IMPROVING THE LAST-ITERATE GUARANTEES OF ANYTIME ALGORITHMS FOR STOCHASTIC MONOTONE VARIATIONAL INEQUALITIES

Anonymous authors
Paper under double-blind review

ABSTRACT

We analyze a stochastic algorithm with Halpern-type anchoring for constrained convex-concave problems and monotone variational inequalities. This single-loop and single-call algorithm uses one unbiased sample of the gradient operator at every iteration, to be applicable to monotone games with noisy feedback. With t denoting the iteration counter, we prove an anytime last-iterate convergence rate of $O(t^{-1/4})$ for both the gradient-mapping norm and restricted gap, bypassing the $O(t^{-1/5})$ constrained-anytime bottleneck in the literature. Specializing then to multi-point oracles, we use variance reduction to achieve the $O(t^{-1/2})$ rate with an anytime single-loop algorithm using 2 samples per iteration. Our results allow constrained problems with a potentially unbounded feasible set; as well as a structured class of stochastic oracles whose variance need not be uniformly bounded.

1 INTRODUCTION

We focus on the unifying framework of monotone variational inequalities (VIs) that cover such problems as monotone games and convex-concave min-max problems. In particular, we focus on the specific inclusion describing this problem, given as

$$0 \in G(\mathbf{z}^\star) + \partial r(\mathbf{z}^\star), \tag{1.1}$$

where $G\colon \mathbb{R}^d \to \mathbb{R}^d$ is monotone, L-Lipschitz; and $r\colon \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ is proper, convex, closed. This is equivalent to the following form of the monotone VI, where the goal is to find $\mathbf{z}^\star$ such that

$$\langle G(\mathbf{z}^\star), \mathbf{z} - \mathbf{z}^\star \rangle + r(\mathbf{z}) - r(\mathbf{z}^\star) \geq 0 \text{ for every } \mathbf{z} \in \mathbb{R}^d. \tag{1.2}$$

An important application of (1.2) is learning in monotone games (Cesa-Bianchi & Lugosi, 2006), where Nash equilibria can be characterized as solutions of a monotone VI, see, e.g., (Mertikopoulos & Zhou, 2019). In this setting, one action profile is played and one noisy payoff feedback is observed at each round. Hence, single-call methods fit this feedback model, whereas two-call methods require an additional gradient evaluation at a different action profile.

Since the current action profile is the one actually played, the horizon is generally unknown, and we are interested in observing the dynamics of the learning process, last-iterate and anytime guarantees are particularly relevant (Cai & Zheng, 2023b). As a result, standard results for stochastic VI algorithms that show convergence rates on the averaged iterate or those that use increasing mini-batches, variance reduction, or multi-loop algorithms are not suitable for this setting. Motivated by this setup, in Section 2, we study an anytime, single-call algorithm with last iterate guarantees.

Another classical application outside of the learning in games setting is constrained convex optimization

$$\min_{\mathbf{x}} f(\mathbf{x})\colon A\mathbf{x} \leq \mathbf{b} \iff \min_{\mathbf{x}} \max_{\mathbf{y} \geq 0} f(\mathbf{x}) + \langle A\mathbf{x} - \mathbf{b}, \mathbf{y} \rangle, \tag{1.3}$$

where single-call algorithms are not necessary in general, and variance reduced methods are applicable. In particular, we often have a large matrix A from which we sample rows or columns to get unbiased estimates for gradients, which is amenable to the oracle access for variance reduction. An important aspect in this setting is that we need algorithms that can be applied for unbounded constraint sets since neither the primal nor the dual domains are bounded in (1.3). In Section 3, we study a single-loop algorithm suitable for the optimization setting by using variance reduction with 2 calls per iteration.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Method	Constraint	Anytime	Last-iterate rate and measure	Variance assumption
(Abe et al., 2025)	Bounded	×	$\widetilde{O}(T^{-1/7})$, gap	Bounded
(Zheng et al., 2026)	Unbounded	×	$O(T^{-1/4})$, gap	Bounded[‡]
(Zheng et al., 2026)	Unbounded	✓	$O(t^{-1/5})$, gap	Bounded[‡]
(Ito et al., 2026)	Bounded	×	$O(T^{-1/4})$, gap	Bounded
(Ito et al., 2026)	Bounded	✓	$O(t^{-1/5})$, gap	Bounded
(Sohrabi et al., 2026)	Unconstrained	✓	$O(t^{-1/4})$, grad. norm	(1.7)
Section 2	Unbounded	✓	$O(t^{-1/4})$, grad. mapping $O(t^{-1/4})$, gap	Blum-Gladyshev

Table 1: Comparison of single-loop and anytime stochastic methods with last-iterate guarantees and without variance reduction with monotone and Lipschitz G. The rates are stated with T for fixed horizon and t for anytime results. The Blum-Gladyshev condition is given in Assumption 1.2. [‡](Zheng et al., 2026, Section 2.1) remarked that their analysis extends to a stronger variant of our Assumption 1.2, but did not provide a proof.

Optimality measures. Two standard optimality measures we consider are gradient mapping norm, also known as the natural residual, and the restricted gap function (Facchinei & Pang, 2003). In particular, for a fixed $\rho > 0$, we define the gradient mapping by

$$\mathcal{G}_\rho(\mathbf{z}) = \frac{1}{\rho}\big(\mathbf{z} - \operatorname{prox}_{\rho r}(\mathbf{z} - \rho G(\mathbf{z}))\big). \tag{1.4}$$

The norm $\|\mathcal{G}_\rho(\mathbf{z})\|$ vanishes if and only if (1.1) holds at $\mathbf{z}$. For simplicity we will set $\rho = \frac{1}{L}$ throughout. This is an optimality measure generalizing the gradient norm in the unconstrained case and is suitable for problems with unbounded feasible sets, in contrast to the gap we introduce next.

We present the gap function for the special case of $r(\mathbf{z}) = \delta_Z(\mathbf{z})$ where δ_Z is the indicator function for a closed and convex set Z:

$$\operatorname{gap}(\mathbf{z}) = \max_{\mathbf{u}\in Z} \{\langle G(\mathbf{z}), \mathbf{z} - \mathbf{u}\rangle\}. \tag{1.5}$$

Since this quantity is generally not suitable when set Z is not bounded, restricted versions of gap functions are often used (Nesterov, 2007). For restricted gap functions, one generally takes the maximum over a compact set $\mathcal{U}$ and this set needs to satisfy certain requirements for restricted gap functions to be valid optimality measures (Nesterov, 2007, Lemma 1). For example, the set needs to contain the output of the algorithm and a solution. Since it is generally not possible to show that the iterates stay bounded for stochastic algorithms, this optimality measure is less meaningful without a bounded domain or for stochastic algorithms. On the other hand, when the domain is bounded, such as in matrix games, the gap functions are well-defined optimality measures.

To compare with some relevant works, we also prove rates for the gap function, which are direct consequences of the rates we prove on the natural residual/gradient mapping norm, given in (1.4). The relationship between the gap function and the gradient mapping is formalized in Corollary D.2 for completeness, which is a standard argument.

Related works. In the deterministic case, optimal last iterate guarantees are obtained recently by anytime algorithms (Yoon & Ryu, 2021; Cai & Zheng, 2023a;b) by using anchoring; and guarantees for the last iterate of standard VI algorithms, such as the classical extragradient method (Korpelevich, 1976), were shown in (Golowich et al., 2020b; Gorbunov et al., 2022; Cai et al., 2022b). In the game setting, last iterate guarantees are proven in (Golowich et al., 2020a) for the optimistic gradient algorithm (Popov, 1980; Rakhlin & Sridharan, 2013). In the stochastic case, last iterate guarantees without using mini-batches or multi-loop algorithms has been of interest for learning in games with noisy feedback, and many works focused on this direction, including, but not limited to, (Chen &

Method	Constraint	Anytime	Batch	Complexity and output	Variance assumption
(Cai et al., 2022a)	Unconstrained	×	large†	$\widetilde{O}(\varepsilon^{-3})$, last-iter.	Bounded
(Tran-Dinh & Nguyen-Trung, 2026)	Unbounded	×	large†	$O(\varepsilon^{-10/3})$, random-iter.	Bounded
(Alacaoglu et al., 2025)	Unbounded	×	2	$O(\varepsilon^{-4})$, random-iter.	Blum-Gladyshev
Section 3	Unbounded	✓	2	$O(\varepsilon^{-2})$, last-iter.	Blum-Gladyshev

Table 2: Comparison of single-loop methods with last-iterate guarantees and variance reduction with monotone and Lipschitz G. †large denotes batch sizes that scale as a power of $1/\varepsilon$.

Mazumdar, 2024; Azizian et al., 2021; Abe et al., 2025; Sohrabi et al., 2026; Zheng et al., 2026; Ito et al., 2026; Abe et al., 2024; Hsieh et al., 2022; Mertikopoulos & Zhou, 2019; Jordan et al., 2025).

The literature on the last iterate guarantees for anytime, single-loop algorithms using a single sample (or 2 samples) at every iteration had recent interesting developments. In particular, for unconstrained problems (Sohrabi et al., 2026) showed the $O(t^{-1/4})$ rate for an anchored algorithm. For constrained problems, two parallel works showed the anytime rate $O(t^{-1/5})$ for the restricted gap function (Zheng et al., 2026; Ito et al., 2026). Surprisingly, both of these works concluded the following separation: when we let go of the *anytime* requirement, that is, when the horizon T is known in advance, it was possible to get the rate $O(T^{-1/4})$ for the last iterate where T is the horizon; yet for the anytime case, both works got the worse $O(t^{-1/5})$ rate.

Other single-loop algorithms in the literature are not compatible with the requirements of our setting: for example the developments in (Cai et al., 2022a) focused on unconstrained problems (for monotone and Lipschitz operators) and used variance reduction along with increasing batches; the works (Pethick et al., 2023a; Alacaoglu & Kim, 2026; Alacaoglu et al., 2025) relied on variance reduction and multi-point oracles and did not prove last iterate guarantees. Classical works obtain optimal rates (in terms of ε independence) on the gap when it is evaluated at the averaged iterate, see for example (Nemirovski et al., 2009).

A large body of literature used increasing (or large depending on $1/\varepsilon$) mini-batch sizes to obtain $O(\varepsilon^{-4})$ complexity (Iusem et al., 2017; Pethick et al., 2023b; Kotsalis et al., 2022; Lee & Kim, 2021). Our complexity in Section 2 matches this by avoiding mini-batches. When we allow increasing mini-batches or multi-loop algorithms, the complexity for residual-type optimality measures was improved to $O(\varepsilon^{-10/3})$ by using mini-batches and variance reduction in the work (Tran-Dinh & Nguyen-Trung, 2026), under Assumption 3.1. Only very recently, methods based on multiple loops obtained the ε-optimal $\widetilde{O}(\varepsilon^{-2})$ complexity for gradient norm for unconstrained problems (Chen & Luo, 2024) and natural residual and tangent residual in the constrained case (Alacaoglu, 2026; Ji et al., 2026), respectively.

To our knowledge, the best complexity with a single-loop method for the residual remained at $O(\varepsilon^{-10/3})$, using large batches. With a constant number of calls, the complexity for the residual is $O(\varepsilon^{-4})$. We obtain the best-known and ε-optimal rate $O(\varepsilon^{-2})$ under Assumption 3.1 with a single-loop algorithm using 2 unbiased oracles at each iteration. Assumption 3.1 is stronger than the assumptions of (Alacaoglu, 2026; Ji et al., 2026), however, unlike these works, we use a single-loop method without using increasing mini-batch sizes and multiple loops.

Context and Contributions. The best anytime rate for the restricted gap stayed at $O\left(t^{-1/5}\right)$ (Zheng et al., 2026; Ito et al., 2026), whereas both works could obtain the horizon-dependent $O(T^{-1/4})$ rate where T is the horizon. At the same time, an anytime $O(t^{-1/4})$ was known for unconstrained problems (Sohrabi et al., 2026; Zheng et al., 2026). The difficulty for extending to the constrained case was discussed in (Zheng et al., 2026, Remark 4.2); open questions to derive an anytime $O(t^{-1/4})$ rate for constrained problems were mentioned in (Sohrabi et al., 2026, Section 7) and (Zheng et al., 2026, Section 6). More technical details are provided in Section 2.3.

We prove that the separation from (Zheng et al., 2026; Ito et al., 2026) can be avoided, that is, we prove the anytime rate $O\left(t^{-1/4}\right)$ for constrained problems. To allow unbounded constraint sets, we

show rates on gradient mapping norm (in addition to the gap function), without the bounded variance assumption, by allowing variance to grow as fast as the displacement of the iterates from an initial point. Our algorithm is of stochastic gradient-type, with anchoring (in view of Halpern (Halpern, 1967)), using a single-call of the unbiased operator at every iteration.

