n}{b} \frac{b-1}{n-1} \ell + \frac{1}{b} \frac{n-b}{n-1} \ell_{\max},$$

where  $\ell_{\max} = \max\{\ell_i\}_{i=1}^n$ .

The specialized expressions of  $\sigma^2$  for the  $b$ -minibatch sampling, can be obtain by following the same steps of Proposition 3.10 of Gower et al. [2019]. Below using our notation, we include this derivation for completeness:

$$\begin{aligned} \sigma^2 = \mathbb{E}_{\mathcal{D}}[\|\xi_v(x^*)\|^2] &= \mathbb{E} \left[ \left\| \frac{1}{n} \sum_{i=1}^n \nabla f_i(x^*) v_i \right\|^2 \right] = \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n \nabla f_i(x^*) v_i \right\|^2 \right] = \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i \in S} \frac{1}{p_i} \xi_i(x^*) \right\|^2 \right] \\ &= \frac{1}{n^2} \mathbb{E} \left[ \left\| \sum_{i=1}^n 1_{i \in S} \frac{1}{p_i} \xi_i(x^*) \right\|^2 \right] = \frac{1}{n^2} \mathbb{E} \left[ \sum_{i=1}^n \sum_{j=1}^n 1_{i \in S} 1_{j \in S} \left\langle \frac{1}{p_i} \xi_i(x^*), \frac{1}{p_j} \xi_j(x^*) \right\rangle \right] \\ &= \frac{1}{n^2} \sum_{i,j} \frac{\mathbb{P}[i, j \in S]}{p_i p_j} \langle \xi_i(x^*), \xi_j(x^*) \rangle. \end{aligned}$$

Recall that for  $b$ -minibatch sampling,  $\mathbb{P}[i \in S] = p_i = \frac{b}{n}$ , and  $\mathbb{P}[i, j \in S] = \frac{b}{n} \frac{b-1}{n-1}$ . Thus,

$$\begin{aligned} \sigma^2 &= \frac{1}{n^2} \sum_{i,j \in [n]} \frac{\mathbb{P}[i, j \in S]}{p_i p_j} \langle \xi_i(x^*), \xi_j(x^*) \rangle \\ &= \frac{1}{n^2} \sum_{i \neq j} \frac{b(b-1)}{n(n-1)} \cdot \frac{n^2}{b^2} \langle \xi_i(x^*), \xi_j(x^*) \rangle + \frac{1}{n^2} \sum_{i \in [n]} \frac{n}{b} \|\xi_i(x^*)\|^2 \\ &= \frac{1}{nb} \left( \sum_{i \neq j} \frac{b-1}{n-1} \langle \xi_i(x^*), \xi_j(x^*) \rangle + \sum_{i \in [n]} \|\xi_i(x^*)\|^2 \right) \\ &= \frac{1}{nb} \left( \sum_{i,j \in [n]} \frac{b-1}{n-1} \langle \xi_i(x^*), \xi_j(x^*) \rangle + \sum_{i \in [n]} \frac{n-b}{n-1} \|\xi_i(x^*)\|^2 \right) \\ &= \frac{1}{nb} \cdot \frac{n-b}{n-1} \sum_{i \in [n]} \|\xi_i(x^*)\|^2 \\ &= \frac{1}{b} \cdot \frac{n-b}{n-1} \sigma_1^2, \end{aligned}$$

where  $\sigma_1^2 := \frac{1}{n} \sum_{i=1}^n \|\xi_i(x^*)\|^2$ . □

#### A.4 Proof of Proposition 3.6

**Proposition A.7.** Let  $\xi$  be  $\mu$ -quasi-strongly monotone and let  $\xi_i$  be  $L_i$ -Lipschitz continuous for all  $i \in [n]$ . Then  $\xi \in EC(\ell_\xi)$ .

For a general proper sampling scheme (Def. A.6), we can provide the following (loose) bound on  $\ell_\xi$ :

$$\ell_\xi \leq \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{E}_{\mathcal{D}} [v_i^2] L_i^2}{\mu}. \quad (24)$$

If, as is the case for standard sampling schemes from Gower et al. [2019], we assume that there exists a  $z > 0$  such that  $z = \frac{\mathbb{P}[i, j \in S]}{\mathbb{P}[i \in S] \mathbb{P}[j \in S]}$  for all  $i, j \in \{1, \dots, n\}$ ,  $i \neq j$ , then we can use the tighter value:

$$\ell_\xi = \left( zL^2 + \sum_{i=1}^n \frac{1}{n^2} \frac{L_i^2}{p_i} (1 - p_i z) \right) \frac{1}{\mu}, \quad (25)$$

where  $L$  is the Lipschitz continuous parameter of operator  $\xi$  and  $p_i = \mathbb{P}[i \in S]$ .

Finally, for  $b$ -minibatch sampling, it holds  $p_i = \frac{b}{n}$  and  $z = \frac{n}{b} \frac{b-1}{n-1}$  and thus  $\xi \in EC(\ell_\xi)$  with

$$\ell_\xi = \left( \frac{n}{b} \frac{b-1}{n-1} L^2 + \left( \frac{1}{n} \sum_{i=1}^n L_i^2 \right) \frac{1}{b} \frac{n-b}{n-1} \right) \frac{1}{\mu}. \quad (26)$$

*Proof.* Since  $\xi_i$  is  $L_i$ -Lipschitz continuous for all  $i$ , then we have that  $\xi$  is  $L$ -Lipschitz continuous (with  $L \leq \frac{1}{n} \sum_i L_i$  by using Jensen's inequality on  $\|\cdot\|^2$ ). That is,  $\forall x, y \in \mathbb{R}^d$  it holds:

$$\|\xi_i(x) - \xi_i(y)\| \leq L_i \|x - y\|, \quad \forall x, y \in \mathbb{R}^d \quad (27)$$

$$\|\xi(x) - \xi(y)\| \leq L \|x - y\|, \quad \forall x, y \in \mathbb{R}^d \quad (28)$$

We first prove the general case:

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} [\|\xi_v(x) - \xi_v(x^*)\|^2] &= \mathbb{E}_{\mathcal{D}} \left[ \left\| \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x) - \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x^*) \right\|^2 \right] \\ &\stackrel{\text{Jensen's}}{\leq} \mathbb{E}_{\mathcal{D}} \left[ \frac{1}{n} \sum_{i=1}^n \|v_i [\xi_i(x) - \xi_i(x^*)]\|^2 \right] \\ &= \mathbb{E}_{\mathcal{D}} \left[ \frac{1}{n} \sum_{i=1}^n v_i^2 \|\xi_i(x) - \xi_i(x^*)\|^2 \right] \\ &\stackrel{(27)}{\leq} \mathbb{E}_{\mathcal{D}} \left[ \frac{1}{n} \sum_{i=1}^n v_i^2 L_i^2 \|x - x^*\|^2 \right] \\ &= \frac{1}{n} \sum_{i=1}^n \mathbb{E}_{\mathcal{D}} [v_i^2] L_i^2 \|x - x^*\|^2 \\ &\stackrel{(2)}{\leq} \frac{1}{n} \sum_{i=1}^n \frac{\mathbb{E}_{\mathcal{D}} [v_i^2] L_i^2}{\mu} \langle \xi(x), x - x^* \rangle, \end{aligned}$$

yielding (24).

The use of Jensen's inequality above is the source of looseness in the bound. With the  $z$  constant property, we can avoid it with the following derivations

By following the same steps to the proof of Proposition 3.5, we obtain (21). That is,

$$\begin{aligned} \mathbb{E} [\|\xi_v(x) - \xi_v(x^*)\|^2] &= \sum_{i, j=1}^n z \left\langle \frac{1}{n} (\xi_i(x) - \xi_i(x^*)), \frac{1}{n} (\xi_j(x) - \xi_j(x^*)) \right\rangle \\ &\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{1}{p_i} (1 - p_i z) \|\xi_i(x) - \xi_i(x^*)\|^2. \end{aligned}$$

Now by using (27) and (28), we obtain the following

$$\begin{aligned}
\mathbb{E}[\|\xi_v(x) - \xi_v(x^*)\|^2] &= \sum_{i,j=1}^n z \left\langle \frac{1}{n}(\xi_i(x) - \xi_i(x^*)), \frac{1}{n}(\xi_j(x) - \xi_j(x^*)) \right\rangle \\
&\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{1}{p_i} (1 - p_i z) \|\xi_i(x) - \xi_i(x^*)\|^2 \\
&\stackrel{(27)}{\leq} z \|\xi(x) - \xi(x^*)\|^2 \\
&\quad + \sum_{i=1}^n \frac{1}{n^2} \frac{L_i^2}{p_i} (1 - p_i z) \|x - x^*\|^2 \\
&\stackrel{(28)}{\leq} \left( zL^2 + \sum_{i=1}^n \frac{1}{n^2} \frac{L_i^2}{p_i} (1 - p_i z) \right) \|x - x^*\|^2. \tag{29}
\end{aligned}$$

Since  $\xi$  is  $\mu$ -quasi strongly monotone, then (29) becomes:

$$\mathbb{E}[\|\xi_v(x) - \xi_v(x^*)\|^2] \leq \left( zL^2 + \sum_{i=1}^n \frac{1}{n^2} \frac{L_i^2}{p_i} (1 - p_i z) \right) \frac{1}{\mu} \langle \xi(x), x - x^* \rangle \tag{30}$$

Thus the expected co-coercivity (EC) is satisfied with

$$\ell_\xi = \left( zL^2 + \sum_{i=1}^n \frac{1}{n^2} \frac{L_i^2}{p_i} (1 - p_i z) \right) \frac{1}{\mu}. \tag{31}$$

Similar to the proof of Proposition 3.5, using that for  $b$ -minibatch sampling,  $\mathbb{P}[i \in S] = p_i = \frac{b}{n}$  and  $z = \frac{n}{b} \frac{b-1}{n-1}$ , we obtain:

$$\ell_\xi = \left( \frac{n}{b} \frac{b-1}{n-1} L^2 + \left( \frac{1}{n} \sum_{i=1}^n L_i^2 \right) \frac{1}{b} \frac{n-b}{n-1} \right) \frac{1}{\mu}. \tag{32}$$

This completes the proof.  $\square$

## A.5 Connections of EC to other Assumptions

In this section, we present some propositions not included in the main paper showing properties and connections between classical assumptions and our proposed expected co-coercivity EC.