Next, by incorporating the STORM variance reduction of (Cutkosky & Orabona, 2019), we improve the rate on the last iterate to $O(t^{-1/2})$. The algorithm is still single-loop and anytime, but not single-call, since it samples two unbiased oracles. This leads to the stochastic first-order complexity $O(\varepsilon^{-2})$ which is optimal in terms of ε-dependence, under the multi-point oracle described in Assumption 3.1, see (Foster et al., 2019). This complexity was recently obtained in works (Alacaoglu, 2026; Ji et al., 2026) without Assumption 3.1, yet, with multi-loop algorithms using the recursive regularization framework. Even though our assumption on Lipschitzness and multi-point access to the oracle is stronger than these works, our algorithm is significantly simpler, since it does not use multiple loops or increasing batch sizes. Our assumptions still allow bilinear problems, see, e.g., Section 1.1.

Even though the analysis frameworks in Sections 2 and 3 are different, Halpern anchoring achieves the same two goals in both cases: the *contraction* brought by β_t is essential to *(i)* get the last iterate rate with the induction argument (which can be seen as taking the place of more standard analyses going through a single Lyapunov function); *(ii)* help cancel the error coming from Assumption 1.2.

1.1 Assumptions

Monotonicity and Lipschitz continuity of G mean that, for all $\mathbf{x}, \mathbf{y}$, we have

$$\langle G(\mathbf{x}) - G(\mathbf{y}), \mathbf{x} - \mathbf{y}\rangle \geq 0, \qquad \|G(\mathbf{x}) - G(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|.$$

We now collect the assumptions made on our main problem.

Assumption 1.1. The function $r\colon \mathbb{R}^d \to \mathbb{R} \cup \{+\infty\}$ is proper, closed, and convex. The operator $G : \mathbb{R}^d \to \mathbb{R}^d$ is monotone and L-Lipschitz, and the set of solutions to (1.1) is nonempty.

We fix a deterministic initial point $\mathbf{z}_0 \in \operatorname{dom} r$. Let $\mathcal{F}_t$ denote the history before the oracle call at iteration t, that is, it is the σ-algebra generated by the randomness of $\xi_0, \dots, \xi_{t-1}$, where $\mathbf{z}_t$ is $\mathcal{F}_t$-measurable, and we write $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot \mid \mathcal{F}_t]$. The oracle returns an unbiased sample whose variance may grow with the distance from the initial point.

Assumption 1.2. At each iteration $t \geq 0$, the sample $\widetilde{G}(\mathbf{z}_t, \xi_t)$ satisfies

$$\mathbb{E}_t[\widetilde{G}(\mathbf{z}_t, \xi_t)] = G(\mathbf{z}_t) \quad \text{and} \quad \mathbb{E}_t\|\widetilde{G}(\mathbf{z}_t, \xi_t) - G(\mathbf{z}_t)\|^2 \leq B^2\|\mathbf{z}_t - \mathbf{z}_0\|^2 + \sigma^2. \tag{1.6}$$

The case $B = 0$ recovers the bounded-variance assumption. Let us compare this to the assumption used in (Sohrabi et al., 2026) who required in the unconstrained case

$$\mathbb{E}\|\widetilde{G}(\mathbf{z}_t, \xi_t) - G(\mathbf{z}_t)\|^2 \leq \sigma^2 + c^2\|G(\mathbf{z}_t)\|^2. \tag{1.7}$$

To see that our assumption is weaker, we focus on the unconstrained case where $G(\mathbf{z}^\star) = 0$. Then, this assumption implies by Lipschitzness of G:

$$\|G(\mathbf{z}_t)\|^2 = \|G(\mathbf{z}_t) - G(\mathbf{z}^\star)\|^2 \leq L^2\|\mathbf{z}_t - \mathbf{z}^\star\|^2 \leq 2L^2\left(\|\mathbf{z}_t - \mathbf{z}_0\|^2 + \|\mathbf{z}_0 - \mathbf{z}^\star\|^2\right).$$

That is, the variance assumption of (Sohrabi et al., 2026) implies Assumption 1.2.

The reason for focusing on Assumption 1.2 is to support our claims for handling problems without bounded domains. In particular, let us assume that we have an affine operator $G(\mathbf{z}) = A\mathbf{z} - \mathbf{b}$ for $A \in \mathbb{R}^{d\times d}$ and we calculate unbiased oracles as $\widetilde{G}(\mathbf{z}) = dA_{:i}z_i - \mathbf{b}$ where i is selected uniformly at random and $A_{:i}$ is the i-th column. Then, Assumption 1.2 holds whereas standard bounded variance or assumptions of the form (1.7) may fail, unless the domain of the problem ($\operatorname{dom} r$ in our notation) is bounded. Since problems with affine G and unbounded domains are common, such as (1.3), our analysis is based on Assumption 1.2 to cover such problems.

Notation. When clear from the context, we shorten $\widetilde{G}(\mathbf{z}_t, \xi_t)$ as $\widetilde{G}(\mathbf{z}_t)$. For $\eta > 0$, we define the proximal operator as $\operatorname{prox}_{\eta r}(\mathbf{x}) = \operatorname{argmin}_{\mathbf{z}\in\mathbb{R}^d}\left\{r(\mathbf{z}) + \frac{1}{2\eta}\|\mathbf{z} - \mathbf{x}\|^2\right\}$.

The optimality condition for the problem appearing in the definition of the proximal operator, under our assumptions on r, gives the well known *prox-inequality*:

$$\mathbf{z} = \operatorname{prox}_{\eta r}(\mathbf{x}) \iff \mathbf{x} - \mathbf{z} \in \eta\partial r(\mathbf{z}) \iff \langle \mathbf{z} - \mathbf{x}, \mathbf{u} - \mathbf{z}\rangle \geq \eta(r(\mathbf{z}) - r(\mathbf{u})) \ \ \forall \mathbf{u} \in \mathbb{R}^d. \tag{1.8}$$

2 A SINGLE-CALL ALGORITHM

Algorithm 1 is a simple extension of stochastic gradient descent-ascent (or stochastic forward-backward method), when Halpern anchoring is introduced (Halpern, 1967). This algorithm appeared many times in the literature, with many different names, see for example (Ito et al., 2026; Neu & Okolo, 2024). It was analyzed by (Lee & Kim, 2021) with increasing mini-batch sizes. The work (Alacaoglu et al., 2025) analyzed this method with a single-call at every iteration and variable parameters, but their analysis followed a more standard *regret-type* framework, which is different than the analysis we employ to get a last iterate rate. This method is a single-call algorithm à la (Hsieh et al., 2019), uses no mini-batch, variance reduction or inner loops.

Algorithm 1 A single-call stochastic Halpern method

Require: Initial point $\mathbf{z}_0 \in \operatorname{dom} r$, $a \geq 2$, $H^2 = L^2 + 2B^2$. Set $\beta_t = \frac{a}{t+a}$ and $\eta_t = \frac{1}{H(t+a)^{3/4}}$.

1: **for** $t = 0, 1, \ldots$ **do**
2: Obtain an unbiased estimate $\widetilde{G}(\mathbf{z}_t)$ of $G(\mathbf{z}_t)$
3: $\mathbf{z}_{t+1} = \operatorname{prox}_{\eta_t r}\left((1-\beta_t)\mathbf{z}_t + \beta_t \mathbf{z}_0 - \eta_t \widetilde{G}(\mathbf{z}_t)\right)$
4: **end for**

We now present our main result which is an anytime convergence rate on the gradient mapping norm, or natural residual. Then, we continue with its implications to anytime rates on gap functions. The proof of the result is given in Appendix A.

Theorem 2.1. *Let Assumptions 1.1 and 1.2 hold. Then, for every $t \geq 1$, Algorithm 1 satisfies*

$$\mathbb{E}\|\mathcal{G}_{1/L}(\mathbf{z}_t)\| \leq \sqrt{\mathbb{E}\|\mathcal{G}_{1/L}(\mathbf{z}_t)\|^2} = O\left(\frac{(L+B)\|\mathbf{z}_0 - \mathbf{z}^\star\| + \sigma}{t^{1/4}}\right).$$

We continue with a direct corollary for a rate on the restricted gap, by using Corollary D.2.

Corollary 2.2. *Let $r = \delta_Z$ for a convex, closed set Z. We then have*

$$\mathbb{E}[\operatorname{gap}_{\mathcal{U}}(\mathbf{z}_t)] = O\left(\frac{1}{t^{1/4}}\right),$$

for any compact set $\mathcal{U} \subseteq Z$, where $\operatorname{gap}_{\mathcal{U}}$ extends (1.5) by taking the maximum over a compact set $\mathcal{U}$.

When we specialize to min-max problems, we can then particularize this result to a rate on the restricted primal-dual gap used in (Zheng et al., 2026). This result is given in the appendix as Corollary C.1 for direct comparison to this work.

Finally, we convert the convergence rate guarantees to iteration and sample complexity results.

Corollary 2.3. *Under the same setup as Theorem 2.1, to obtain $\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\| \leq \varepsilon$ or $\mathbb{E}[\operatorname{gap}_{\mathcal{U}}(\mathbf{z}_t)] \leq \varepsilon$ on the last iterate $\mathbf{z}_t$, the required total number of stochastic oracles is of the order $O(\varepsilon^{-4})$.*

The proof of Corollary 2.2 is in App. C and the proof of Corollary 2.3 is omitted since it is immediate.

2.1 APPLICATION FOR LEARNING IN GAMES

We now apply our anytime guarantee for monotone games (Cesa-Bianchi & Lugosi, 2006). Following the setup in (Cai & Zheng, 2023b; Ito et al., 2026), we have N players where i-th player is denoted as z_i and $\mathbf{z} = (z^i; z^{-i}) = (z^1, \ldots, z^N)$. Each player selects actions from sets Z^i which are assumed to be compact, and $Z = Z_1 \times \cdots \times Z_N$. Each player minimizes a convex loss function $\ell_i(z^i; x^{-i})$. The operator G is defined as

$$G(\mathbf{z}) = \begin{bmatrix} \nabla_{z_1} \ell_1(\mathbf{z}) \\ \vdots \\ \nabla_{z_N} \ell_N(\mathbf{z}) \end{bmatrix}.$$

We assume access to noisy estimates of $G(\mathbf{z})$ such that $\widetilde{G}(\mathbf{z}_t) = G(\mathbf{z}_t) + \zeta_t$ where $\mathbb{E}[\zeta_t] = 0$ and $\mathbb{E}\|\zeta_t\|^2 \leq B^2\|\mathbf{z}_t - \mathbf{z}_0\|^2 + \sigma^2$. Even though this allows handling problems without bounded feasible

sets, for this section, we assume that Z is compact, for direct comparison with (Ito et al., 2026). We define the constants

$$D = \max_{\mathbf{x},\mathbf{y}\in Z} \|\mathbf{x}-\mathbf{y}\|, \qquad U = \max_{\mathbf{z}\in Z} \|G(\mathbf{z})\|.$$

Set $a = 2$ in Algorithm 1 and define, for $\mathbf{u} \in \operatorname{dom} r$, the regret $\mathrm{Reg}_T(\mathbf{u}) = \sum_{t=0}^{T-1} \langle G(\mathbf{z}_t), \mathbf{z}_t - \mathbf{u}\rangle$.

A direct consequence of Theorem 2.1, Corollary 2.2, and (Ito et al., 2026, Section 7) (by selecting our algorithm parameters as (Ito et al., 2026, Section 7)) is that we obtain the convergence and regret guarantees

$$\mathbb{E}[\mathrm{gap}(\mathbf{z}_t)] = O\left(\frac{D(LD+U+\sigma)}{t^{1/4}}\right) \quad \text{and} \quad \mathbb{E}[\mathrm{Reg}_T(\mathbf{u})] = O\left(D(LD+U+\sigma)T^{3/4}\right). \tag{2.1}$$

Remark 2.4. Under the same bounded-domain and bounded-variance setting, (Ito et al., 2026, Theorem 5.1) obtained the gap bound $O(D(LD+U+\sigma)t^{-1/5})$. Thus our result improves the anytime rate from $O(t^{-1/5})$ to $O(t^{-1/4})$. Importantly, this improvement does not follow simply by using our parameter schedule in the analysis of Ito et al. (2026). See (Zheng et al., 2026) (who used a similar analysis as (Ito et al., 2026)) who showed that the best anytime rate achievable with this analysis framework is $O(t^{-1/5})$. As a direct consequence, our analysis also improves the regret guarantee obtained for the anytime algorithm in (Ito et al., 2026) which was $O(T^{4/5})$, to $O(T^{3/4})$. This regret guarantee is definitely suboptimal, however, our aim is to argue that Alg. 1 has sublinear regret, while improving the last iterate anytime guarantee for approaching a Nash equilibrium.

We continue with the two required results for the proof of Thm. 2.1, which is in App A. Proof structure of Thm. 2.1 is inspired by the simple and powerful idea that first appeared in the work (Sohrabi et al., 2026) and used also for the unconstrained result of (Zheng et al., 2026). The idea is to analyze the deterministic shadow sequence given in (2.2) and then provide a pointwise bound between the deterministic sequence and the original stochastic sequence.

Instead of the gradient norm used in these works, we use the gradient mapping, which satisfies the necessary Lipschitz properties and then we show that the gradient mapping norm majorizes the gap function (see Corollary D.2), to go around the difficulty of working directly with the gap function.

2.2 Convergence of the Deterministic Sequence

Starting from $\bar{\mathbf{z}}_0 = \mathbf{z}_0$, define the deterministic reference iterates by

$$\bar{\mathbf{z}}_{t+1} = \mathrm{prox}_{\eta_t r}\left((1-\beta_t)\bar{\mathbf{z}}_t + \beta_t \mathbf{z}_0 - \eta_t G(\bar{\mathbf{z}}_t)\right). \tag{2.2}$$

These iterates are used only in the analysis. This sequence is indeed the algorithm analyzed by (Cai & Zheng, 2026), referred to as the anchored gradient descent. We follow the arguments of (Cai & Zheng, 2026), by adapting the parameter choices. Their $t^{-1/2}$ stepsize is replaced here by $t^{-3/4}$; the estimates below require only $H \geq L$.

We now show the uniform boundedness of this deterministic sequence, with its proof in App. A.1.

Lemma 2.5. *Under Assumption 1.1, for the iterates generated by (2.2), we have for any $t \geq 0$ that*

$$\|\bar{\mathbf{z}}_t - \mathbf{z}^\star\| \leq \frac{3}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|, \qquad \|\bar{\mathbf{z}}_t - \mathbf{z}_0\| \leq \frac{5}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|. \tag{2.3}$$

Next, using the ideas from (Cai & Zheng, 2026), we get a guarantee on the optimality measure $\|G(\mathbf{z}) + \mathbf{s}\|$ where $\mathbf{s} \in \partial r(\mathbf{z})$. However, we later use this bound to transfer the guarantee to the gradient mapping norm since $\|G(\mathbf{z}) + \mathbf{s}\|$, just defined, does not admit the Lipschitzness properties required for the analysis. We now continue with this analysis of the deterministic references iterates. Its proof, following (Cai & Zheng, 2026) is given in Appendix A.2.

Lemma 2.6. *Let Assumption 1.1 hold. For every $t \geq 0$, we have*

$$\|\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t\| \leq \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t+a)}. \tag{2.4}$$

Moreover, for every $t \geq 1$, we find $\mathbf{s}_t \in \partial r(\bar{\mathbf{z}}_t)$ such that

$$\|G(\bar{\mathbf{z}}_t) + \mathbf{s}_t\| \leq \frac{5aH\|\mathbf{z}_0 - \mathbf{z}^\star\|}{2(t-1+a)^{1/4}}\left(\frac{3}{2} + \frac{L}{2H(t-1+a)^{3/4}}\right). \tag{2.5}$$

This is a rather *unsurprising* result since it is now classical. Indeed, starting from the work of (Yoon & Ryu, 2021) which is extended to the constrained case by (Kovalev & Gasnikov, 2022; Cai & Zheng, 2023a; Cai et al., 2024), it is known how to prove that the residual has the last iterate rate of $O(1/k)$ (with more sophisticated algorithms). In our case, since we are limited by the pointwise bounds in Section 2.3, we seek a simple bound for this sequence that will be *just* sufficient for our purposes with the parameter choices for β_t, η_t.

2.3 Pointwise Bounds between Stochastic and Deterministic Sequences

We follow the stochastic comparison argument of (Sohrabi et al., 2026), bounding the distance between the stochastic and deterministic iterates. For the constrained problem, we adapt the argument using Lipschitzness of the gradient mapping (which is due to the nonexpansiveness of the proximal operator), under Assumption 1.2. Since this variance assumption is weaker than the assumption in (Sohrabi et al., 2026) (see Section 1.1), our result extends this work even in the unconstrained case.

Interestingly, a similar analysis was used in (Zheng et al., 2026, Sec. 5.4) in the unconstrained case, but not in the constrained case where the strategy was characterizing the solution of a perturbed problem, similar to (Ito et al., 2026). Both works obtained the $O(t^{-1/5})$ anytime rate for the gap.

Our insight for the constrained extension departs from the discussions in (Sohrabi et al., 2026, Section 7) and (Zheng et al., 2026, Remark 4.2) in that we do not directly analyze the gap function in the constrained setting, but go though the gradient mapping. The advantages are three-fold: *(i)* gradient mapping norm is Lipschitz; *(ii)* gradient mapping norm is a classical optimality measure for problems without bounded constraint sets, making our result valid for problems with unbounded constraint sets; *(iii)* since gradient mapping norm majorizes the (restricted) gap (see Corollary D.2), the anytime rate we prove for the gradient mapping norm directly transfers to the same rate on the gap.

Proof of the next result is given in Appendix A.3, which tells us that the expected squared distance between the stochastic and deterministic trajectories decay as $1/\sqrt{t}$.

Lemma 2.7. *Under Assumptions 1.1 and 1.2, for every $t \geq 0$, we have*

$$\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 \leq \frac{\sigma^2 + \frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2}{H^2(a-\frac{1}{2})\sqrt{t+a}}. \tag{2.6}$$

3 A Halpern Method with STORM Variance Reduction

In this section, we depart from the game setting which came with specific requirements on access to the oracle. In particular, in the previous section, it was essential that the algorithm only accessed one unbiased oracle per random seed ξ, that is $\widetilde{G}(z_t, \xi_t)$. However, in optimization setting, such as for linearly constrained optimization (1.3), it is possible to access two-point oracles $(\widetilde{G}(z_t, \xi_t), \widetilde{G}(z_{t-1}, \xi_t))$. As we now prove, this oracle will helps us obtain a better, in fact ε-optimal, complexity.

Algorithm 2 is simple combination of Algorithm 1 and the STORM estimator (Cutkosky & Orabona, 2019). STORM has previously been combined with Halpern anchoring for extragradient methods (Alacaoglu et al., 2025) and forward-backward-forward methods (Alacaoglu & Kim, 2026; Pethick et al., 2023a), yet only suboptimal complexity results with $\widetilde{O}(\varepsilon^{-4})$ were shown for a randomly selected iterate, where the analysis followed a Lyapunov framework. In this section, we establish three recursions for three sequences that determine the residual norm and then apply induction, inspired by the deterministic analyses of (Cai & Zheng, 2026; Surina et al., 2026).

We replace the single-call stochastic gradient in Algorithm 1 with the STORM estimator, without an additional extragradient correction. This method uses two stochastic-calls and one proximal update, with no mini-batch or inner loops. We will prove the ε-optimal complexity of $O(\varepsilon^{-2})$ for the last iterate. Because we use variance reduction, we will use the mean-square smoothness property for the operator. This assumption in fact implies Assumption 1.2. Moreover, for the example sketched in Section 1.1, the following assumption is also satisfied.

Assumption 3.1. For some ξ, we can obtain the multi-point oracle $\widetilde{G}(\mathbf{x}, \xi)$ and $\widetilde{G}(\mathbf{y}, \xi)$. Let also $\mathbb{E}\|\widetilde{G}(\mathbf{z}_0, \xi)\|^2 < +\infty$ and suppose that G is Lipschitz in mean square-sense, that is:

$$\mathbb{E}_\xi\|\widetilde{G}(\mathbf{x}, \xi) - \widetilde{G}(\mathbf{y}, \xi)\|^2 \leq \Lambda^2\|\mathbf{x} - \mathbf{y}\|^2.$$

Algorithm 2 A STORM proximal Halpern method

Require: Initial point $\mathbf{z}_0 \in \operatorname{dom} r$, and $H = \max\left\{2L,\ 4\sqrt{1024\Lambda^2 + 4 + 16B^2}\right\}$. Set $\alpha_t = \beta_t = \frac{1}{t+1}$ and $\eta_t = \frac{1}{H\sqrt{t+1}}$. Draw ξ_0, set $\mathbf{v}_0 = \widetilde{G}(\mathbf{z}_0, \xi_0)$ and $\mathbf{z}_1 = \operatorname{prox}_{\eta_0 r}(\mathbf{z}_0 - \eta_0 \mathbf{v}_0)$

1: **for** $t = 1, 2, \ldots$ **do**
2: Draw one sample ξ_t and obtain unbiased estimates $\widetilde{G}(\mathbf{z}_t, \xi_t)$ and $\widetilde{G}(\mathbf{z}_{t-1}, \xi_t)$
3: $\mathbf{v}_t = \widetilde{G}(\mathbf{z}_t, \xi_t) + (1-\alpha_t)\big(\mathbf{v}_{t-1} - \widetilde{G}(\mathbf{z}_{t-1}, \xi_t)\big)$
4: $\mathbf{z}_{t+1} = \operatorname{prox}_{\eta_t r}\left((1-\beta_t)\mathbf{z}_t + \beta_t \mathbf{z}_0 - \eta_t \mathbf{v}_t\right)$
5: **end for**

We now show a last-iterate guarantee for the tangent residual, which is a slightly stronger optimality measure than the gradient mapping (1.4). That is, *tangent residual* (sometimes called *residual*) is

$$\operatorname{res}(\mathbf{z}) = \min_{\mathbf{s} \in \partial r(\mathbf{z})} \|G(\mathbf{z}) + \mathbf{s}\|.$$

In fact, we prove the rate for the quantity $\|G(\mathbf{z}) + \mathbf{s}\|$, where $\mathbf{s} \in \partial r(\mathbf{z})$ which majorizes the tangent residual, as well as the gradient mapping norm, see for example (Cai & Zheng, 2023a, Prop. 2).

Theorem 3.2. *Let Assumptions 1.1, 1.2 and 3.1 hold. Then, for every $t \geq 1$, Algorithm 2 satisfies*

$$\mathbb{E}[\operatorname{res}(\mathbf{z}_t)] \leq \sqrt{\mathbb{E}[\operatorname{res}(\mathbf{z}_t)^2]} \leq \frac{27H\max\{\sigma, 4\|\mathbf{z}_0 - \mathbf{z}^\star\|\}}{\sqrt{t}} = O\left(t^{-1/2}\right).$$

The proof appears in App. B. Unlike the analysis of Algorithm 1, we are not using the deterministic-tracking argument of (Sohrabi et al., 2026). Instead, the proof will follow from induction, inspired by the proof in (Cai & Zheng, 2026; Surina et al., 2026). The main distinction is that variance reduction brings a new recursion that is coupled with other recursions. In particular, we obtain recursions depending on the sequences $\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t+1}\|^2, \mathbb{E}\|\mathbf{z}_t - \mathbf{z}^\star\|^2, \mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2$. After showing that all these sequences have certain bounds, we can arrive at the claimed rate on the residual, due to (B.2).

3.1 Three Lemmas for Three Sequences

We start with a recursion showing the evolution of the difference of iterates of the method. The explicit form is rather involved, however, we include the orders of the terms in the statement below.

Lemma 3.3 (See Lemma B.4 for the full expression)**.** *Under Assumptions 1.1, 1.2 and 3.1, for every $t \geq 1$ and for some $a > 2$, we have*

$$\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 \leq \Theta\left(1 - \frac{a}{t+1}\right)\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 + \Theta\left(\frac{1}{t^2}\right)\mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 + \Theta\left(\frac{1}{t^3}\right)(\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2 + \sigma^2).$$

On a high level, thanks to the precise form of the contraction on the first term on the right-hand side (see Lem B.4), one can see that if $\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2$ was bounded and if $\mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2$ had a rate $O(1/t)$, then inducting on this inequality would give $\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 = O(1/t^2)$. A priori, we do not have such a boundedness for $\mathbf{z}_t$ or a bound for $\mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2$, but, we will characterize the evolution of these three sequences jointly and then use induction. See App. B for further commentary.

We follow with the recursion for the variance reduced estimator $\mathbf{v}_t$ with its proof in App. B.1 which is rather classical and due to (Cutkosky & Orabona, 2019).

Lemma 3.4. *Under Assumption 1.1, 1.2 and 3.1, for every $t \geq 1$, the estimator in Alg. 2 satisfies*

$$\mathbb{E}_t\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 \leq (1-\alpha_t)^2\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 + 2\Lambda^2\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 + 2\alpha_t^2\big(\sigma^2 + B^2\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2\big).$$

We next have a recursion on the expected distance of the iterates to a solution, to show that the iterates are bounded in expectation, even when the problem domain is unbounded. Its proof is in App. B.2.