**Proposition A.8.** In the unconstrained setting, if  $\xi$  is  $\mu$ -quasi strongly monotone, then it is not possible to satisfy the bounded operator assumption that there exists a finite  $c$  such that  $\mathbb{E}\|\xi_v(x)\|^2 \leq c$  for every  $x$  in  $\mathbb{R}^d$ .

*Proof.* Let us assume that  $\mathbb{E}\|\xi_v(x)\|^2 \leq c, \forall x \in \mathbb{R}^d$ . Note also that if an operator satisfies the  $\mu$ -quasi strongly monotone property (2), then by using the Cauchy–Schwarz inequality, it satisfies the error bound condition

$$\|\xi(x)\| \geq \mu \|x - x^*\|.$$

By combining the above two inequalities, it holds that:

$$\mu^2 \|x - x^*\|^2 \leq \|\xi(x)\|^2 = \|\mathbb{E}[\xi_v(x)]\|^2 \leq \mathbb{E}[\|\xi_v(x)\|^2] \leq c$$

which means that:

$$\|x - x^*\|^2 \leq \frac{c}{\mu^2}.$$

However, for the *unconstrained* stochastic variational inequality problems (1), a point  $x$  can be very far from the optimum point  $x^*$  and as a result  $\|x - x^*\|^2 \geq \frac{c}{\mu^2}$ . This leads to a contradiction.  $\square$

**Proposition A.9.** In the single-objective optimization when the stochastic problem

$$\min_x \left[ f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x) \right]$$

has convex and smooth functions  $f_i$ , then expected smoothness and expected co-coercivity are equivalent (see last row of Table 1).

*Proof.* For simplicity of exposition, let us focus on single-element sampling.

According to Theorem 2.1.5 in [Nesterov \[2013\]](#), if  $f_i$  is convex and  $L_i$  smooth, then the following two conditions are equivalent:

$$\|\nabla f_i(x) - \nabla f_i(y)\|^2 \leq 2L_i(f_i(x) - f_i(y) - \langle \nabla f_i(y), x - y \rangle). \quad (33)$$

$$\|\nabla f_i(x) - \nabla f_i(y)\|^2 \leq L_i \langle \nabla f_i(x) - \nabla f_i(y), x - y \rangle \quad (34)$$

If in the above two condition we select  $y = x^*$  and take the expectations with respect to  $i$ , then we obtain the following two equivalent conditions:

$$\mathbb{E} [\|\nabla f_i(x) - \nabla f_i(x^*)\|^2] \leq 2L_{\max} [f(x) - f(x^*)]. \quad (35)$$

$$\mathbb{E} [\|\nabla f_i(x) - \nabla f_i(x^*)\|^2] \leq L_{\max} \langle \nabla f(x), x - x^* \rangle \quad (36)$$

Note that in the above, (35) is the expected smoothness as proposed in [Gower et al. \[2019\]](#) while (36) is our expected co-coercivity (EC) for the single element sampling. Note that for single-objective optimization problems, the operator  $\xi$  is simply the gradient vector. As we mentioned in Section 3, in single-objective optimization, a function is  $L$ -co-coercive if and only if it is convex and  $L$ -smooth (i.e.  $L$ -Lipschitz gradients) [[Bauschke et al., 2011](#)]. Thus, the co-coercivity constant is equivalent to the smoothness parameter.  $\square$

Let us also add a simple remark highlighting the weakness of EC compare to other previously used assumptions in the literature of stochastic algorithms for solving (1).

**Remark A.10.** As we show in Lemma (3.4), by assuming EC we obtain the following bound

$$\mathbb{E} \|\xi_v(x)\|^2 \leq 2\ell_\xi \langle \xi(x), x - x^* \rangle + 2\sigma^2. \quad (37)$$

Let us now compare this bound to the assumption of growth condition  $\mathbb{E} \|\xi_i(x)\|^2 \leq c_1 \|\xi(x)\|^2 + c_2$  (weakest among the other assumptions).

Note that if an operator  $\xi$  is  $\ell$ -co-coercive, then the growth condition implies:

$$\mathbb{E} \|\xi_i(x)\|^2 \leq c_1 \ell \langle \xi(x), x - x^* \rangle + c_2.$$

This has the same form to the bound (37), obtained by EC. However, the parameters  $c_1$  and  $c_2$  have unknown values while using EC these parameters are closed-form problem-dependent expressions.

Thus, expected co-coercivity is weaker than the growth condition and at the same time more powerful, as in many case it is not really an assumption, but a condition that is satisfied for free. See for example, Propositions 3.5 and 3.6.

## A.6 Example: Quasi-strongly Monotone Operator that is not Monotone nor Lipschitz

An operator that is  $\mu$ -quasi strongly monotone may not even be monotone. We now give a simple example of such an operator, with the additional property that it is  $L$ -co-coercive around  $x^*$  but is *not* Lipschitz continuous. This highlights the generality of our convergence results beyond the standard monotone setting. Let  $L > \mu > 0$ , we define  $\xi(x) = x(\frac{L-\mu}{2} \cos(\|x\|_2) + \frac{L+\mu}{2})$ . We have that  $x^* = 0$  and

$$\langle \xi(x), x - x^* \rangle = \|x - x^*\|_2^2 \left( \frac{L-\mu}{2} \cos(\|x\|_2) + \frac{L+\mu}{2} \right) \geq \mu \|x - x^*\|_2^2, \quad (38)$$

and is thus  $\mu$ -quasi strongly monotone. However, it is *not* monotone. To see this, let us consider the one dimensional case  $x \in \mathbb{R}$ . In this case, we get,  $\xi(x) = x(\frac{L-\mu}{2} \cos(|x|) + \frac{L+\mu}{2})$ . To formally violate the monotonicity inequality, we can for instance consider  $x = 2\pi k + \frac{\pi}{2}$  and  $x' = 2\pi k$  to get

$$\langle \xi(x) - \xi(x'), x - x' \rangle = (\frac{\pi}{2} \frac{L+\mu}{2} - 2\pi k \frac{L-\mu}{2}) \frac{\pi}{2} = \frac{\pi^2}{4} (L + \mu - 4k(L - \mu)). \quad (39)$$

This quantity is negative for  $k > \frac{L+\mu}{4(L-\mu)}$ .

This operator is also  $L$ -co-coercive with respect to  $x^*$  since,

$$\langle \xi(x), x - x^* \rangle = \|\xi(x)\|_2^2 (\frac{L-\mu}{2} \cos(\|x\|_2) + \frac{L+\mu}{2})^{-1} \geq L^{-1} \|\xi(x)\|_2^2. \quad (40)$$

This operator is *not* Lipschitz continuous for all  $x \in \mathbb{R}$ , as its derivative is unbounded over  $\mathbb{R}$ .

### A.7 Example: Co-coercivity for Quadratic Games

For quadratic games, it is relatively easy to characterize co-coercivity.

**Proposition A.11.** If  $\xi(x) = Ax$  where  $A \in \mathbb{R}^{d \times d}$ , we have that  $\xi$  is co-coercive if and only if  $\langle x, Ax \rangle > 0, \forall x \in \mathbb{R}^d \setminus \text{null}(A)$ . In that case, we have  $\ell = \sup_{\|x\|=1} \frac{\|Ax\|_2^2}{\langle x, Ax \rangle}$ .

*Proof.* When  $\xi(x) = Ax$ , the co-coercivity condition (Definition 3.1)

$$\|A(x - x')\|^2 \leq \ell \langle A(x - x'), x - x' \rangle, \quad \forall x, x' \in \mathbb{R}^d. \quad (41)$$

Now, we can note that the variable of interest is  $x - x' \in \mathbb{R}^d$ . Thus we get equivalently,

$$\|Ax\|^2 \leq \ell \langle Ax, x \rangle, \quad \forall x \in \mathbb{R}^d. \quad (42)$$

This inequality is valid for any  $x \in \mathbb{R}^d$  such that  $Ax = 0$ . Now, let us consider  $x \in \mathbb{R}^d$  such that  $Ax \neq 0$ .

If there exist  $x \in \mathbb{R}^d$  such that  $Ax \neq 0$  and  $\langle x, Ax \rangle \leq 0$  then, we have that

$$0 < \|Ax\|^2 \leq \ell \langle Ax, x \rangle \leq 0 \quad (43)$$

which is not valid. Thus  $\xi$  is *not* co-coercive.

On the other hand, if for all  $x \in \mathbb{R}^d$  such that  $Ax \neq 0$ , we have  $\langle x, Ax \rangle > 0$  then, the co-coercivity condition would demand

$$\frac{\|Ax\|^2}{\langle Ax, x \rangle} \leq \ell, \quad \forall x \in \mathbb{R}^d, Ax \neq 0. \quad (44)$$

Finally, since the left-hand side of the previous equation is scale invariant ( $x \mapsto \lambda x$  does not change the LHS), we have

$$\sup_{x \in \mathbb{R}^d} \frac{\|Ax\|^2}{\langle Ax, x \rangle} = \sup_{x \in \mathbb{R}^d \setminus \{0\}} \frac{\|Ax\|^2}{\langle Ax, x \rangle} = \sup_{\|x\|=1} \frac{\|Ax\|^2}{\langle Ax, x \rangle} \quad (45)$$

So we have  $\ell \geq \sup_{\|x\|=1} \frac{\|Ax\|^2}{\langle Ax, x \rangle}$ . Because the RHS is continuous in  $x$ , and the unit ball is a compact, this quantity is achieved. Thus there exists  $x \in \mathbb{R}^d$  such that

$$\|Ax\|^2 = \sup_{\|x\|=1} \frac{\|Ax\|^2}{\langle Ax, x \rangle} \langle Ax, x \rangle \quad (46)$$

Thus  $\ell \leq \sup_{\|x\|=1} \frac{\|Ax\|^2}{\langle Ax, x \rangle}$ , which concludes the proof.  $\square$

For instance, a class of quadratic games that are *not* co-coercive are the ones where  $A$  is anti-symmetric ( $A^\top = -A$ ). On the other hand, there is a large class of games that are *not* strongly monotone, like for instance any quadratic game induced by a matrix with a non-zero nullspace.



## B Proofs of Main Convergence Analysis Results

### B.1 Proof of Theorem 4.1

In the main paper, we present the update rule of SGDA in (5). Let us also present here the pseudo-code of SGDA:

---

#### Algorithm 2 Stochastic Gradient Descent Ascent (SGDA)

---

**Input:** Starting stepsize  $\gamma_0 > 0$ . Choose initial points  $x^0 \in \mathbb{R}^d$ . Distribution  $\mathcal{D}$  of samples.  
**for**  $k = 0, 1, 2, \dots, K$  **do**  
    Sample  $v^k \sim \mathcal{D}$   
    Set step-size  $\gamma_k$  following one of the selected choices (constant, decreasing)  
    Set  $x^{k+1} = x^k - \gamma_k \xi_{v^k}(x)$   
**end for**  
**Output:** The last iterate  $x^k$

---

Let us present a more general version of Theorem 4.1 that allows convergence with a larger step-size. Due to space limitations, we focus only on the important regime in the main paper, as selecting a larger step-size gives a worse convergence rate.

**Theorem B.1** (Constant Step-size). Assume that  $\xi$  is  $\mu$ -quasi strongly monotone and that  $\xi \in EC(\ell_\xi)$ . Choose  $\alpha_k = \alpha < \frac{1}{\ell_\xi}$  for all k. Then, the iterates of SGDA, given by (5), satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq [1 - 2\alpha\mu(1 - \alpha\ell_\xi)]^k \|x^0 - x^*\|^2 + \frac{\alpha\sigma^2}{\mu(1 - \alpha\ell_\xi)} \quad (47)$$

and if  $\alpha_k = \alpha \in (0, \frac{1}{2\ell_\xi}]$  then the iterates of SGDA satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \alpha\mu)^k \|x^0 - x^*\|^2 + \frac{2\alpha\sigma^2}{\mu} \quad (48)$$

*Proof.*

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|x^k - \alpha\xi_{v^k}(x^k) - x^*\|^2 \\ &= \|x^k - x^*\|^2 - 2\langle x^k - x^*, \alpha\xi_{v^k}(x^k) \rangle + \|\alpha\xi_{v^k}(x^k)\|^2 \\ &= \|x^k - x^*\|^2 - 2\alpha\langle x^k - x^*, \xi_{v^k}(x^k) \rangle + \alpha^2\|\xi_{v^k}(x^k)\|^2 \end{aligned} \quad (49)$$

By taking expectation condition on  $x^k$ :

$$\begin{aligned} \mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] &= \|x^k - x^*\|^2 - 2\alpha\langle x^k - x^*, \xi(x^k) \rangle + \alpha^2\mathbb{E}_{\mathcal{D}} [\|\xi_{v^k}(x^k)\|^2] \\ &\stackrel{\text{Lemma 3.4}}{\leq} \|x^k - x^*\|^2 - 2\alpha\langle x^k - x^*, \xi(x^k) \rangle + 2\alpha^2\ell_\xi\langle x^k - x^*, \xi(x^k) \rangle + 2\alpha^2\sigma^2 \\ &\stackrel{(2), \alpha < \frac{1}{\ell_\xi}}{\leq} \|x^k - x^*\|^2 - 2\alpha\mu(1 - \alpha\ell_\xi)\|x^k - x^*\|^2 + 2\alpha^2\sigma^2 \end{aligned} \quad (50)$$

Recursively applying the above and summing up the resulting geometric series gives:

$$\begin{aligned} \mathbb{E} [\|x^k - x^*\|^2] &\leq [1 - 2\alpha\mu(1 - \alpha\ell_\xi)]^k \|x^0 - x^*\|^2 + 2\sum_{j=0}^{k-1} (1 - 2\alpha\mu(1 - \alpha\ell_\xi))^j \alpha^2 \sigma^2 \\ &\leq [1 - 2\alpha\mu(1 - \alpha\ell_\xi)]^k \|x^0 - x^*\|^2 + \frac{\alpha\sigma^2}{\mu(1 - \alpha\ell_\xi)} \end{aligned} \quad (51)$$

If we further take  $\alpha \leq \frac{1}{2\ell_\xi}$  then (50) becomes:

$$\mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] \leq (1 - \alpha\mu)\|x^k - x^*\|^2 + 2\alpha^2\sigma^2 \quad (52)$$

and by recursively applying the above and summing up the resulting geometric series gives:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \alpha\mu)^k \|x^0 - x^*\|^2 + \frac{2\alpha\sigma^2}{\mu} \quad (53)$$

□

**Comment on the convergence deterministic Gradient Descent Ascent:** In the main paper, to highlight the generality of Theorem 4.1, we present Corollary 4.2 on the convergence of deterministic gradient descent ascent. Let us provide some more details of how one can obtain such a result through Proposition 3.5.

Let us select the sampling vector  $v = (1, 1, \dots, 1) \in R^n$  with probability 1 in each step. Note that this is still a sampling vector as  $\mathbb{E}_{\mathcal{D}}[v_i] = 1$ . In this case, at iteration  $k$ ,  $\xi_{v^k}(x^k) := \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x^k) \stackrel{v_i=1}{=} \frac{1}{n} \sum_{i=1}^n \xi_i(x^k) = \xi(x^k)$  and the update rule becomes equivalent to the deterministic GDA:

$$x^{k+1} = x^k - \alpha_k \xi(x^k).$$

In addition, by Proposition 3.5 we have that if  $|S| = n$  with probability one (each iteration of SGDA uses a full batch gradient), then  $\ell_\xi = \ell$  and  $\sigma^2 = 0$ . Thus, by combining (7) of Theorem 4.1 with Proposition 3.5 we obtain the convergence given in Corollary 4.2 for the deterministic gradient descent ascent. We highlight that for this case, the expected co-coercivity condition (EC) is equivalent to assuming that operator  $\xi$  is  $\ell$ -co-coercive.

## B.2 Proof of Theorem 4.3

**Theorem B.2.** Assume  $\xi$  is  $\mu$ -quasi-strongly monotone and that  $\xi \in EC(\ell_\xi)$ . Let  $\mathcal{K} := \ell_\xi / \mu$  and let

$$\alpha_k = \begin{cases} \frac{1}{2\ell_\xi} & \text{for } k \leq 4\lceil\mathcal{K}\rceil \\ \frac{2k+1}{(k+1)^2\mu} & \text{for } k > 4\lceil\mathcal{K}\rceil. \end{cases} \quad (54)$$

If  $k \geq 4\lceil\mathcal{K}\rceil$ , then iterates of SGDA, given by (5) satisfy:

$$\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma^2}{\mu^2} \frac{8}{k} + \frac{16\lceil\mathcal{K}\rceil^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right). \quad (55)$$

*Proof.* Let  $\alpha_k := \frac{2k+1}{(k+1)^2\mu}$  and let  $k^*$  be an integer that satisfies  $\alpha_{k^*} \leq \frac{1}{2\ell_\xi}$ . Note that  $\alpha_k$  is decreasing in  $k$  and consequently  $\alpha_k \leq \frac{1}{2\ell_\xi}$  for all  $k \geq k^*$ . This in turn guarantees that (52) holds for all  $k \geq k^*$  with  $\alpha_k$  in place of  $\alpha$ , that is

$$\mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] \leq (1 - \alpha_k \mu) \|x^k - x^*\|^2 + 2\alpha_k^2 \sigma^2 \quad (56)$$

Hence, if we take expectations and replace  $\alpha_k := \frac{2k+1}{(k+1)^2\mu}$  then

$$\mathbb{E}\|x^{k+1} - x^*\|^2 \leq \frac{k^2}{(k+1)^2} \mathbb{E}\|x^k - x^*\|^2 + \frac{2\sigma^2}{\mu^2} \frac{(2k+1)^2}{(k+1)^4}. \quad (57)$$

Multiplying both sides by  $(k+1)^2$  we obtain

$$\begin{aligned} (k+1)^2 \mathbb{E}\|x^{k+1} - x^*\|^2 &\leq k^2 \mathbb{E}\|x^k - x^*\|^2 + \frac{2\sigma^2}{\mu^2} \left(\frac{2k+1}{k+1}\right)^2 \\ &\leq k^2 \mathbb{E}\|x^k - x^*\|^2 + \frac{8\sigma^2}{\mu^2}, \end{aligned}$$

where the second inequality holds because  $\frac{2k+1}{k+1} < 2$ . Rearranging and summing from  $t = k^* \dots k$  we obtain:

$$\sum_{t=k^*}^k [(t+1)^2 \mathbb{E}\|x^{t+1} - x^*\|^2 - t^2 \mathbb{E}\|x^t - x^*\|^2] \leq \sum_{t=k^*}^k \frac{8\sigma^2}{\mu^2}. \quad (58)$$

Using telescopic cancellation gives

$$(k+1)^2 \mathbb{E} \|x^{k+1} - x^*\|^2 \leq (k^*)^2 \mathbb{E} \|x^{k^*} - x^*\|^2 + \frac{8\sigma^2(k-k^*)}{\mu^2}.$$