Lemma 3.5. *Let Assumption 1.1 hold. For every $t \geq 1$ we have that*

$$\|\mathbf{z}_{t+1} - \mathbf{z}^\star\|^2 \leq \left(1 - \frac{\beta_t}{2}\right)\|\mathbf{z}_t - \mathbf{z}^\star\|^2 + 2\beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2\eta_t^2}{\beta_t}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2.$$

3.2 Convergence Analysis

Proof sketch of Theorem 3.2. We use induction on the three estimates given in Lemmas 3.3, 3.4, and 3.5. This is carried out in Lemma B.1 and we obtain

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2 = O(1), \quad \mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 = O(1/t), \quad \mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t+1}\|^2 = O(1/t^2). \tag{3.1}$$

Moreover, expanding on the definition of the tangent residual, we obtain

$$\sqrt{\mathbb{E}\|G(\mathbf{z}_t) + \mathbf{s}_t\|^2} \le (L + \eta_{t-1}^{-1})\sqrt{\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2} + \left(\frac{\beta_{t-1}}{\eta_{t-1}}\right)\sqrt{\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2} + \sqrt{\mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2},$$

where the precise estimation appears in Appendix B. Since $\eta_t = \Theta(1/\sqrt{t})$ and $\beta_t = \Theta(1/t)$, plugging in the bounds from (3.1) gives the assertion. ■

4 Preliminary Numerical Results

Our numerical experiments evaluate two constrained problems to complement our theoretical results. We first consider the problem of matrix games given as $\min_{\mathbf{x}\in\Delta^m} \max_{\mathbf{y}\in\Delta^n} \mathbf{x}^\top A\mathbf{y}$, where Δ^d denotes the d-dimensional simplex. For this problem, we set $m = n = 50$ and set A by using the random payoff game of (Abe et al., 2025), replacing their addictive noise with row and column sampling for computing the unbiased oracles. We report the primal-dual gap. Next we consider the robust least squares problem studied by (Cai et al., 2022a): $\min_{\mathbf{x}\in\mathcal{X}} \max_{\mathbf{y}\in\mathcal{Y}} \frac{1}{2n}\|A\mathbf{x} - \mathbf{y}\|^2 - \frac{1.5}{2n}\|\mathbf{y} - \mathbf{b}\|^2$, where $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\|_\infty \le 1\}$ and $\mathcal{Y} = \{\mathbf{y} : \|\mathbf{y} - \mathbf{b}\|_\infty \le 1/2\}$. Following (Cai et al., 2022a), we use the same real dataset size of 21263×81 from (Hamidieh, 2018; Dua et al., 2017). We estimate the operator by uniformly sampling a data row and report the gradient mapping norm.

We compare Algorithm 1 and Algorithm 2 with RG and ROG of (Ito et al., 2026), PS-OGDA of (Zheng et al., 2026). All methods considered are anytime algorithms with last-iterate convergence guarantees. For each method, the hyperparameters are selected using three validations and then fixed for evaluation. The reported curves are averages over ten independent runs. We observe that Algorithm 1 performs comparable to or slightly better than the baselines eventually, while Algorithm 2 outperforms over the baselines on both cases. See Appendix E for hyperparameter choices, additional measurements, and an additional experiment with initialization sensitivity of our methods.

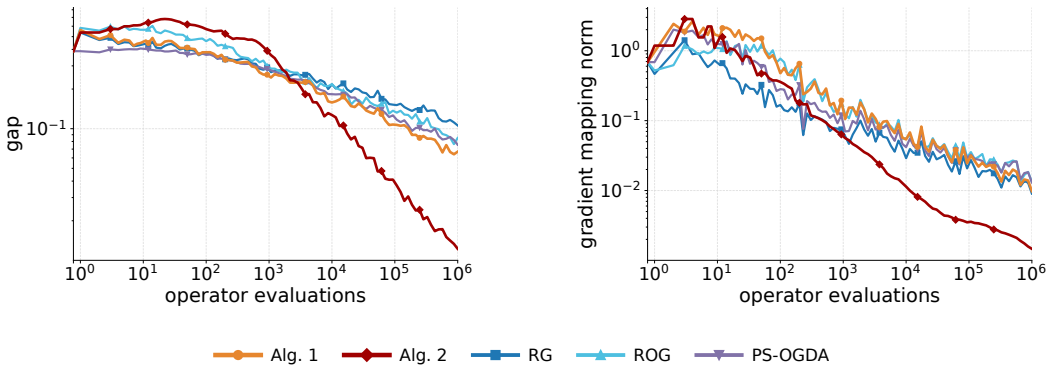

Figure 1: Left: matrix games. Right: robust least squares problem.

5 Conclusions

We studied single-loop, anytime algorithms and proved the rates of convergence $O(t^{-1/4})$ and $O(t^{-1/2})$, where the former is by using a single call to the unbiased oracle at every iteration with no variance reduction and the latter is by using two calls, with variance reduction. Directions for future work include handling nonmonotonicity, getting optimal dependence on non-dominant terms in convergence bounds, and improving regret bounds for the learning in games setting.

AI USE STATEMENT

The authors used generative AI tools for the straightforward but tedious numerical estimations used in Facts D.4 and D.5, in particular to obtain precise constants. The authors then verified, simplified, and rewritten all of these estimations. The authors set up the algorithm and the proof template for Section 3, and verified the orders of the terms in the analysis for η_t of order $1/t^{1/3}$; then, the authors asked these tools to check the orders of the terms when η_t has order $1/t^{1/2}$ to check feasibility for an extension. After this, the authors wrote the proof for Section 3. The authors used generative AI tools for coding for the numerical experiments.

REFERENCES

Kenshi Abe, Kaito Ariu, Mitsuki Sakamoto, and Atsushi Iwasaki. Adaptively perturbed mirror descent for learning in games. In *International Conference on Machine Learning*, 2024.

Kenshi Abe, Mitsuki Sakamoto, Kaito Ariu, and Atsushi Iwasaki. Boosting perturbed gradient ascent for last-iterate convergence in games. In *International Conference on Learning Representations*, volume 2025, pp. 93223–93253, 2025.

Ahmet Alacaoglu. How to make the gradient mapping small for constrained stochastic min-max problems and beyond. *arXiv preprint arXiv:2609.08380*, 2026.

Ahmet Alacaoglu and Jun-Hyun Kim. Solving stochastic variational inequalities without the bounded variance assumption. In *International Conference on Machine Learning*, 2026.

Ahmet Alacaoglu, Yura Malitsky, and Stephen J Wright. Towards weaker variance assumptions for stochastic optimization. *arXiv preprint arXiv:2504.09951*, 2025.

Waïss Azizian, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. The last-iterate convergence rate of optimistic mirror descent in stochastic variational inequalities. In *Conference on Learning Theory*, pp. 326–358. PMLR, 2021.

Xufeng Cai, Chaobing Song, Cristóbal Guzmán, and Jelena Diakonikolas. Stochastic halpern iteration with variance reduction for stochastic monotone inclusions. *Advances in Neural Information Processing Systems*, 35:24766–24779, 2022a.

Yang Cai and Weiqiang Zheng. Accelerated single-call methods for constrained min-max optimization. In *International Conference on Learning Representations*, 2023a.

Yang Cai and Weiqiang Zheng. Doubly optimal no-regret learning in monotone games. In *International Conference on Machine Learning*, pp. 3507–3524. PMLR, 2023b.

Yang Cai and Weiqiang Zheng. Last-iterate convergence of anchored gradient descent. *arXiv preprint arXiv:2604.12235*, 2026.

Yang Cai, Argyris Oikonomou, and Weiqiang Zheng. Finite-time last-iterate convergence for learning in multi-player games. *Advances in Neural Information Processing Systems*, 35:33904–33919, 2022b.

Yang Cai, Argyris Oikonomou, and Weiqiang Zheng. Accelerated algorithms for constrained nonconvex-nonconcave min-max optimization and comonotone inclusion. In *International Conference on Machine Learning*, 2024.

Yair Carmon, Yujia Jin, Aaron Sidford, and Kevin Tian. Variance reduction for matrix games. *Advances in Neural Information Processing Systems*, 32, 2019.

Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*, volume 1. Cambridge university press Cambridge, 2006.

Lesi Chen and Luo Luo. Near-optimal algorithms for making the gradient small in stochastic minimax optimization. *Journal of Machine Learning Research*, 25(387):1–44, 2024.

Zaiwei Chen and Eric Mazumdar. Last-iterate convergence for generalized frank-wolfe in monotone variational inequalities. *Advances in Neural Information Processing Systems*, 37:115440–115467, 2024.

Ashok Cutkosky and Francesco Orabona. Momentum-based variance reduction in non-convex SGD. *Advances in neural information processing systems*, 32, 2019.

Dheeru Dua, Casey Graff, et al. Uci machine learning repository, 2017. *URL http://archive. ics. uci. edu/ml*, 2017.

Francisco Facchinei and Jong-Shi Pang. *Finite-dimensional variational inequalities and complementarity problems*. Springer, 2003.

Dylan J Foster, Ayush Sekhari, Ohad Shamir, Nathan Srebro, Karthik Sridharan, and Blake Woodworth. The complexity of making the gradient small in stochastic convex optimization. In *Conference on Learning Theory*, pp. 1319–1345. PMLR, 2019.

Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. *Advances in neural information processing systems*, 33:20766–20778, 2020a.

Noah Golowich, Sarath Pattathil, Constantinos Daskalakis, and Asuman Ozdaglar. Last iterate is slower than averaged iterate in smooth convex-concave saddle point problems. In *Conference on Learning Theory*, pp. 1758–1784. PMLR, 2020b.

Eduard Gorbunov, Adrien Taylor, and Gauthier Gidel. Last-iterate convergence of optimistic gradient method for monotone variational inequalities. *Advances in neural information processing systems*, 35:21858–21870, 2022.

Benjamin Halpern. Fixed points of nonexpanding maps. *Bulletin of the American Mathematical Society*, 73(6):957–961, 1967.

Kam Hamidieh. A data-driven statistical model for predicting the critical temperature of a superconductor. *Computational Materials Science*, 154:346–354, 2018.

Yu-Guan Hsieh, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. On the convergence of single-call stochastic extra-gradient methods. *Advances in Neural Information Processing Systems*, 32, 2019.

Yu-Guan Hsieh, Kimon Antonakopoulos, Volkan Cevher, and Panayotis Mertikopoulos. No-regret learning in games with noisy feedback: Faster rates and adaptivity via learning rate separation. *Advances in Neural Information Processing Systems*, 35:6544–6556, 2022.

Shinji Ito, Taira Tsuchiya, Kaito Ariu, and Kenshi Abe. Last-iterate convergence of regularized gradient methods for stochastic monotone variational inequalities. In *International Conference on Machine Learning*, 2026.

Alfredo N Iusem, Alejandro Jofré, Roberto Imbuzeiro Oliveira, and Philip Thompson. Extragradient method with variance reduction for stochastic variational inequalities. *SIAM Journal on Optimization*, 27(2):686–724, 2017.

Yao Ji, Guanghui Lan, and Jason Zhu. Computation of strong solutions to stochastic variational inequalities. *arXiv preprint arXiv:2609.04188*, 2026.

Michael Jordan, Tianyi Lin, and Zhengyuan Zhou. Adaptive, doubly optimal no-regret learning in strongly monotone and exp-concave games with gradient feedback. *Operations Research*, 73(3): 1675–1702, 2025.

Galina M Korpelevich. The extragradient method for finding saddle points and other problems. *Matecon*, 12:747–756, 1976.

Georgios Kotsalis, Guanghui Lan, and Tianjiao Li. Simple and optimal methods for stochastic variational inequalities, i: operator extrapolation. *SIAM Journal on Optimization*, 32(3):2041–2073, 2022.

Dmitry Kovalev and Alexander Gasnikov. The first optimal algorithm for smooth and strongly-convex-strongly-concave minimax optimization. *Advances in Neural Information Processing Systems*, 35:14691–14703, 2022.

Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. *Advances in Neural Information Processing Systems*, 34:22588–22600, 2021.

Panayotis Mertikopoulos and Zhengyuan Zhou. Learning in games with continuous action sets and unknown payoff functions. *Mathematical Programming*, 173(1):465–507, 2019.

Arkadi Nemirovski, Anatoli Juditsky, Guanghui Lan, and Alexander Shapiro. Robust stochastic approximation approach to stochastic programming. *SIAM Journal on optimization*, 19(4):1574–1609, 2009.

Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2):319–344, 2007.

Gergely Neu and Nneka Okolo. Dealing with unbounded gradients in stochastic saddle-point optimization. In *International Conference on Machine Learning*, 2024.

Thomas Pethick, Olivier Fercoq, Puya Latafat, Panagiotis Patrinos, and Volkan Cevher. Solving stochastic weak minty variational inequalities without increasing batch size. In *International Conference on Learning Representations*, 2023a.

Thomas Pethick, Wanyun Xie, and Volkan Cevher. Stable nonconvex-nonconcave training via linear interpolation. *Advances in Neural Information Processing Systems*, 36:49830–49841, 2023b.

Leonid Denisovich Popov. A modification of the arrow-hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.

Sasha Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. *Advances in Neural Information Processing Systems*, 26, 2013.

Motahareh Sohrabi, Jianxin You, Simon Lacoste-Julien, Eduard Gorbunov, and Gauthier Gidel. Accelerated and stable convergence with anchored generalized optimistic method. In *International Conference on Machine Learning*, 2026.

Anja Surina, Arun Suggala, George Tsoukalas, Anton Kovsharov, Sergey Shirobokov, Francisco JR Ruiz, Pushmeet Kohli, and Swarat Chaudhuri. An improved last-iterate convergence rate for anchored gradient descent ascent. *arXiv preprint arXiv:2604.03782*, 2026.

Quoc Tran-Dinh and Nghia Nguyen-Trung. Unbiased and biased variance-reduced forward-reflected-backward splitting methods for stochastic composite inclusions. *arXiv preprint arXiv:2603.15576*, 2026.