Dividing the above by  $(k+1)^2$  gives

$$\mathbb{E} \|x^{k+1} - x^*\|^2 \leq \frac{(k^*)^2}{(k+1)^2} \mathbb{E} \|x^{k^*} - x^*\|^2 + \frac{8\sigma^2(k-k^*)}{\mu^2(k+1)^2}. \quad (59)$$

For  $k \leq k^*$  we have that (53) holds with  $\alpha_k = \frac{1}{2\ell_\xi}$ , which combined with (59), gives

$$\begin{aligned} \mathbb{E} \|x^{k+1} - x^*\|^2 &\leq \frac{(k^*)^2}{(k+1)^2} \left(1 - \frac{\mu}{2\ell_\xi}\right)^{k^*} \|x^0 - x^*\|^2 \\ &\quad + \frac{\sigma^2}{\mu^2(k+1)^2} \left(8(k-k^*) + \frac{(k^*)^2}{\mathcal{K}}\right). \end{aligned} \quad (60)$$

Choosing  $k^*$  that minimizes the second line of the above gives  $k^* = 4\lceil\mathcal{K}\rceil$ , which when inserted into (60) becomes

$$\begin{aligned} \mathbb{E} \|x^{k+1} - x^*\|^2 &\leq \frac{16\lceil\mathcal{K}\rceil^2}{(k+1)^2} \left(1 - \frac{1}{2\mathcal{K}}\right)^{4\lceil\mathcal{K}\rceil} \|x^0 - x^*\|^2 \\ &\quad + \frac{\sigma^2}{\mu^2} \frac{8(k-2\lceil\mathcal{K}\rceil)}{(k+1)^2} \\ &\leq \frac{16\lceil\mathcal{K}\rceil^2}{e^2(k+1)^2} \|x^0 - x^*\|^2 + \frac{\sigma^2}{\mu^2} \frac{8}{k+1}, \end{aligned} \quad (61)$$

where we have used that  $\left(1 - \frac{1}{2x}\right)^{4x} \leq e^{-2}$  for all  $x \geq 1$ .  $\square$

### B.3 Proof of Theorem 5.1

Before providing the proof of Theorem 5.1, let us present the definitions of quasi-strong convexity and expected smoothness condition, together with a lemma that provides a bound to the expected norm of the stochastic gradients when a function satisfies the expected smoothness. Recall that in the main paper, we assume that the Hamiltonian function  $\mathcal{H}(x)$  is quasi-strongly convex and  $\mathcal{L}$ —expected smooth (satisfies the expected smoothness condition). Thus, these assumptions are vital for the convergence guarantees of SCO presented in Section 5.

**Technical Background on Optimization.** Let us consider the optimization problem

$$x^* = \operatorname{argmin}_{x \in \mathbb{R}^d} \left[ f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x) \right], \quad (62)$$

where each  $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$  is smooth and  $f$  has a unique global minimizer  $x^*$ .

**Definition B.3** (Quasi-strong convexity). We say that a function  $f : \mathbb{R}^d \rightarrow \mathbb{R}$  is  $\mu$ —strongly quasi-convex [Karimi et al., 2016, Necoara et al., 2018] if there is  $\mu > 0$  such that:

$$f(x^*) \geq f(x) + \langle \nabla f(x), x^* - x \rangle + \frac{\mu}{2} \|x^* - x\|^2 \quad (63)$$

for all  $x \in \mathbb{R}^d$ . Here  $x^*$  is the global minimizer of  $f^a$ .

Note that we have already presented the expected smoothness condition in Table 1. Below we present its formal definition. For this definition we use the stochastic reformulation of the finite-sum problem  $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ . That is, we define:  $f_v(x) := \frac{1}{n} \sum_{i=1}^n v_i f_i(x)$  where  $v \sim \mathcal{D}$  is a random sampling vector (see Section 2 for more details on sampling vectors).

<sup>a</sup>In our setting we assume that  $x^*$  is unique, but in the more general setting,  $x^*$  is the projection of point  $x$  onto the solution set  $X^*$  minimizing  $f$ .

**Definition B.4** (Expected Smoothness). We say that  $f$  is  $\mathcal{L}$ —expected smooth with respect to a distribution  $\mathcal{D}$  if there exists  $\mathcal{L} = \mathcal{L}(f, \mathcal{D}) > 0$  such that

$$\mathbb{E}_{\mathcal{D}} \left[ \|\nabla f_v(x) - \nabla f_v(x^*)\|^2 \right] \leq 2\mathcal{L}(f(x) - f(x^*)), \quad (64)$$

for all  $x \in \mathbb{R}^d$ .

In the next lemma, by assuming that a function  $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$  satisfies the expected smoothness we are able to bound the expected norm of its stochastic gradients. This is precisely the result we use in our proofs on the convergence of SCO, to upper bound  $\mathbb{E}_{\mathcal{D}} [\|\nabla \mathcal{H}_{v^k, u^k}(x^k)\|^2]$ . This bound allows us to avoid the much stronger bounded gradient or bounded variance assumptions.

**Lemma B.5** (Lemma 2.4 in Gower et al. [2019]). If  $f$  is  $\mathcal{L}$ —expected smooth, then

$$\mathbb{E}_{\mathcal{D}} [\|\nabla f_v(x)\|^2] \leq 4\mathcal{L}(f(x) - f(x^*)) + 2\sigma^2. \quad (65)$$

where  $\sigma^2 := \mathbb{E}_{\mathcal{D}} [\|\nabla f_v(x^*)\|^2]$ .

*Proof.* The proof can be easily obtained by following the same steps of the proof of Lemma 3.4. See also the proof of Lemma 2.4 in Gower et al. [2019].  $\square$

Let us now present a more general version of Theorem 5.1 that allows convergence with a larger step-size  $\alpha$ . Due to space limitations, in the main paper, we focus only on the important regime of step-size  $\alpha$  as selecting a larger step-size gives a worse convergence rate.

**Theorem B.6** (Constant Step-size). Assume  $\xi$  is  $\mu$ -quasi-strongly monotone with  $\mu \geq 0$  and that  $\xi \in EC(\ell_\xi)$ . Let us also assume that the Hamiltonian function  $\mathcal{H}$  is  $\mu_{\mathcal{H}}$ -quasi strongly convex and  $\mathcal{L}_{\mathcal{H}}$ -expected smooth. Then, for  $\gamma_k = \gamma \leq 1/4\mathcal{L}_{\mathcal{H}}$  and  $\alpha_k = \alpha < 1/2\ell_\xi$  it holds that :

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}} - 2\alpha\mu + 4\alpha^2\ell_\xi\mu)^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu_{\mathcal{H}} + 2\alpha\mu - 4\alpha^2\ell_\xi\mu}, \quad (66)$$

and if  $\alpha_k = \alpha \in (0, \frac{1}{4\ell_\xi}]$  then the iterates of SCO satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}} - \alpha\mu)^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu_{\mathcal{H}} + \alpha\mu}. \quad (67)$$

If  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma\mu_{\mathcal{H}})^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2\sigma^2 + \gamma^2\sigma_{\mathcal{H}}^2]}{\gamma\mu_{\mathcal{H}}}. \quad (68)$$

*Proof.*

$$\begin{aligned} \|x^{k+1} - x^*\|^2 &= \|x^k - \alpha\xi_{v^k}(x^k) - \gamma\nabla\mathcal{H}_{v^k, u^k}(x^k) - x^*\|^2 \\ &= \|x^k - x^*\|^2 - 2\langle x^k - x^*, \alpha\xi_{v^k}(x^k) + \gamma\nabla\mathcal{H}_{v^k, u^k}(x^k) \rangle \\ &\quad + \|\alpha\xi_{v^k}(x^k) + \gamma\nabla\mathcal{H}_{v^k, u^k}(x^k)\|^2 \\ &= \|x^k - x^*\|^2 - 2\alpha\langle x^k - x^*, \xi_{v^k}(x^k) \rangle - 2\gamma\langle x^k - x^*, \nabla\mathcal{H}_{v^k, u^k}(x^k) \rangle \\ &\quad + \|\alpha\xi_{v^k}(x^k) + \gamma\nabla\mathcal{H}_{v^k, u^k}(x^k)\|^2 \\ &\stackrel{\text{Young's}}{\leq} \|x^k - x^*\|^2 - 2\alpha\langle x^k - x^*, \xi_{v^k}(x^k) \rangle - 2\gamma\langle x^k - x^*, \nabla\mathcal{H}_{v^k, u^k}(x^k) \rangle \\ &\quad + 2\alpha^2\|\xi_{v^k}(x^k)\|^2 + 2\gamma^2\|\nabla\mathcal{H}_{v^k, u^k}(x^k)\|^2 \end{aligned} \quad (69)$$

By taking expectation condition on  $x^k$ :

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] &\leq \|x^k - x^*\|^2 - 2\alpha \langle x^k - x^*, \xi(x^k) \rangle - 2\gamma \langle x^k - x^*, \nabla \mathcal{H}(x^k) \rangle \\
&\quad + 2\alpha^2 \mathbb{E}_{\mathcal{D}} [\|\xi_{v^k}(x^k)\|^2] + 2\gamma^2 \mathbb{E}_{\mathcal{D}} [\|\nabla \mathcal{H}_{v^k, u^k}(x^k)\|^2] \\
&\stackrel{(63)}{\leq} \|x^k - x^*\|^2 - 2\alpha \langle x^k - x^*, \xi(x^k) \rangle - 2\gamma [\mathcal{H}(x^k) - \mathcal{H}(x^*)] - \gamma \mu_{\mathcal{H}} \|x^k - x^*\|^2 \\
&\quad + 2\alpha^2 \mathbb{E}_{\mathcal{D}} [\|\xi_{v^k}(x^k)\|^2] + 2\gamma^2 \mathbb{E}_{\mathcal{D}} [\|\nabla \mathcal{H}_{v^k, u^k}(x^k)\|^2] \\
&\stackrel{(65)}{\leq} \|x^k - x^*\|^2 - 2\alpha \langle x^k - x^*, \xi(x^k) \rangle - 2\gamma [\mathcal{H}(x^k) - \mathcal{H}(x^*)] - \gamma \mu_{\mathcal{H}} \|x^k - x^*\|^2 \\
&\quad + 2\alpha^2 \mathbb{E}_{\mathcal{D}} [\|\xi_{v^k}(x^k)\|^2] + 8\gamma^2 \mathcal{L}_{\mathcal{H}}(\mathcal{H}(x) - \mathcal{H}(x^*)) + 4\gamma^2 \sigma_{\mathcal{H}}^2 \\
&\stackrel{(16)}{\leq} \|x^k - x^*\|^2 - 2\alpha \langle x^k - x^*, \xi(x^k) \rangle - 2\gamma [\mathcal{H}(x^k) - \mathcal{H}(x^*)] - \gamma \mu_{\mathcal{H}} \|x^k - x^*\|^2 \\
&\quad + 4\alpha^2 \ell_{\xi} \langle \xi(x), x - x^* \rangle + 4\alpha^2 \sigma^2 + 8\gamma^2 \mathcal{L}_{\mathcal{H}}(\mathcal{H}(x) - \mathcal{H}(x^*)) + 4\gamma^2 \sigma_{\mathcal{H}}^2 \quad (70)
\end{aligned}$$

Recall that  $\xi$  is  $\mu$ -quasi strongly monotone,  $\langle x^k - x^*, \xi(x^k) \rangle \geq \mu \|x^k - x^*\|^2$ . Thus, for  $\alpha_k < \frac{1}{2\ell_{\xi}}$ , it holds that:

$$(-2\alpha_k + 4\alpha_k^2 \ell_{\xi}) \langle \xi(x^k), x - x^* \rangle \leq (-2\alpha_k + 4\alpha_k^2 \ell_{\xi}) \mu \|x^k - x^*\|^2,$$

and the inequality (70) takes the following form:

$$\begin{aligned}
\mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] &\leq (1 - \gamma \mu_{\mathcal{H}}) \|x^k - x^*\|^2 + (-2\alpha + 4\alpha^2 \ell_{\xi}) \mu \|x^k - x^*\|^2 \\
&\quad + (-2\gamma + 8\gamma_k^2 \mathcal{L}_{\mathcal{H}}) [\mathcal{H}(x^k) - \mathcal{H}(x^*)] + 4\alpha^2 \sigma^2 + 4\gamma^2 \sigma_{\mathcal{H}}^2 \\
&= (1 - \gamma \mu_{\mathcal{H}} - 2\alpha \mu + 4\alpha^2 \ell_{\xi} \mu) \|x^k - x^*\|^2 \\
&\quad + (-2\gamma + 8\gamma^2 \mathcal{L}_{\mathcal{H}}) [\mathcal{H}(x^k) - \mathcal{H}(x^*)] + 4\alpha^2 \sigma^2 + 4\gamma^2 \sigma_{\mathcal{H}}^2 \\
&\stackrel{\gamma_k < \frac{1}{4\mathcal{L}_{\mathcal{H}}}}{\leq} (1 - \gamma \mu_{\mathcal{H}} - 2\alpha \mu + 4\alpha^2 \ell_{\xi} \mu) \|x^k - x^*\|^2 + 4[\alpha^2 \sigma^2 + \gamma^2 \sigma_{\mathcal{H}}^2] \quad (71)
\end{aligned}$$

By taking expectations again and by recursively applying the above and summing up the resulting geometric series gives:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma \mu_{\mathcal{H}} - 2\alpha \mu + 4\alpha^2 \ell_{\xi} \mu)^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2 \sigma^2 + \gamma^2 \sigma_{\mathcal{H}}^2]}{\gamma \mu_{\mathcal{H}} + 2\alpha \mu - 4\alpha^2 \ell_{\xi} \mu}$$

If we further assume that  $\alpha \leq \frac{1}{4\ell_{\xi}}$  then  $1 - \gamma \mu_{\mathcal{H}} - 2\alpha \mu + 4\alpha^2 \ell_{\xi} \mu \leq 1 - \gamma \mu_{\mathcal{H}} - \alpha \mu$  and the iterates of SCO satisfy:

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma \mu_{\mathcal{H}} - \alpha \mu)^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2 \sigma^2 + \gamma^2 \sigma_{\mathcal{H}}^2]}{\gamma \mu_{\mathcal{H}} + \alpha \mu}. \quad (72)$$

In addition, if  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then

$$\mathbb{E} [\|x^k - x^*\|^2] \leq (1 - \gamma \mu_{\mathcal{H}})^k \|x^0 - x^*\|^2 + \frac{4[\alpha^2 \sigma^2 + \gamma^2 \sigma_{\mathcal{H}}^2]}{\gamma \mu_{\mathcal{H}}}. \quad (73)$$

This completes the proof.  $\square$

**On Deterministic Consensus Optimization.** In Corollary 5.2 we show the convergence of Deterministic CO as special case of our main Theorem. Here we provide few more details to understand exactly this convergence.

Let us select the sampling vectors  $v = u = (1, 1, \dots, 1) \in R^n$  with probability 1 in each step. Note that these are still sampling vectors as  $\mathbb{E}_{\mathcal{D}}[v_i] = \mathbb{E}_{\mathcal{D}}[u_i] = 1$ . In this case, at iteration  $k$ ,  $\xi_{v^k}(x^k) := \frac{1}{n} \sum_{i=1}^n v_i \xi_i(x^k) \stackrel{v_i=1}{=} \frac{1}{n} \sum_{i=1}^n \xi_i(x^k) = \xi(x^k)$  and  $\nabla \mathcal{H}_{u,v}(x) =$

$\frac{1}{2} [\mathbf{J}_u^\top(x)\xi_v(x) + \mathbf{J}_v^\top(x)\xi_u(x)] \stackrel{v_i=1, u_i=1}{=} \frac{1}{2} [\mathbf{J}^\top(x)\xi(x) + \mathbf{J}^\top(x)\xi(x)] = \mathbf{J}^\top(x)\xi(x) = \nabla\mathcal{H}(x)$ . Thus, the update rule becomes equivalent to the deterministic CO:

$$x^{k+1} = x^k - \alpha\xi(x^k) - \gamma\nabla\mathcal{H}(x^k).$$

In addition, by Proposition 3.5 we have that if  $|S| = n$  with probability one, then  $\ell_\xi = \ell$  and  $\sigma^2 = 0$ . Using also the properties of expected smoothness (see Proposition 3.8 in Gower et al. [2019]) we have that  $\mathcal{L}_\mathcal{H} = L_\mathcal{H}$ , where  $L_\mathcal{H}$  is the smoothness parameter of the Hamiltonian function, and  $\sigma_\mathcal{H}^2 = 0$ . By simply substituting these values to the main theorem we are able to obtain the convergence result presented in Corollary 5.2.

**Convergence of SGDA and SHGD as special cases of our Analysis.** As we mentioned in the main paper, SCO is a weighted combination of SGDA (Algorithm 2) and the stochastic Hamiltonian gradient descent (SHGD) of Loizou et al. [2020]. Thus, it is clear, that if one selects  $\alpha^k = 0, \forall k > 0$  then the method is equivalent to SHGD and if  $\gamma^k = 0, \forall k > 0$  then the method becomes equivalent to the SGDA. Here we highlight that in these cases the Young's inequality used in (69) of the proof of main theorem is not necessary. This is exactly why by specifying the update rule to SHGD we are able to have convergence using the larger bound on the step-size  $\gamma \leq 1/2\mathcal{L}_\mathcal{H}$  (see Corollary 5.3). Following similar argument the convergence of SGDA presented in Theorem 4.1 can also be obtained as special case of Theorem 5.1.

#### B.4 Proof of Theorem 5.4

**Theorem B.7.** Assume  $\xi$  is  $\mu$ -quasi-strongly monotone and that  $\xi \in EC(\ell_\xi)$ . Assume that the Hamiltonian function  $\mathcal{H}$  is  $\mu_\mathcal{H}$ -quasi strongly convex and  $\mathcal{L}_\mathcal{H}$ -expected smooth. Let  $\alpha_k = \gamma_k$ ,  $\psi = \max\{\ell_\xi, \mathcal{L}_\mathcal{H}\}$  and  $k^* := 8 \frac{\psi}{\mu_\mathcal{H} + \mu}$ . Let also,

$$\gamma_k = \begin{cases} \frac{1}{4\psi} & \text{for } k \leq \lceil k^* \rceil \\ \frac{2k+1}{(k+1)^2[\mu_\mathcal{H} + \mu]} & \text{for } k > \lceil k^* \rceil. \end{cases} \quad (74)$$

If  $k \geq \lceil k^* \rceil$ , then SCO iterates satisfy:

$$\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma_\mathcal{H}^2 + \sigma^2}{[\mu + \mu_\mathcal{H}]^2} \frac{16}{k} + \frac{(k^*)^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right) \quad (75)$$

If  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then SCO is still able to converge sublinearly with  $O\left(\frac{1}{k}\right)$ , to  $x^*$ .