TaeHo Yoon and Ernest K Ryu. Accelerated algorithms for smooth convex-concave minimax problems with $O(1/k^2)$ rate on squared gradient norm. In *International conference on machine learning*, pp. 12098–12109. PMLR, 2021.

Taoli Zheng, Jiajin Li, and Anthony Man-Cho So. Last-iterate convergence of single-loop stochastic methods for constrained convex-concave minimax problems. *arXiv preprint arXiv:2607.11056*, 2026.

A Analysis of Algorithm 1

Proof of Theorem 2.1. Recall that $\rho = 1/L$. By adding and subtracting $\mathcal{G}_\rho(\bar{\mathbf{z}}_t)$ and applying Young's inequality, we obtain

$$\begin{aligned}\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\|^2 &\leq 2\|\mathcal{G}_\rho(\bar{\mathbf{z}}_t)\|^2 + 2\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t) - \mathcal{G}_\rho(\bar{\mathbf{z}}_t)\|^2 \\ &\leq 2\|\mathcal{G}_\rho(\bar{\mathbf{z}}_t)\|^2 + 2(\rho^{-2} + L^2)\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2, \end{aligned} \tag{A.1}$$

where the last line is by (D.6). We bound the first term on the right-hand side. By Lemmas 2.6 and D.3, we get

$$\|\mathcal{G}_\rho(\bar{\mathbf{z}}_t)\| \leq \frac{5aH\|\mathbf{z}_0 - \mathbf{z}^\star\|}{2(t-1+a)^{1/4}}\left(\frac{3}{2} + \frac{L}{2H(t-1+a)^{3/4}}\right).$$

Taking the square of both sides and multiplying by 2 give us

$$2\|\mathcal{G}_\rho(\bar{\mathbf{z}}_t)\|^2 \leq \frac{25a^2H^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2}{2\sqrt{t-1+a}}\left(\frac{3}{2} + \frac{L}{2Ha^{3/4}}\right)^2,$$

where we used $t - 1 + a \geq a$.

For the second term on the right-hand side of (A.1), by Lemma 2.7, we have

$$2(\rho^{-2} + L^2)\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 \leq \frac{2(\rho^{-2} + L^2)}{H^2(a - \frac{1}{2})\sqrt{t-1+a}}\left(\sigma^2 + \frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2\right),$$

where we used $t + a \geq t - 1 + a$.

Combining the last two estimates in (A.1) yields

$$\begin{aligned}\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\|^2 \leq \frac{1}{\sqrt{t-1+a}}\Bigg[&\frac{25a^2H^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2}{8}\left(3 + \frac{L}{Ha^{3/4}}\right)^2 \\ &+ \frac{2(\rho^{-2} + L^2)}{H^2(a - \frac{1}{2})}\left(\sigma^2 + \frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2\right)\Bigg].\end{aligned}$$

Using $H \geq L$ and $a \geq 2$, absorbing numerical constants and Jensen's inequality, that is,

$$\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\| \leq \left(\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\|^2\right)^{1/2},$$

provides the assertion. ■

A.1 Proof of Lemma 2.5

Proof of Lemma 2.5. By the definition of the solution, we have $-G(\mathbf{z}^\star) \in \partial r(\mathbf{z}^\star)$ and $\mathbf{z}^\star = \mathrm{prox}_{\eta_t r}(\mathbf{z}^\star - \eta_t G(\mathbf{z}^\star))$. With this, nonexpansiveness of the proximal operator, and triangle inequality, we obtain

$$\begin{aligned}\|\bar{\mathbf{z}}_{t+1} - \mathbf{z}^\star\| &\leq \|(1-\beta_t)\bar{\mathbf{z}}_t + \beta_t\mathbf{z}_0 - \mathbf{z}^\star - \eta_t G(\bar{\mathbf{z}}_t) + \eta_t G(\mathbf{z}^\star)\| \\ &\leq \|(1-\beta_t)(\bar{\mathbf{z}}_t - \mathbf{z}^\star) - \eta_t(G(\bar{\mathbf{z}}_t) - G(\mathbf{z}^\star))\| + \beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|. \end{aligned} \tag{A.2}$$

For the first norm on the right-hand side, we use Lemma D.1 with $\mathbf{u} = \bar{\mathbf{z}}_t$, $\mathbf{s} = \mathbf{z}^\star$, $\alpha = (1-\beta_t)$, $\gamma = \eta_t$ to get

$$\|\bar{\mathbf{z}}_{t+1} - \mathbf{z}^\star\| \leq \sqrt{(1-\beta_t)^2 + \eta_t^2L^2}\,\|\bar{\mathbf{z}}_t - \mathbf{z}^\star\| + \beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|. \tag{A.3}$$

We first verify the induction base case $t = 1$. At the first step, $\bar{\mathbf{z}}_0 = \mathbf{z}_0$, $\beta_0 = 1$ and $\eta_0 = 1/(Ha^{3/4})$. By definition of $\bar{\mathbf{z}}_1$ and Lemma D.1 with $\mathbf{u} = \mathbf{z}_0$, $\mathbf{s} = \mathbf{z}^\star$, $\alpha = 1$, $\gamma = \eta_0$, we have

$$\|\bar{\mathbf{z}}_1 - \mathbf{z}^\star\|^2 \leq \|\mathbf{z}_0 - \mathbf{z}^\star - \eta_0(G(\mathbf{z}_0) - G(\mathbf{z}^\star))\|^2 \leq (1 + \eta_0^2L^2)\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \leq \frac{9}{4}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2.$$

Indeed, $\eta_0^2L^2 = L^2/(H^2a^{3/2}) \leq a^{-3/2} \leq 5/4$ since $H \geq L$, and $a \geq 2$.

We now prove the first bound in (2.3) by induction. The case $t = 0$ follows from $\bar{\mathbf{z}}_0 = \mathbf{z}_0$, and the calculation above proves the case $t = 1$. Suppose that, for some $t \geq 1$, $\|\bar{\mathbf{z}}_t - \mathbf{z}^\star\| \leq \frac{3}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|$. For $t \geq 1$, Fact D.4 gives $\sqrt{(1-\beta_t)^2 + \eta_t^2 L^2} \leq 1 - \frac{2}{3}\beta_t$. Using this and the inductive assumption in (A.3), we obtain

$$\begin{aligned}\|\bar{\mathbf{z}}_{t+1} - \mathbf{z}^\star\| &\leq \left(1 - \frac{2}{3}\beta_t\right)\|\bar{\mathbf{z}}_t - \mathbf{z}^\star\| + \beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\| \\ &\leq \left(1 - \frac{2}{3}\beta_t\right)\frac{3}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\| + \beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\| \\ &= \frac{3}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|.\end{aligned}$$

This completes the induction and proves the first bound in (2.3). Finally, using the triangle inequality and the first bound in (2.3) gives the second bound in (2.3). ■

A.2 Proof of Lemma 2.6

Proof of Lemma 2.6. This proof follows the structure of (Cai & Zheng, 2026) only by changing the parameter schedules.

By the definition in (2.2) and the definition of the proximal operator gives

$$\mathbf{s}_{t+1} = \frac{(1-\beta_t)\bar{\mathbf{z}}_t + \beta_t\mathbf{z}_0 - \eta_t G(\bar{\mathbf{z}}_t) - \bar{\mathbf{z}}_{t+1}}{\eta_t} \in \partial r(\bar{\mathbf{z}}_{t+1}). \tag{A.4}$$

In particular, $\bar{\mathbf{z}}_{t+1} = \mathrm{prox}_{\eta_{t+1}r}(\bar{\mathbf{z}}_{t+1} + \eta_{t+1}\mathbf{s}_{t+1}) \iff \bar{\mathbf{z}}_{t+1} + \eta_{t+1}\partial r(\bar{\mathbf{z}}_{t+1}) \ni \bar{\mathbf{z}}_{t+1} + \eta_{t+1}\mathbf{s}_{t+1}$. We compare this expression with the update $\bar{\mathbf{z}}_{t+2} = \mathrm{prox}_{\eta_{t+1}r}((1-\beta_{t+1})\bar{\mathbf{z}}_{t+1} + \beta_{t+1}\mathbf{z}_0 - \eta_{t+1}G(\bar{\mathbf{z}}_{t+1}))$. By Fact D.4, we have $\eta_{t+1}/\eta_t - \beta_{t+1} \geq 0$. Nonexpansiveness of the proximal operator and substitution of (A.4) give the first inequality below. Collecting the coefficients of $\bar{\mathbf{z}}_{t+1}$, $\bar{\mathbf{z}}_t$, and $\mathbf{z}_0$ gives the equality, while the last inequality follows from the triangle inequality.

$$\begin{aligned}&\|\bar{\mathbf{z}}_{t+2} - \bar{\mathbf{z}}_{t+1}\| \\ &\leq \left\| -\beta_{t+1}\bar{\mathbf{z}}_{t+1} + \beta_{t+1}\mathbf{z}_0 - \eta_{t+1}G(\bar{\mathbf{z}}_{t+1}) - \frac{\eta_{t+1}}{\eta_t}\big((1-\beta_t)\bar{\mathbf{z}}_t + \beta_t\bar{\mathbf{z}}_0 - \eta_t G(\bar{\mathbf{z}}_t) - \bar{\mathbf{z}}_{t+1}\big)\right\| \\ &= \left\| \left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)(\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t) - \eta_{t+1}\big(G(\bar{\mathbf{z}}_{t+1}) - G(\bar{\mathbf{z}}_t)\big) + \left(\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right)(\bar{\mathbf{z}}_0 - \bar{\mathbf{z}}_t)\right\| \\ &\leq \left\| \left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)(\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t) - \eta_{t+1}\big(G(\bar{\mathbf{z}}_{t+1}) - G(\bar{\mathbf{z}}_t)\big)\right\| + \left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right|\|\bar{\mathbf{z}}_0 - \bar{\mathbf{z}}_t\|.\end{aligned} \tag{A.5}$$

We next estimate the first norm on the right-hand side. Using Lemma D.1 with $\mathbf{u} = \bar{\mathbf{z}}_{t+1}$, $\mathbf{s} = \bar{\mathbf{z}}_t$, $\alpha = \frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}$, $\gamma = \eta_{t+1}$ yields

$$\begin{aligned}&\left\| \left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)(\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t) - \eta_{t+1}\big(G(\bar{\mathbf{z}}_{t+1}) - G(\bar{\mathbf{z}}_t)\big)\right\|^2 \\ &\leq \left[\left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)^2 + \eta_{t+1}^2 L^2\right]\|\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t\|^2.\end{aligned}$$

Taking square roots, substituting this estimate into (A.5), and using Lemma 2.5 gives

$$\begin{aligned}\|\bar{\mathbf{z}}_{t+2} - \bar{\mathbf{z}}_{t+1}\| \leq &\sqrt{\left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)^2 + \eta_{t+1}^2 L^2}\,\|\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t\| \\ &+ \frac{5}{2}\left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right|\|\mathbf{z}_0 - \mathbf{z}^\star\|.\end{aligned} \tag{A.6}$$

By fact D.4, we have

$$\left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)^2 + \eta_{t+1}^2 L^2 \le \left(1 - \frac{3}{2s}\right)^2, \quad \text{and} \quad \left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right| \le \frac{a}{4(t+a)(t+a+1)}.$$

Substituting these bounds into (A.6) gives

$$\|\bar{\mathbf{z}}_{t+2} - \bar{\mathbf{z}}_{t+1}\| \le \left(1 - \frac{3}{2(t+a+1)}\right)\|\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t\| + \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{8(t+a)(t+a+1)}. \tag{A.7}$$

For $t = 0$, since $\bar{\mathbf{z}}_1 = \operatorname{prox}_{\eta_0 r}(\mathbf{z}_0 - \eta_0 G(\mathbf{z}_0))$, the definition of the proximal operator gives

$$\langle \bar{\mathbf{z}}_1 - \mathbf{z}_0 + \eta_0 G(\mathbf{z}_0), \mathbf{z}^\star - \bar{\mathbf{z}}_1\rangle \ge \eta_0\left(r(\bar{\mathbf{z}}_1) - r(\mathbf{z}^\star)\right)$$

By $2\langle \mathbf{u}, \mathbf{s}\rangle = \|\mathbf{u}+\mathbf{s}\|^2 - \|\mathbf{u}\|^2 - \|\mathbf{s}\|^2$ and adding and subtracting $\eta_0\langle G(\mathbf{z}^\star), \mathbf{z}^\star - \bar{\mathbf{z}}_1\rangle$, we get

$$\begin{aligned}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 - \|\bar{\mathbf{z}}_1 - \mathbf{z}_0\|^2 - \|\bar{\mathbf{z}}_1 - \mathbf{z}^\star\|^2 + 2\eta_0\langle G(\mathbf{z}_0) - G(\mathbf{z}^\star), \mathbf{z}^\star - \bar{\mathbf{z}}_1\rangle \\ \ge 2\eta_0\left(r(\bar{\mathbf{z}}_1) - r(\mathbf{z}^\star) + \langle G(\mathbf{z}^\star), \bar{\mathbf{z}}_1 - \mathbf{z}^\star\rangle\right)\end{aligned}$$

The right-hand side is nonnegative by (1.2) and rearranging gives

$$\begin{aligned}\|\bar{\mathbf{z}}_1 - \mathbf{z}_0\|^2 &\le \|\mathbf{z}_0 - \mathbf{z}^\star\|^2 - \|\bar{\mathbf{z}}_1 - \mathbf{z}^\star\|^2 - 2\eta_0\langle \bar{\mathbf{z}}_1 - \mathbf{z}^\star, G(\mathbf{z}_0) - G(\mathbf{z}^\star)\rangle \\ &\le \|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \eta_0^2\|G(\mathbf{z}_0) - G(\mathbf{z}^\star)\|^2 \le (1 + a^{-3/2})\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \le \frac{25}{16}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2.\end{aligned}$$

Where the last two inequalities use $H \ge L$ and $a \ge 2$, respectively. Taking square roots in the preceding estimate gives $\|\bar{\mathbf{z}}_1 - \bar{\mathbf{z}}_0\| \le \frac{5}{4}\|\mathbf{z}_0 - \mathbf{z}^\star\|$, which proves (2.4) for $t = 0$. Suppose that, for some $t \ge 0$,

$$\|\bar{\mathbf{z}}_{t+1} - \bar{\mathbf{z}}_t\| \le \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t+a)}.$$

Substituting this bound into (A.7) gives

$$\begin{aligned}\|\bar{\mathbf{z}}_{t+2} - \bar{\mathbf{z}}_{t+1}\| &\le \left(1 - \frac{3}{2(t+a+1)}\right)\frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t+a)} + \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{8(t+a)(t+a+1)} \\ &= \left(1 - \frac{3}{2(t+a+1)} + \frac{1}{2(t+a+1)}\right)\frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t+a)} \\ &= \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t+a+1)}.\end{aligned}$$

This completes the induction and proves (2.4).