*Proof.* Let  $\gamma_k := \frac{2k+1}{(k+1)^2[\mu_\mathcal{H} + \mu]}$  and let  $k^*$  be an integer that satisfies

$$\gamma_{k^*} \leq \min\left\{\frac{1}{4\ell_\xi}, \frac{1}{4\mathcal{L}_\mathcal{H}}\right\} = \frac{1}{4\psi}.$$

Note that  $\gamma_k$  is decreasing in  $k$  and consequently  $\gamma_k \leq \frac{1}{4\psi}$  for all  $k \geq k^*$ . This in turn guarantees that (71) holds for all  $k \geq k^*$  with  $\gamma_k$  in place of  $\gamma$  and  $\alpha_k = \gamma_k$  in place of  $\alpha$ , that is

$$\mathbb{E}_\mathcal{D} [\|x^{k+1} - x^*\|^2] \leq (1 - \gamma_k\mu_\mathcal{H} - 2\gamma_k\mu + 4\gamma_k^2\ell_\xi\mu)\|x^k - x^*\|^2 + 4[\gamma_k^2\sigma^2 + \gamma_k^2\sigma_\mathcal{H}^2]$$

and since  $\gamma_k \leq \frac{1}{4\ell_\xi}$  we obtain:

$$\mathbb{E}_\mathcal{D} [\|x^{k+1} - x^*\|^2] \leq (1 - \gamma_k[\mu_\mathcal{H} + \mu])\|x^k - x^*\|^2 + 4\gamma_k^2[\sigma^2 + \sigma_\mathcal{H}^2]. \quad (76)$$

For simplicity of presentation let us denote  $\bar{\mu} = [\mu_\mathcal{H} + \mu]$  and  $\bar{\sigma}^2 = [\sigma^2 + \sigma_\mathcal{H}^2]$  then (76) can be written as:

$$\mathbb{E}_\mathcal{D} [\|x^{k+1} - x^*\|^2] \leq (1 - \gamma_k\bar{\mu})\|x^k - x^*\|^2 + 4\gamma_k^2\bar{\sigma}^2. \quad (77)$$

Now let us follow similar steps to the proof of Theorem 4.3.

By taking expectations and replacing  $\gamma_k := \frac{2k+1}{(k+1)^2 \lceil \mu\gamma + \mu \rceil} = \frac{2k+1}{(k+1)^2 \bar{\mu}}$  we obtain,

$$\mathbb{E}\|x^{k+1} - x^*\|^2 \leq \frac{k^2}{(k+1)^2} \mathbb{E}\|x^k - x^*\|^2 + \frac{4\bar{\sigma}^2}{\bar{\mu}^2} \frac{(2k+1)^2}{(k+1)^4}. \quad (78)$$

Multiplying both sides by  $(k+1)^2$  we obtain

$$\begin{aligned} (k+1)^2 \mathbb{E}\|x^{k+1} - x^*\|^2 &\leq k^2 \mathbb{E}\|x^k - x^*\|^2 + \frac{4\bar{\sigma}^2}{\bar{\mu}^2} \left( \frac{2k+1}{k+1} \right)^2 \\ &\leq k^2 \mathbb{E}\|x^k - x^*\|^2 + \frac{16\bar{\sigma}^2}{\bar{\mu}^2}, \end{aligned}$$

where the second inequality holds because  $\frac{2k+1}{k+1} < 2$ . Rearranging and summing from  $t = k^* \dots k$  we obtain:

$$\sum_{t=k^*}^k [(t+1)^2 \mathbb{E}\|x^{t+1} - x^*\|^2 - t^2 \mathbb{E}\|x^t - x^*\|^2] \leq \sum_{t=k^*}^k \frac{16\bar{\sigma}^2}{\bar{\mu}^2}. \quad (79)$$

Using telescopic cancellation gives

$$(k+1)^2 \mathbb{E}\|x^{k+1} - x^*\|^2 \leq (k^*)^2 \mathbb{E}\|x^{k^*} - x^*\|^2 + \frac{8\bar{\sigma}^2(k - k^*)}{\bar{\mu}^2}.$$

Dividing the above by  $(k+1)^2$  gives

$$\mathbb{E}\|x^{k+1} - x^*\|^2 \leq \frac{(k^*)^2}{(k+1)^2} \mathbb{E}\|x^{k^*} - x^*\|^2 + \frac{8\bar{\sigma}^2(k - k^*)}{\bar{\mu}^2(k+1)^2}. \quad (80)$$

At this point note that for  $k \leq k^*$  we have that (77) holds and by using our step-size selection  $\gamma_k = \gamma = \frac{1}{4\psi}$  for  $k \leq \lceil k^* \rceil$  we obtain

$$\mathbb{E}_{\mathcal{D}} [\|x^{k+1} - x^*\|^2] \leq (1 - \gamma\bar{\mu})\|x^k - x^*\|^2 + 4\gamma^2\bar{\sigma}^2, \quad (81)$$

which by taking expectations again and by recursively applying the above and summing up the resulting geometric series gives (for  $k \leq k^*$ ):

$$\mathbb{E} [\|x^{k+1} - x^*\|^2] \leq (1 - \gamma\bar{\mu})^k \|x^0 - x^*\|^2 + \frac{4\gamma\bar{\sigma}^2}{\bar{\mu}}, \quad (82)$$

Thus, for  $k \leq k^*$  we have that (82) holds with  $\gamma_k = \frac{1}{4\psi}$ , which combined with (80), gives

$$\begin{aligned} \mathbb{E}\|x^{k+1} - x^*\|^2 &\leq \frac{(k^*)^2}{(k+1)^2} \left(1 - \frac{\bar{\mu}}{4\psi}\right)^{k^*} \|x^0 - x^*\|^2 \\ &\quad + \frac{\bar{\sigma}^2}{\bar{\mu}^2(k+1)^2} \left(16(k - k^*) + \frac{(k^*)^2\bar{\mu}}{\psi}\right). \end{aligned} \quad (83)$$

Choosing  $k^*$  that minimizes the second line of the above gives  $k^* = 8\frac{\psi}{\bar{\mu}}$ , which when inserted into (83) becomes

$$\begin{aligned} \mathbb{E}\|x^{k+1} - x^*\|^2 &\leq \frac{(k^*)^2}{(k+1)^2} \left(1 - \frac{2}{k^*}\right)^{k^*} \|x^0 - x^*\|^2 \\ &\quad + \frac{\bar{\sigma}^2}{\bar{\mu}^2(k+1)^2} 8(2k - k^*) \\ &\leq \frac{(k^*)^2}{(k+1)^2 e^2} \|x^0 - x^*\|^2 + \frac{\bar{\sigma}^2}{\bar{\mu}^2(k+1)^2} 8(2k - k^*) \\ &\leq \frac{(k^*)^2}{(k+1)^2 e^2} \|x^0 - x^*\|^2 + \frac{\bar{\sigma}^2}{\bar{\mu}^2} \frac{16}{k+1}. \end{aligned} \quad (84)$$

where in the second inequality we have used that  $(1 - \frac{1}{2x})^{4x} \leq \frac{1}{e^2}$  for all  $x \geq 1$  and in the last inequality we used that  $\frac{2k-k^*}{k+1} \leq \frac{2k}{k+1} \leq 2$ .

Thus by replacing  $\bar{\mu} = [\mu_{\mathcal{H}} + \mu]$  and  $\bar{\sigma}^2 = [\sigma^2 + \sigma_{\mathcal{H}}^2]$  we obtain:

$$\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma_{\mathcal{H}}^2 + \sigma^2}{[\mu + \mu_{\mathcal{H}}]^2} \frac{16}{k} + \frac{(k^*)^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right).$$

As we mentioned in the statement of the Theorem, if  $\mu = 0$ , that is  $\xi$  only satisfies the variational stability condition  $\langle \xi(x), x - x^* \rangle \geq 0$ , then SCO is still able to converge sublinearly with  $O\left(\frac{1}{k}\right)$ , to  $x^*$ . In this case, the proof will be exactly the same as above but  $k^* := 8 \frac{\psi}{\mu_{\mathcal{H}}}$  and  $\gamma_k = \frac{2k+1}{(k+1)^2 \mu_{\mathcal{H}}}$  for  $k > \lceil k^* \rceil$ . And the convergence will be  $\mathbb{E}\|x^k - x^*\|^2 \leq \frac{\sigma_{\mathcal{H}}^2 + \sigma^2}{\mu_{\mathcal{H}}^2} \frac{16}{k} + \frac{(k^*)^2}{e^2 k^2} \|x^0 - x^*\|^2 = O\left(\frac{1}{k}\right)$ .  $\square$



## C On Experiments

### C.1 Properties of Hamiltonian Function for Quadratic Games

In the next proposition, we explain how the assumptions on the Hamiltonian function used in the main theorems of Section 5 (convergence analysis of SCO) are satisfied for the quadratic min-max problems.

**Proposition C.1.** For quadratic games of the form (13) with  $\mathbf{A}_i$  and  $\mathbf{C}_i$  symmetric with at least one solution  $x^*$ , the Hamiltonian function  $\mathcal{H}(x)$  is a  $L_{\mathcal{H}}$ -smooth and  $\mu_{\mathcal{H}}$ -quasi-strongly convex quadratic function with constants  $L_{\mathcal{H}} = \sigma_{\max}^2(\mathbf{J})$  and  $\mu_{\mathcal{H}} = \sigma_{\min}^2(\mathbf{J})$  where  $\sigma_{\max}$  and  $\sigma_{\min}$  are the maximum and minimum non-zero singular values of  $\mathbf{J}$ , and  $\mathbf{J} = \nabla \xi$  is the Jacobian matrix of the game.

*Proof.* Our approach follows closely the proof of Proposition 4.3 of Loizou et al. [2020] where the properties of the Hamiltonian function for stochastic bilinear games were presented. Let  $\mathbf{A} = \frac{1}{n} \sum_{i=1}^n \mathbf{A}_i$ ,  $\mathbf{B} = \frac{1}{n} \sum_{i=1}^n \mathbf{B}_i$  and  $\mathbf{C} = \frac{1}{n} \sum_{i=1}^n \mathbf{C}_i$  and let  $a = \frac{1}{n} \sum_{i=1}^n a_i$  and  $c = \frac{1}{n} \sum_{i=1}^n c_i$ .

Firstly, note that the stochastic Hamiltonian function of (13) has the following form:

$$\mathcal{H}(x) = \frac{1}{2} x^\top \mathbf{Q} x + q^\top x + \ell$$

where  $\mathbf{Q} = \begin{pmatrix} \mathbf{A}^2 + \mathbf{B}\mathbf{B}^\top & \mathbf{A}\mathbf{B} - \mathbf{B}\mathbf{C} \\ \mathbf{B}^\top \mathbf{A} - \mathbf{C}\mathbf{B}^\top & \mathbf{C}^2 + \mathbf{B}^\top \mathbf{B} \end{pmatrix} = \mathbf{J}^\top \mathbf{J}$ ,  $q = \begin{pmatrix} \mathbf{A}a - \mathbf{B}c \\ \mathbf{B}^\top a + \mathbf{C}c \end{pmatrix}$ , and  $\ell = \frac{1}{2}(\|a\|^2 + \|b\|^2)$

The Hamiltonian is thus a smooth quadratic function and the matrix  $\mathbf{Q}$  is symmetric and positive semi-definite. In addition, as we assumed that there was at least one solution  $x^*$  for the game, i.e.  $\xi(x^*) = 0$ , we have that  $x^*$  is also a global minimum of the Hamiltonian function  $\mathcal{H}(x)$ , and thus we have that  $q = -\mathbf{Q}x^* = -\mathbf{J}^\top \mathbf{J}x^*$ . Using this, we can rewrite the Hamiltonian as:

$$\mathcal{H}(x) = \phi(\mathbf{J}x),$$

where the function  $\phi(y) := \frac{1}{2}\|y\|^2 - (\mathbf{J}x^*)^\top y + \ell$  is 1-strongly convex with 1-Lipschitz continuous gradient.

Thus, using Lemma D.1 in Loizou et al. [2020], we have that the the Hamiltonian function is a  $L_{\mathcal{H}}$ -smooth,  $\mu_{\mathcal{H}}$ -quasi-strongly convex function with constants  $L_{\mathcal{H}} = \lambda_{\max}(\mathbf{J}^\top \mathbf{J}) = \lambda_{\max}(\mathbf{Q}) = \sigma_{\max}^2(\mathbf{J})$  and  $\mu_{\mathcal{H}} = \sigma_{\min}^2(\mathbf{J}) = \lambda_{\min}^+(\mathbf{Q})$ .  $\square$

### C.2 Experimental Details

We describe here in more details the exact settings we use for evaluating the different algorithms. As mentioned in Section 6, we evaluate the different algorithms on the class of quadratic games:

$$\min_{x_1 \in \mathbb{R}^d} \max_{x_2 \in \mathbb{R}^p} \frac{1}{n} \sum_i \frac{1}{2} x_1^\top \mathbf{A}_i x_1 + x_1^\top \mathbf{B}_i x_2 - \frac{1}{2} x_2^\top \mathbf{C}_i x_2 + a_i^\top x_1 - c_i^\top x_2$$

In all our experiments we choose  $d = p = 100$  and  $n = 100$ . To sample the matrices  $\mathbf{A}_i$  (resp.  $\mathbf{C}_i$ ) we first generate a random orthogonal matrix  $\mathbf{Q}_i$  (resp.  $\mathbf{Q}'_i$ ), we then sample a random diagonal matrix  $\mathbf{D}_i$  (resp.  $\mathbf{D}'_i$ ) where the elements on the diagonal are sampled uniformly in  $[\mu_A, L_A]$  (resp.  $[\mu_C, L_C]$ ), such that at least one of the matrices has a minimum eigenvalue equal to  $\mu_A$  (resp.  $\mu_C$ ) and one matrix has a maximum eigenvalue equal to  $L_A$  (resp.  $L_B$ ). Finally we construct the matrices by computing  $\mathbf{A}_i = \mathbf{Q}_i \mathbf{D}_i \mathbf{Q}_i^\top$  (resp.  $\mathbf{C}_i = \mathbf{Q}'_i \mathbf{D}'_i \mathbf{Q}'_i{}^\top$ ). This ensures that the matrices  $\mathbf{A}_i$  and  $\mathbf{C}_i$  for all  $i \in [n]$ , are symmetric and positive definite. We sample the matrices  $\mathbf{B}_i$  in a similar fashion with the diagonal matrix  $\mathbf{D}_i$  to lie between  $[\mu_B, L_B]$ <sup>9</sup>. In all our experiments we choose  $\mu_A = \mu_C$  and  $L_A = L_C$ . By varying the different constants  $\mu_A, L_A, \mu_C, L_C, \mu_B, L_B$  we can get a variety of games with different properties  $\mu, \ell_\xi, \mu_{\mathcal{H}}, L_{\mathcal{H}}$ . The bias terms  $a_i, c_i$  are sampled from a normal distribution. For further details, see also our source code<sup>10</sup>.

<sup>9</sup>We highlight that matrices  $\mathbf{B}_i$  are not necessarily symmetric.

<sup>10</sup><https://github.com/hugobb/StochasticGamesOpt>

As we have already mentioned in Section 6, we pick the step-sizes for the different methods according to our theoretical findings. That is, for constant step-size, we select  $\alpha = \frac{1}{2\ell_\xi}$  for SGDA (Theorem 4.1),  $\alpha = \frac{1}{4\ell_\xi}, \gamma = \frac{1}{4\mathcal{L}_\mathcal{H}}$  for SCO (Theorem 5.1), and  $\gamma = \frac{1}{2\mathcal{L}_\mathcal{H}}$  for SHGD (Corollary 5.3). For the stepsize-switching rule that guarantees convergence to  $x^*$ , we use the step-sizes proposed in Theorem 4.3 for SGDA and Theorem 5.4 for SCO.

In the experiments, we run all methods (SGDA, SCO and SHGD) using uniform single-element sampling. That is,  $|S| = 1$ , according to the Definition 2.1. Thus,  $\mathbb{P}[i \in S] = p_i = \frac{1}{n}$ . By Proposition 3.5, this means that  $\ell_\xi = \ell_{\max} = \max\{\ell_i\}_{i=1}^n$ . In addition, by Proposition 3.8 in Gower et al. [2019] and the structure of the stochastic Hamiltonian function,  $\mathcal{H}(x) = \frac{1}{2}\|\xi(x)\|^2 = \frac{1}{n} \sum_{i=1}^n \frac{1}{n} \sum_{j=1}^n \frac{1}{2} \langle \xi_i(x), \xi_j(x) \rangle$ , we have that  $\mathcal{L}_\mathcal{H} = \max\{L_{\mathcal{H}_{i,j}}\}_{i=1,j=1}^n$ , where  $L_{\mathcal{H}_{i,j}}$  is the smoothness parameter of  $\mathcal{H}_{i,j} = \frac{1}{2} \langle \xi_i(x), \xi_j(x) \rangle$ .

The values of the co-coercive parameters  $\ell_i$  for all  $i \in [n]$  are computed using  $\frac{1}{\ell_i} = \min_{\lambda \in Sp(\mathbf{J}_i)} \Re(\frac{1}{\lambda})$  [Azizian et al., 2020]. Here  $Sp(\mathbf{J}_i)$  denotes the spectrum of the Jacobian matrix  $\mathbf{J}_i$  for all  $i \in [n]$  and  $\Re$  denotes the real part of a complex number.

### C.3 Additional Experiment: Influence of the Step-size on Convergence

In this section we provide further experiments exploring the performance of SGDA, SHGD and SCO for different constant step-sizes (see Fig. 3). This experiment aims to understand better how the convergence rate and the size of the neighborhood the algorithms converge to depend on the step size. This also enables us to assess if the optimal step-size suggested by the theory is tight. When the step-size is too large we observe that the methods diverges, we do not include them in the plots but provide them below for completeness. We observe that SHGD diverges with a step-size of  $\gamma = 3\gamma^*$  with  $\gamma^* = \frac{1}{2\mathcal{L}_\mathcal{H}}$ , SGDA diverges with a step-size of  $\alpha = 4\alpha^*$  with  $\alpha^* = \frac{1}{2\ell_\xi}$ , SCO diverges with a step-size of  $\alpha = 4\alpha^*$  and  $\gamma = 4\gamma^*$  with  $\alpha^* = \frac{1}{4\ell_\xi}$  and  $\gamma^* = \frac{1}{4\mathcal{L}_\mathcal{H}}$ . We also observe that SCO is less sensitive to the choice of  $\alpha$  than to the choice of  $\gamma$

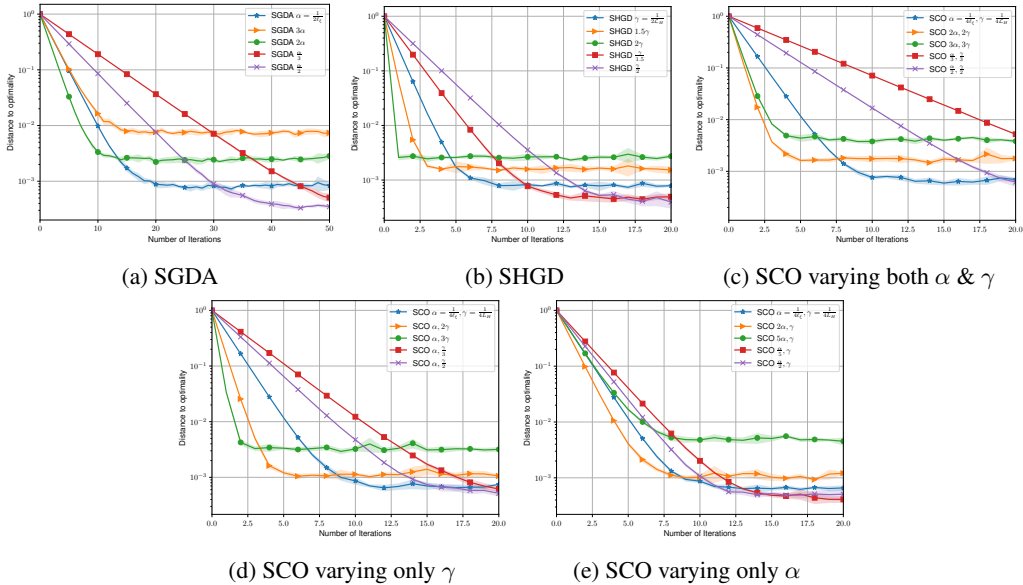


Figure 3: Performance of the methods, SGDA, SHGD and SCO with different constant step-sizes. Distance to optimality:  $\frac{\|x^k - x^*\|^2}{\|x^0 - x^*\|^2}$ . The step-sizes used are expressed as the optimal theoretical step-size predicted by the theory times a constant. The step-sizes for which the methods diverge are not included in the plots.