Finally, setting $t \leftarrow t - 1$ in (A.4) and rearranging gives

$$G(\bar{\mathbf{z}}_T) + \mathbf{s}_T = G(\bar{\mathbf{z}}_T) - G(\bar{\mathbf{z}}_{T-1}) + \frac{\beta_{T-1}(\mathbf{z}_0 - \bar{\mathbf{z}}_{T-1}) + \bar{\mathbf{z}}_{T-1} - \bar{\mathbf{z}}_T}{\eta_{T-1}}.$$

Therefore, the triangle inequality and Lipschitzness gives

$$\begin{aligned}\|G(\bar{\mathbf{z}}_t) + \mathbf{s}_t\| &\le L\|\bar{\mathbf{z}}_t - \bar{\mathbf{z}}_{t-1}\| + \frac{\beta_{t-1}\|\mathbf{z}_0 - \bar{\mathbf{z}}_{t-1}\| + \|\bar{\mathbf{z}}_t - \bar{\mathbf{z}}_{t-1}\|}{\eta_{t-1}} \\ &\le \frac{5aL\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t-1+a)} + H(t-1+a)^{3/4}\left(\frac{a}{t-1+a}\cdot\frac{5}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\| + \frac{5a\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t-1+a)}\right) \\ &= \frac{5aL\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t-1+a)} + \frac{15aH\|\mathbf{z}_0 - \mathbf{z}^\star\|}{4(t-1+a)^{1/4}} \\ &= \frac{5aH\|\mathbf{z}_0 - \mathbf{z}^\star\|}{2(t-1+a)^{1/4}}\left(\frac{3}{2} + \frac{L}{2H(t-1+a)^{3/4}}\right).\end{aligned}$$

The second inequality uses Lemma 2.5, (2.4), $\beta_{t-1} = a/(t-1+a)$, and $\eta_{t-1}^{-1} = H(t-1+a)^{3/4}$. Since $\mathbf{s}_t \in \partial r(\bar{\mathbf{z}}_t)$, this proves (2.5). ■

A.3 Proof of Lemma 2.7

Proof of Lemma 2.7. By the Young's inequality and Lemma 2.5, we have

$$\|\mathbf{z}_t - \mathbf{z}_0\|^2 \le 2\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + 2\|\mathbf{z}_0 - \bar{\mathbf{z}}_t\|^2 \le 2\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + \frac{25}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2. \tag{A.8}$$

Using the definitions of $\mathbf{z}_{t+1}$, $\bar{\mathbf{z}}_{t+1}$, and the nonexpansiveness of the proximal operator, we obtain

$$\begin{aligned}\|\mathbf{z}_{t+1} - \bar{\mathbf{z}}_{t+1}\|^2 &= \|\operatorname{prox}_{\eta_t r}\left((1-\beta_t)\mathbf{z}_t + \beta_t\mathbf{z}_0 - \eta_t\widetilde{G}(\mathbf{z}_t)\right) \\ &\quad - \operatorname{prox}_{\eta_t r}\left((1-\beta_t)\bar{\mathbf{z}}_t + \beta_t\mathbf{z}_0 - \eta_t G(\bar{\mathbf{z}}_t)\right)\|^2 \\ &\le \left\|(1-\beta_t)(\mathbf{z}_t - \bar{\mathbf{z}}_t) - \eta_t\big(\widetilde{G}(\mathbf{z}_t) - G(\bar{\mathbf{z}}_t)\big)\right\|^2.\end{aligned}$$

We add and subtract $G(\mathbf{z}_t)$ inside the norm and take conditional expectation. Since the stochastic oracle is unbiased, the noise cross term vanishes, yielding

$$\mathbb{E}_t\|\mathbf{z}_{t+1} - \bar{\mathbf{z}}_{t+1}\|^2 \le \left\|(1-\beta_t)(\mathbf{z}_t - \bar{\mathbf{z}}_t) - \eta_t\big(G(\mathbf{z}_t) - G(\bar{\mathbf{z}}_t)\big)\right\|^2 + \eta_t^2\mathbb{E}_t\|\widetilde{G}(\mathbf{z}_t) - G(\mathbf{z}_t)\|^2. \tag{A.9}$$

We use Lemma D.1 with $\mathbf{u} = \mathbf{z}_t$, $\mathbf{s} = \bar{\mathbf{z}}_t$, $\alpha = 1-\beta_t$, $\gamma = \eta_t$ to get

$$\left\|(1-\beta_t)(\mathbf{z}_t - \bar{\mathbf{z}}_t) - \eta_t(G(\mathbf{z}_t) - G(\bar{\mathbf{z}}_t))\right\| \le \sqrt{(1-\beta_t)^2 + \eta_t^2 L^2}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|.$$

Next, Assumption 1.2 and (A.8) bound the last term on the right-hand side as

$$\mathbb{E}_t\|\widetilde{G}(\mathbf{z}_t) - G(\mathbf{z}_t)\|^2 \le B^2\|\mathbf{z}_t - \mathbf{z}_0\|^2 + \sigma^2 \le 2B^2\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + \frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \sigma^2. \tag{A.10}$$

Substituting these two bounds into (A.9) gives

$$\mathbb{E}_t\|\mathbf{z}_{t+1} - \bar{\mathbf{z}}_{t+1}\|^2 \le \left((1-\beta_t)^2 + \eta_t^2(L^2 + 2B^2)\right)\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + \eta_t^2\left(\frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \sigma^2\right). \tag{A.11}$$

For the rest of the proof let us set

$$V_\star = \frac{25}{2}B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \sigma^2. \tag{A.12}$$

For $t \ge 1$, Fact D.4 gives

$$(1-\beta_t)^2 + \eta_t^2(L^2 + 2B^2) \le 1 - \frac{a}{t+a}.$$

Taking expectation in (A.11) and using the bound together with (A.12) yields

$$\mathbb{E}\|\mathbf{z}_{t+1} - \bar{\mathbf{z}}_{t+1}\|^2 \le \left(1 - \frac{a}{t+a}\right)\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + \frac{V_\star}{H^2(t+a)^{3/2}}, \qquad t \ge 1. \tag{A.13}$$

At $t = 0$, $\mathbf{z}_0 = \bar{\mathbf{z}}_0$. Substituting this into (A.9) and using $\eta_0^2 = 1/(H^2a^{3/2})$ gives

$$\mathbb{E}\|\mathbf{z}_1 - \bar{\mathbf{z}}_1\|^2 \le \eta_0^2\mathbb{E}\|\widetilde{G}(\mathbf{z}_0) - G(\mathbf{z}_0)\|^2 \le \eta_0^2 V_\star = \frac{V_\star}{H^2a^{3/2}},$$

where the second inequality used (A.10) and (A.12). Thus, (A.13) also holds at $t = 0$.

We now prove (2.6) by induction. The case $t = 0$ follows from $\mathbf{z}_0 = \bar{\mathbf{z}}_0$. Suppose that, for some $t \ge 0$,

$$\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 \le \frac{V_\star}{H^2(a - \frac{1}{2})\sqrt{t+a}}.$$

Since $1 - \frac{a}{t+a} \geq 0$, substituting the inductive assumption into (A.13) gives

$$\begin{aligned}\mathbb{E}\|\mathbf{z}_{t+1} - \bar{\mathbf{z}}_{t+1}\|^2 &\leq \left(1 - \frac{a}{t+a}\right) \frac{V_\star}{H^2(a-\frac{1}{2})\sqrt{t+a}} + \frac{V_\star}{H^2(t+a)^{3/2}} \\ &= \frac{V_\star}{H^2(a-\frac{1}{2})} \left(\frac{1}{\sqrt{t+a}} - \frac{1}{2(t+a)^{3/2}}\right) \\ &\leq \frac{V_\star}{H^2(a-\frac{1}{2})\sqrt{t+a+1}}.\end{aligned}$$

The last inequality follows from convexity of $s \mapsto s^{-1/2}$ for $s = t + a$, which gives $(s+1)^{-1/2} \geq s^{-1/2} - \frac{1}{2}s^{-3/2}$. This completes the induction and proves (2.6) after using (A.12).

Finally, taking expectations in (A.8) and using (2.6) gives

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2 \leq 2\mathbb{E}\|\mathbf{z}_t - \bar{\mathbf{z}}_t\|^2 + \frac{25}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \leq \frac{25}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2V_\star}{H^2(a-\frac{1}{2})\sqrt{t+a}},$$

which gives (2.6) after using (A.12). ■

B Analysis for Algorithm 2

Let us remark that (Alacaoglu et al., 2025) analyzed a similar method by using recursions depending on $\mathbf{z}_{t+1}, \mathbf{z}_t, \mathbf{z}^\star, \mathbf{v}_t - G(\mathbf{z}_t)$, but by combining the recursions with a single Lyapunov-type function and summing the resulting inequality, using the β_t anchoring to handle Assumption 1.2. The analysis framework used in this earlier work then seems to miss the benefit of anchoring with β_t to improve the rate and to get a guarantee on the last iterate. Instead, the analysis we use gets further mileage from anchoring with the induction argument, inspired by (Cai & Zheng, 2026; Surina et al., 2026). This leads to the double benefit from anchoring, described in Section 1.

Proof of Theorem 3.2. By the update in Algorithm 2 and the prox-inequality in (1.8), for every $t \geq 1$, we have

$$\mathbf{s}_t = \frac{(1-\beta_{t-1})\mathbf{z}_{t-1} + \beta_{t-1}\mathbf{z}_0 - \mathbf{z}_t}{\eta_{t-1}} - \mathbf{v}_{t-1} \in \partial r(\mathbf{z}_t). \tag{B.1}$$

The tangent residual satisfies

$$\begin{aligned}
\operatorname{dist}\big(0, G(\mathbf{z}_t) + \partial r(\mathbf{z}_t)\big) &\leq \|G(\mathbf{z}_t) + \mathbf{s}_t\| \\
&= \left\|G(\mathbf{z}_t) + \frac{(1-\beta_{t-1})\mathbf{z}_{t-1} + \beta_{t-1}\mathbf{z}_0 - \mathbf{z}_t}{\eta_{t-1}} - \mathbf{v}_{t-1}\right\| \\
&\leq \|G(\mathbf{z}_t) - G(\mathbf{z}_{t-1})\| + \frac{1}{\eta_{t-1}}\|\mathbf{z}_{t-1} - \mathbf{z}_t\| + \frac{\beta_{t-1}}{\eta_{t-1}}\|\mathbf{z}_0 - \mathbf{z}_{t-1}\| \\
&\quad + \|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\| \\
&\leq (L + \eta_{t-1}^{-1})\|\mathbf{z}_t - \mathbf{z}_{t-1}\| + \frac{\beta_{t-1}}{\eta_{t-1}}\|\mathbf{z}_{t-1} - \mathbf{z}_0\| + \|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|.
\end{aligned} \tag{B.2}$$

The equality follows from (B.1). The second and last inequalities follow from the triangle inequality and Lipschitzness, respectively. Squaring (B.2), taking expectations and applying Cauchy-Schwarz to the cross terms and completing the square, gives

$$\begin{aligned}
\sqrt{\mathbb{E}[\operatorname{res}(\mathbf{z}_t)^2]} &\leq (L + \eta_{t-1}^{-1})\sqrt{\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2} \\
&\quad + \frac{\beta_{t-1}}{\eta_{t-1}}\sqrt{\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2} + \sqrt{\mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2}.
\end{aligned}$$

By the parameter definitions, we have $\eta_{t-1} = 1/(H\sqrt{t})$ and $\beta_{t-1}/\eta_{t-1} = H/\sqrt{t}$. Substituting (B.3) and (B.4) from Lemma B.1 therefore give us

$$\begin{aligned}
\sqrt{\mathbb{E}[\operatorname{res}(\mathbf{z}_t)^2]} &\leq \frac{16(L + H\sqrt{t})\sqrt{C}}{t+1} + \frac{2H\sqrt{C}}{\sqrt{t}} + \frac{\sqrt{(1024\Lambda^2 + 4 + 16B^2)C}}{\sqrt{t}} \\
&\leq \frac{\big(16L + 18H + \sqrt{1024\Lambda^2 + 4 + 16B^2}\big)\sqrt{C}}{\sqrt{t}} \\
&\leq \frac{27H\sqrt{C}}{\sqrt{t}}.
\end{aligned}$$

The second inequality uses $t + 1 \geq t \geq \sqrt{t}$ for $t \geq 1$. The last inequality follows the definition of H, which gives $L \leq H/2$ and $\sqrt{1024\Lambda^2 + 4 + 16B^2} \leq H/4$. Thus, we have $16L + 18H + \sqrt{1024\Lambda^2 + 4 + 16B^2} \leq 27H$.

Applying the definition of C finishes the proof. ■

We now give the details of the analysis of Algorithm 2. It remains to control the three term, $\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_\star\|$, $\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|$ and $\mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|$. The following lemma gives the bounds used in the proof of Theorem 3.2.

Lemma B.1. *Let Assumptions 1.1, 1.2, and 3.1 hold. Then, for every $t \geq 0$,*

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}^\star\|^2 \leq C, \quad \mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2 \leq 4C, \quad \mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 \leq \frac{(1024\Lambda^2 + 4 + 16B^2)C}{t+1}, \tag{B.3}$$

and, for every $t \geq 1$,

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 \leq \frac{256C}{(t+1)^2}, \tag{B.4}$$

where we define C as $C = \max\left\{\sigma^2, 16\|\mathbf{z}_0 - \mathbf{z}^\star\|^2\right\}$.

Proof of Lemma B.1. By the definition of C, we have $C \geq \sigma^2$, $\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \leq C/16$ and

$$\mathbb{E}\|\mathbf{v}_0 - G(\mathbf{z}_0)\|^2 \leq \sigma^2 \leq C \leq (1024\Lambda^2 + 4 + 16B^2)C,$$

where the first inequality follows from Assumption 1.2 with definition of $\mathbf{v}_0 = \widetilde{G}(\mathbf{z}_0)$, while the last inequality follows from $1024\Lambda^2 + 4 + 16B^2 \geq 4$.

We first verify the base case at $t = 1$. Together with prox-inequality of $\mathbf{z}_1 = \mathrm{prox}_{\eta_0 r}(\mathbf{z}_0 - \eta_0\mathbf{v}_0)$ and $\mathbf{z}^\star = \mathrm{prox}_{\eta_t r}(\mathbf{z}^\star - \eta_t G(\mathbf{z}^\star))$, using nonexpansiveness of proximal operator gives,

$$\begin{aligned}
\mathbb{E}\|\mathbf{z}_1 - \mathbf{z}^\star\|^2 &\leq \mathbb{E}\Big\|\mathbf{z}_0 - \frac{1}{H}\mathbf{v}_0 - \mathbf{z}^\star + \eta_0 G(\mathbf{z}^\star)\Big\|^2 \\
&\leq 2\Big\|\mathbf{z}_0 - \mathbf{z}^\star - \frac{1}{H}\big(G(\mathbf{z}_0) - G(\mathbf{z}^\star)\big)\Big\|^2 + \frac{2}{H^2}\mathbb{E}\|\mathbf{v}_0 - G(\mathbf{z}_0)\|^2 \\
&\leq 2\left(1 + \frac{L^2}{H^2}\right)\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2}{H^2}\mathbb{E}\|\mathbf{v}_0 - G(\mathbf{z}_0)\|^2 \\
&\leq 2\left(1 + \frac{1}{4}\right)\frac{C}{16} + \frac{C}{8} = \frac{9C}{32} \leq C.
\end{aligned}$$

Here, the second inequality follows from Young's inequality, while the third inequality can be obtained from Lemma D.1 with $\alpha = 1, \gamma = 1/H$, $\mathbf{u} = \mathbf{z}_0$ and $\mathbf{s} = \mathbf{z}^\star$. The last inequality uses the definition of H, which implies $L^2/H^2 \leq 1/4$ and $(1024\Lambda^2 + 4 + 16B^2)/H^2 \leq 1/16$.

Using Young's inequality implies

$$\begin{aligned}
\mathbb{E}\|\mathbf{z}_1 - \mathbf{z}_0\|^2 &\leq 2\mathbb{E}\|\mathbf{z}_1 - \mathbf{z}^\star\|^2 + 2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \\
&\leq 2C + \frac{C}{8} \leq 64C = \frac{256C}{(1+1)^2}.
\end{aligned}$$

These bounds finish the base case of induction at $t = 1$.

Now suppose for some $t \geq 1$, that

$$\begin{aligned}
&\mathbb{E}\|\mathbf{z}_t - \mathbf{z}^\star\|^2 \leq C, \quad \mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}^\star\|^2 \leq C, \quad \mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 \leq \frac{256C}{(t+1)^2}, \\
&\qquad \text{and} \quad \mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 \leq \frac{(1024\Lambda^2 + 4 + 16B^2)C}{t}.
\end{aligned} \tag{B.5}$$

We first estimate the second term in (B.3). By Young's inequality, we have

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2 \leq 2\mathbb{E}\|\mathbf{z}_t - \mathbf{z}^\star\|^2 + 2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \leq 2C + \frac{C}{8} \leq 4C, \tag{B.6}$$

where the second inequality follows the property in (B.5) and definition of C. The same argument can be applied to $\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2 \leq 4C$.

We next estimate the last term in (B.3). We recall the recursion from Lemma 3.4:

$$\begin{aligned}
\mathbb{E}_t\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 \leq (1-\alpha_t)^2\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 + 2\Lambda^2\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 \\
+ 2\alpha_t^2\big(\sigma^2 + B^2\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2\big).
\end{aligned}$$

Taking full expectations and using (B.5) and (B.6) gives,

$$\begin{aligned}\mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 &\leq \left(1 - \frac{1}{t+1}\right)^2 \frac{(1024\Lambda^2 + 4 + 16B^2)C}{t} \\ &\quad + \frac{512\Lambda^2 C}{(t+1)^2} + \frac{2(C + 4B^2C)}{(t+1)^2} \\ &= \frac{(1024\Lambda^2 + 4 + 16B^2)C}{t+1} - \frac{(1024\Lambda^2 + 4 + 16B^2)C}{2(t+1)^2}, \end{aligned} \tag{B.7}$$

where the equality follows from Fact D.5.

Next, we verify first term in (B.3). Taking expectations in Lemma 3.5 give us

$$\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}^\star\|^2 \leq \left(1 - \frac{\beta_t}{2}\right)\mathbb{E}\|\mathbf{z}_t - \mathbf{z}^\star\|^2 + 2\beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2\eta_t^2}{\beta_t}\mathbb{E}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2. \tag{B.8}$$

Using (B.7) and definition of C which implies $\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 \leq C/16$, and $\eta_t^2/\beta_t = 1/H^2$,

$$\begin{aligned}\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}^\star\|^2 &\leq \left(1 - \frac{1}{2(t+1)}\right)C + \frac{C}{8(t+1)} + \frac{2(1024\Lambda^2 + 4 + 16B^2)C}{H^2(t+1)} \\ &\leq \left(1 - \frac{1}{2(t+1)} + \frac{1}{8(t+1)} + \frac{1}{8(t+1)}\right)C \\ &= \left(1 - \frac{1}{4(t+1)}\right)C \leq C. \end{aligned} \tag{B.9}$$

The second inequality used $(1024\Lambda^2 + 4 + 16B^2)/H^2 \leq 1/16$.

Finally, we estimate (B.4). Bringing the recursion in Lemma 3.3, we have

$$\begin{aligned}\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 &\leq \left(1 + \frac{1}{4(t+1)}\right)\left(\left(\frac{\eta_t}{\eta_{t-1}} - \beta_t\right)^2 + \eta_t^2 L^2\right)\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 \\ &\quad + 10(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1} - \beta_t\right)^2 \mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2 \\ &\quad + 10(t+1)\eta_t^2\alpha_t^2\mathbb{E}\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 \\ &\quad + \eta_t^2\Big(2\Lambda^2\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 + 2\alpha_t^2\big(\sigma^2 + B^2\mathbb{E}\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2\big)\Big). \end{aligned}$$

Applying (B.5), (B.6), and definitions of η_t and β_t with $\sigma^2 \leq C$, together with Fact D.5 gives,

$$\begin{aligned}\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 &\leq \frac{256C}{(t+1)^2}\left(1 + \frac{1}{4(t+1)}\right)\left(\left(1 - \frac{3}{2(t+1)}\right)^2 + \frac{1}{4(t+1)}\right) \\ &\quad + \left(40 + \frac{41(1024\Lambda^2 + 4 + 16B^2)}{2H^2}\right)\frac{C}{(t+1)^3} \\ &\leq \frac{256C}{(t+1)^2}\left(1 + \frac{1}{4(t+1)}\right)\left(\left(1 - \frac{3}{2(t+1)}\right)^2 + \frac{1}{4(t+1)}\right) + \frac{64C}{(t+1)^3}. \end{aligned}$$

Moreover, applying Fact D.5(5) give us,

$$\mathbb{E}\|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2 \leq \frac{256C}{(t+2)^2}. \tag{B.10}$$

This completes the induction and proves (B.3) and (B.4). ■

B.1 Proof of Lemma 3.4

The following lemma and proof are using the ideas from the work (Cutkosky & Orabona, 2019) that introduced the STORM estimator. The lemma with the variance required by Assumption 2 appears in the work (Alacaoglu et al., 2025).

Lemma B.2. *Under Assumption 1.1, 1.2, and 3.1, for every $t \geq 1$, the estimator in algorithm 2 satisfies*

$$\mathbb{E}_t\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 \leq (1-\alpha_t)^2\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2 + 2\Lambda^2\|\mathbf{z}_t - \mathbf{z}_{t-1}\|^2 + 2\alpha_t^2\big(\sigma^2 + B^2\|\mathbf{z}_{t-1} - \mathbf{z}_0\|^2\big).$$

Proof of Lemma 3.4. Throughout the proof, we use the notations

$$\widetilde{G}(\mathbf{z}_t) = \widetilde{G}(\mathbf{z}_t, \xi_t) \quad \text{and} \quad \widetilde{G}(\mathbf{z}_{t-1}) = \widetilde{G}(\mathbf{z}_{t-1}, \xi_t).$$

By the estimator update, we have

$$\begin{aligned}\mathbf{v}_t - G(\mathbf{z}_t) &= \widetilde{G}(\mathbf{z}_t) + (1-\alpha_t)\big(\mathbf{v}_{t-1} - \widetilde{G}(\mathbf{z}_{t-1})\big) - G(\mathbf{z}_t)\\ &= (1-\alpha_t)\big(\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\big) + (1-\alpha_t)\big(G(\mathbf{z}_{t-1}) - \widetilde{G}(\mathbf{z}_{t-1})\big)\\ &\quad + \widetilde{G}(\mathbf{z}_t) - G(\mathbf{z}_t).\end{aligned}$$

Since the last two terms on the right-hand side is 0 under conditional expectation and since $\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})$ is measurable under the conditioning of $\mathbb{E}_t$,we have

$$\mathbb{E}_t\langle \mathbf{v}_{t-1} - G(\mathbf{z}_{t-1}), \widetilde{G}(\mathbf{z}_t) - G(\mathbf{z}_t) + (1-\alpha_t)(G(\mathbf{z}_{t-1}) - \widetilde{G}(\mathbf{z}_{t-1}))\rangle = 0.$$

Thus, expanding the squared norm and taking expectations give

$$\begin{aligned}\mathbb{E}_t\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2 &= (1-\alpha_t)^2\|\mathbf{v}_{t-1} - G(\mathbf{z}_{t-1})\|^2\\ &\quad + \mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t) - G(\mathbf{z}_t) + (1-\alpha_t)(G(\mathbf{z}_{t-1}) - \widetilde{G}(\mathbf{z}_{t-1})\right\|^2.\end{aligned}$$

Applying Lemma B.5 to last term on right-hand gives the assertion. ∎

B.2 Proof of Lemma 3.5

Lemma B.3 (Restatement of Lemma 3.5)**.** *Let Assumption 1.1 hold. For every $t \geq 0$ we have that*

$$\|\mathbf{z}_{t+1} - \mathbf{z}^\star\|^2 \leq \left(1 - \frac{\beta_t}{2}\right)\|\mathbf{z}_t - \mathbf{z}^\star\|^2 + 2\beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2\eta_t^2}{\beta_t}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2.$$