## D Beyond Finite-sum Structure

We explain in this section how all our convergence results also hold for the more general (non finite-sum) stochastic approximation setting [Nemirovski et al., 2009], where we define  $\xi : \mathbb{R}^d \rightarrow \mathbb{R}^d$  as:

$$\xi(x) := \mathbb{E}_U \tilde{\xi}(x, U), \quad (85)$$

where  $U$  is a random vector in  $\mathbb{R}^p$ , and  $\tilde{\xi} : \mathbb{R}^d \times \mathbb{R}^p \rightarrow \mathbb{R}^d$  is used to define  $\xi$  and is assumed to be well-behaved enough so that (85) exists for all  $x$ . Setting  $U$  to be a uniform random variable in  $\{1, \dots, n\}$  gives back the finite sum setting with uniform weights which was covered in the paper in (1).

We present here only results for the singleton sampling regime, as generalizing to mini-batching and other sampling schemes in the continuous regime is non-trivial and beyond the scope of this paper. This means that we restrict here our estimator for  $\xi(x)$  to simply be  $\tilde{\xi}(x, U)$ , with  $U$  sampled according to its distribution. Any appearance of  $\xi_v$  in the paper can then be replaced with  $\tilde{\xi}(x, U)$  to be able to re-interpret the algorithms and proofs in this setting. In particular, we have that our estimator is  $g(x) = \xi_v(x) := \tilde{\xi}(x, U)$ , and so  $\mathbb{E}[g(x)] = \mathbb{E}_U \tilde{\xi}(x, U) = \xi(x)$  trivially. Also, the expected co-coercivity of  $\xi$  is defined as the existence of a  $\ell_\xi > 0$  such that

$$\mathbb{E}_U \left[ \|\tilde{\xi}(x, U) - \tilde{\xi}(x^*, U)\|^2 \right] \leq \ell_\xi \langle \xi(x), x - x^* \rangle \quad \forall x \in \mathbb{R}^d. \quad (86)$$

An important quantity appearing in our convergence results is

$$\sigma^2 := \mathbb{E}_U \|\tilde{\xi}(x, U)\|^2, \quad (87)$$

and we assume that  $U$  and  $\tilde{\xi}$  are such that  $\sigma^2 < \infty$ .<sup>11</sup> Similarly,  $\sigma_{\mathcal{H}}^2$  is defined by letting  $u$  and  $v$  represent independent singleton sampling in (10), and is assumed to be finite.

Under the assumption of expected co-coercivity of  $\xi$  and that  $\sigma^2$  is finite, all the convergence theorems of Section 4 and 5 hold as is for  $\xi$  more generally defined by (85). This is because all the convergence proofs did not use the finite sum structure explicitly, only the linearity of the expectation operator.

Finally, we can easily generalize Proposition 3.5 for  $b = 1$  with  $\ell_\xi = \ell_{\max}$  (we assume each  $\tilde{\xi}(x, u)$  is  $\ell_u$ -co-coercive in  $x$  around  $x^*$ , and assume that  $\ell_{\max} := \sup_u \ell_u$  is finite).

We can also generalize Proposition 3.6 (for singleton sampling) with  $\ell_\xi := \mathbb{E}_U L_U^2 / \mu$  (we assume each  $\tilde{\xi}(x, u)$  is  $L_u$ -Lipschitz continuous in  $x$ , for each  $u$ ; and that  $\mathbb{E}_U L_U^2$  is finite).

## E More Related Work

The references necessary to motivate our work and connect it to the most relevant literature is included in the appropriate sections of the main body of the paper. Here we present a broader view of the literature, including some more references to papers of the area that are not directly related with our work.

**Smooth monotone games.** The deterministic version of the problem has been studied extensively both in past and recent work, initially focusing on strongly monotone problems [Tseng, 1995, Gidel et al., 2018, Liang and Stokes, 2019, Azizian et al., 2020, Zhou et al., 2021]. Mokhtari et al. [2020] recently gave rates for extragradient and optimistic gradient through a proximal point approach.

For monotone problems, in the absence of strong monotonicity but assuming a lower bound of the singular values of the coupling between players, Azizian et al. [2020] produce tight results for the extragradient and optimistic gradient methods.

For general monotone problems, Mertikopoulos and Zhou [2019] establish last-iterate convergence, but requires a decreasing step-size schedule; it also does not guarantee convergence to fixed points of non-strictly monotone problems like bilinears. Golowich et al. [2020b] show that last-iterate methods are slower than averaging in this setting. Overcoming the above limitations, Golowich et al. [2020b,a] establish tight upper bounds,  $O(1/\sqrt{T})$ , for general monotone problems under a weak smoothness assumption for extra-gradient and optimistic gradient respectively.

<sup>11</sup>In the finite sum setting, this is always true. But for general random variable  $U$ , this is not always the case.

**Stochastic smooth monotone games.** Contrary to the deterministic family of problems, in the stochastic setting (where available gradients are noisy), progress has occurred mostly in the past five years. Without variance reduction, some work establishes last-iterate convergence results with a slow (sublinear) rate [Rosasco et al., 2014, Mishchenko et al., 2020]. For *pseudomonotone problems*, Kannan and Shanbhag [2019] give last-iterate convergence of stochastic extragradient with decreasing step sizes, assuming a uniform bound on the variance. Using variance reduction techniques, Palaniappan and Bach [2016] yield fast linear rates without requiring a bound on the variance, however the step size and inner loop length need to be tuned using the modulus of local strong monotonicity. This limitation was later lifted with a variance reduced extragradient method proposed by Chavdarova et al. [2019], which is adaptive to local strong monotonicity.

For monotone problems, in the absence of strong monotonicity but assuming a lower bound of the singular values of the coupling between players, Chavdarova et al. [2019] show that stochastic extragradient with constant step-size does not always yield last-iterate convergence. Such a result for extragradient was shown to be possible using a *double step-size scheme* in Hsieh et al. [2020], using one step size for the extrapolation step and a different step size for the update step.

For general monotone problems, Mertikopoulos and Zhou [2019] use a dual averaging approach to get last-iterate convergence for no-regret algorithms making a bounded variance assumption. This bounded variance assumption is lifted in Lin et al. [2020b]; the authors give finite-time last-iterate convergence of optimistic gradient under other strong conditions: i) for a *relative random noise* model, the noise variance is assumed to be proportional to the squared norm of the vector of gradients, according to the sequence of multiplicative coefficients,  $\tau_k$ . That is,  $\mathbb{E}_i \|\xi_i(x_k) - \xi(x_k)\|^2 < \tau_k \|\xi(x_k)\|^2$ . This kind of condition with a constant  $\tau_k = \tau$  can occur in some machine learning problems, in particular in supervised learning problems when the model is overparametrized and capable to perfectly fitting all training points (so the gradient variance becomes zero at the stationary point of the solution). On the other hand, this assumption is satisfied rarely for adversarial formulations. Even more, Theorem 4.6 in Lin et al. [2020b] requires that the sequence,  $\tau_k$ , goes to zero with  $k$  in order to get convergence; ii) for an *absolute random noise* model, noise variance is uniformly bounded by  $\sigma_k$ , a bounded variance assumption. Again, for the result of Theorem A.4 in Lin et al. [2020b] to hold, the sequence  $\sigma_k$  is assumed to go to zero with  $k$ . To the best of our knowledge, both assumptions are very strong: in order to get either  $\tau_k$  or  $\sigma_k$  to decay, one will have to use a mini-batch size that grows up to  $n$ , which is impractical. It should be noted that the same work contains almost sure (non finite-time) last-iterate results without the above restrictive assumptions.

**Structured non-monotone problems.** The recent works of Daskalakis et al. [2021] and Diakonikolas et al. [2021] show that, for general smooth objectives, the computation of even approximate first-order locally optimal min-max solutions is intractable, motivating the identification of structural assumptions on the objective function for which these intractability barriers can be bypassed.

In this work we focus on a setting (structured non-monotone operators) that one can provide tight convergence guarantees and avoid the standard issues (cycling and divergence of the methods) appearing in the more general non-monotone regime, like . That is, problems satisfying quasi-strong monotonicity (2) or the variational stability condition.

The two classes of VI problems we consider, quasi-strongly monotone and problems satisfying the variational stability condition, have already been used in several papers (under different names). For example, Song et al. [2020] also focuses on quasi-strongly monotone VI but they called them strong coherent VI and show that quasi strong monotonicity is weaker than the strongly pseudomonotone [Kannan and Shanbhag, 2019] and strongly monotone assumption. In addition, as presented in Song et al. [2020], the non-monotone subsets of quasi-strong monotonicity (a.k.a., strongly pseudomonotone), have many real applications in competitive exchange economy [Brighi and John, 2002], fractional programming [Elizarov and Kalimullina, 2009, Rousseau et al., 2005], and product pricing [Choi et al., 1990]. Meanwhile, the restriction of quasi-strong monotonicity in minimization problems such as one-point convexity [Li and Yuan, 2017] is also used in analyzing neural networks.

In addition in Zhou et al. [2017], the strongly coherent optimization problems have been studied, which are special cases of the quasi-strongly monotone VI. We refer the interested reader to Zhou et al. [2017] for a list of examples of strongly coherent optimization problems. As we mentioned in the main paper, the assumption of quasi-strong monotonicity is equivalent to the concept of strong stability described in Mertikopoulos and Zhou [2019].