Proof. By the definition of the solution, we have $-G(\mathbf{z}^\star) \in \partial r(\mathbf{z}^\star)$ and $\mathbf{z}^\star = \operatorname{prox}_{\eta_t r}(\mathbf{z}^\star - \eta_t G(\mathbf{z}^\star))$. Comparing with the update in algorithm 2, nonexpansiveness of the proximal operator gives

$$\begin{aligned}\|\mathbf{z}_{t+1} - \mathbf{z}^\star\| &\leq \|(1-\beta_t)\mathbf{z}_t + \beta_t\mathbf{z}_0 - \eta_t\mathbf{v}_t - (\mathbf{z}^\star - \eta_t G(\mathbf{z}^\star))\|\\ &= \|(1-\beta_t)(\mathbf{z}_t - \mathbf{z}^\star) - \eta_t\big(G(\mathbf{z}_t) - G(\mathbf{z}^\star)\big) + \beta_t(\mathbf{z}_0 - \mathbf{z}^\star) - \eta_t\big(\mathbf{v}_t - G(\mathbf{z}_t)\big)\|.\end{aligned}$$

Taking the square of this inequality and using Young's inequality with parameter $\beta_t/(1-\beta_t)$ gives

$$\begin{aligned}\|\mathbf{z}_{t+1} - \mathbf{z}^\star\|^2 &\leq \frac{1}{1-\beta_t}\|(1-\beta_t)(\mathbf{z}_t - \mathbf{z}^\star) - \eta_t\big(G(\mathbf{z}_t) - G(\mathbf{z}^\star)\big)\|^2\\ &\quad + \frac{1}{\beta_t}\big\|\beta_t(\mathbf{z}_0 - \mathbf{z}^\star) - \eta_t\big(\mathbf{v}_t - G(\mathbf{z}_t)\big)\big\|^2. \qquad \text{(B.11)}\end{aligned}$$

We first estimate the first term on the right-hand side, using Lemma D.1 with $\alpha = 1-\beta_t$, $\gamma = \eta_t$, $\mathbf{u} = \mathbf{z}_t$ and $\mathbf{s} = \mathbf{z}^\star$ yield

$$\begin{aligned}\frac{1}{1-\beta_t}\|(1-\beta_t)(\mathbf{z}_t - \mathbf{z}^\star) - \eta_t\big(G(\mathbf{z}_t) - G(\mathbf{z}^\star)\big)\|^2 &\leq \left(1 - \beta_t + \frac{\eta_t^2 L^2}{1-\beta_t}\right)\|\mathbf{z}_t - \mathbf{z}^\star\|^2\\ &\leq \left(1 - \frac{\beta_t}{2}\right)\|\mathbf{z}_t - \mathbf{z}^\star\|^2,\end{aligned}$$

where the second inequality follows from Fact D.5.

For the second term on the right-hand side of (B.11), Young's inequality gives

$$\frac{1}{\beta_t}\big\|\beta_t(\mathbf{z}_0 - \mathbf{z}^\star) - \eta_t\big(\mathbf{v}_t - G(\mathbf{z}_t)\big)\big\|^2 \leq 2\beta_t\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2\eta_t^2}{\beta_t}\|\mathbf{v}_t - G(\mathbf{z}_t)\|^2.$$

Combining these two estimates with (B.11) proves the claim. ∎

B.3 Full Expression of Lemma 3.3 and Its Proof

Lemma B.4. *Under Assumptions 1.1, 1.2 and 3.1, for every $t \geq 1$, we have*

$$\begin{aligned}\mathbb{E}\|\mathbf{z}_{t+1}-\mathbf{z}_t\|^2 \leq \left(1+\frac{1}{4(t+1)}\right)\left(\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)^2+\eta_t^2L^2\right)\mathbb{E}\|\mathbf{z}_t-\mathbf{z}_{t-1}\|^2 \\ +10(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)^2\mathbb{E}\|\mathbf{z}_{t-1}-\mathbf{z}_0\|^2 \\ +10(t+1)\eta_t^2\alpha_t^2\mathbb{E}\|\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\|^2 \\ +\eta_t^2\Big(2\Lambda^2\mathbb{E}\|\mathbf{z}_t-\mathbf{z}_{t-1}\|^2+2\alpha_t^2\big(\sigma^2+B^2\mathbb{E}\|\mathbf{z}_{t-1}-\mathbf{z}_0\|^2\big)\Big).\end{aligned}$$

Proof. Throughout the proof, we use the notations

$$\widetilde{G}(\mathbf{z}_t)=\widetilde{G}(\mathbf{z}_t,\xi_t) \quad \text{and} \quad \widetilde{G}(\mathbf{z}_{t-1})=\widetilde{G}(\mathbf{z}_{t-1},\xi_t).$$

By the previous update and the prox-inequality in (1.8), we have

$$\mathbf{s}_t=\frac{(1-\beta_{t-1})\mathbf{z}_{t-1}+\beta_{t-1}\mathbf{z}_0-\mathbf{z}_t}{\eta_{t-1}}-\mathbf{v}_{t-1}\in\partial r(\mathbf{z}_t).$$

Consequently, with stepsize η_t, we have $\mathbf{z}_t=\operatorname{prox}_{\eta_t r}(\mathbf{z}_t+\eta_t\mathbf{s}_t)$. Comparing this with the update of $\mathbf{z}_{t+1}$ and using nonexansiveness of the proximal operator gives

$$\begin{aligned}\|\mathbf{z}_{t+1}-\mathbf{z}_t\| &\leq \|(1-\beta_t)\mathbf{z}_t+\beta_t\mathbf{z}_0-\eta_t\mathbf{v}_t-(\mathbf{z}_t+\eta_t\mathbf{s}_t)\| \\ &= \left\|\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)(\mathbf{z}_t-\mathbf{z}_{t-1})\right. \\ &\quad \left.+\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)(\mathbf{z}_{t-1}-\mathbf{z}_0)-\eta_t(\mathbf{v}_t-\mathbf{v}_{t-1})\right\|. \qquad \text{(B.12)}\end{aligned}$$

Here the last equality follows from the definition of $\mathbf{s}_t$.

We first rewrite the last term inside the norm. Using the definition of estimator $\mathbf{v}_t$ in Algorithm 2, we have,

$$\begin{aligned}\mathbf{v}_t-\mathbf{v}_{t-1} &= \widetilde{G}(\mathbf{z}_t)-\widetilde{G}(\mathbf{z}_{t-1})-\alpha_t(\mathbf{v}_{t-1}-\widetilde{G}(\mathbf{z}_{t-1})) \\ &= \widetilde{G}(\mathbf{z}_t)-\widetilde{G}(\mathbf{z}_{t-1})-\alpha_t(\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1}))+\alpha_t(\widetilde{G}(\mathbf{z}_{t-1},\xi_t)-G(\mathbf{z}_{t-1})).\end{aligned}$$

Substituting this into (B.12) yields

$$\begin{aligned}\|\mathbf{z}_{t+1}-\mathbf{z}_t\| \leq \left\|\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)(\mathbf{z}_t-\mathbf{z}_{t-1})-\eta_t\big(G(\mathbf{z}_t)-G(\mathbf{z}_{t-1})\big)\right. \\ +\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)(\mathbf{z}_{t-1}-\mathbf{z}_0)+\eta_t\alpha_t\big(\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\big) \\ \left.-\eta_t\left(\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right)\right\|.\end{aligned}$$

Because $\mathbb{E}_t\left[\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right]=0$ and because the terms in the first and second lines of the right-hand side in the display equation above are measurable under the conditioning of $\mathbb{E}_t$, we have

$$\begin{aligned}\mathbb{E}_t\|\mathbf{z}_{t+1}-\mathbf{z}_t\|^2 \leq \left\|\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)(\mathbf{z}_t-\mathbf{z}_{t-1})-\eta_t\big(G(\mathbf{z}_t)-G(\mathbf{z}_{t-1})\big)\right. \\ \left.+\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)(\mathbf{z}_{t-1}-\mathbf{z}_0)+\eta_t\alpha_t\big(\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\big)\right\|^2 \\ +\eta_t^2\mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right\|^2.\end{aligned}$$

We now apply Young's inequality to the first term on the right-hand side with the parameter $1/(4(t+1))$ to retain a sufficient contraction.

$$\begin{aligned}\mathbb{E}_t\|\mathbf{z}_{t+1}-\mathbf{z}_t\|^2 \le \left(1+\frac{1}{4(t+1)}\right)\left\|\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)(\mathbf{z}_t-\mathbf{z}_{t-1})-\eta_t\big(G(\mathbf{z}_t)-G(\mathbf{z}_{t-1})\big)\right\|^2 \\ +(1+4(t+1))\left\|\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)(\mathbf{z}_{t-1}-\mathbf{z}_0)+\eta_t\alpha_t\big(\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\big)\right\|^2 \\ +\eta_t^2\mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right\|^2. \qquad \text{(B.13)}\end{aligned}$$

For the first term on the right-hand side, using Lemma D.1 with $\alpha=\eta_t/\eta_{t-1}-\beta_t\ge 0$ and $\gamma=\eta_t$ gives

$$\begin{aligned}\left(1+\frac{1}{4(t+1)}\right)\left\|\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)(\mathbf{z}_t-\mathbf{z}_{t-1})-\eta_t\big(G(\mathbf{z}_t)-G(\mathbf{z}_{t-1})\big)\right\|^2 \\ \le\left(1+\frac{1}{4(t+1)}\right)\left(\left(\frac{\eta_t}{\eta_{t-1}}-\beta_t\right)^2+\eta_t^2L^2\right)\|\mathbf{z}_t-\mathbf{z}_{t-1}\|^2. \qquad \text{(B.14)}\end{aligned}$$

We next estimate second term on the right-hand side of (B.13), using Young's inequality with the fact that $2(1+4(t+1))\le 10(t+1)$ for $t\ge 0$.

$$\begin{aligned}(1+4(t+1))\left\|\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)(\mathbf{z}_{t-1}-\mathbf{z}_0)+\eta_t\alpha_t\big(\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\big)\right\|^2 \\ \le 10(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)^2\|\mathbf{z}_{t-1}-\mathbf{z}_0\|^2+10(t+1)\eta_t^2\alpha_t^2\|\mathbf{v}_{t-1}-G(\mathbf{z}_{t-1})\|^2. \qquad \text{(B.15)}\end{aligned}$$

Applying Lemma B.5 to the last term in equation (B.13) and substituting (B.14) and (B.15) into equation (B.13), followed by taking expectation, proves the claim. ■

The following lemma is a simple consequence of Assumptions 1.2 and 3.1 and is given as a separate result since it is repeated in our proofs.

Lemma B.5. *Suppose Assumptions 1.1, 1.2, and 3.1 hold. Then,*

$$\begin{aligned}\mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right\|^2 \le 2\Lambda^2\|\mathbf{z}_t-\mathbf{z}_{t-1}\|^2 \\ +2\alpha_t^2\big(\sigma^2+B^2\|\mathbf{z}_{t-1}-\mathbf{z}_0\|^2\big).\end{aligned}$$

Proof. Using Young's inequality gives

$$\begin{aligned}&\mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)+(1-\alpha_t)\big(G(\mathbf{z}_{t-1})-\widetilde{G}(\mathbf{z}_{t-1})\big)\right\|^2 \\ &\quad\le 2\mathbb{E}_t\left\|\widetilde{G}(\mathbf{z}_t)-G(\mathbf{z}_t)-\widetilde{G}(\mathbf{z}_{t-1})+G(\mathbf{z}_{t-1})\right\|^2+2\alpha_t^2\mathbb{E}_t\|\widetilde{G}(\mathbf{z}_{t-1})-G(\mathbf{z}_{t-1})\|^2 \\ &\quad\le 2\Lambda^2\|\mathbf{z}_t-\mathbf{z}_{t-1}\|^2+2\alpha_t^2\big(\sigma^2+B^2\|\mathbf{z}_{t-1}-\mathbf{z}_0\|^2\big),\end{aligned}$$

where the last inequality follows from Assumption 1.1 and Assumption 3.1. ■

C Deferred Proofs for Corollaries

Proof of Corollary 2.2. Setting $\mathbf{z} = \mathbf{z}_t$ in (D.2), taking expectations, and applying Cauchy–Schwarz give

$$\mathbb{E}\,\mathrm{gap}(\mathbf{z}_t) \le \left[\max_{\mathbf{u}\in\mathcal{B}} \|\mathbf{z}_0 - \mathbf{u}\| + \rho\|G(\mathbf{z}_0)\|\right] \mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\| + (1+\rho L)\sqrt{\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2}\sqrt{\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\|^2} + \rho\mathbb{E}\|\mathcal{G}_\rho(\mathbf{z}_t)\|^2. \tag{C.1}$$

By (2.6), because $t \ge 1$, we have $1/(\sqrt{t+a}) \le 1/(\sqrt{a+1})$,

$$\mathbb{E}\|\mathbf{z}_t - \mathbf{z}_0\|^2 \le \frac{25}{2}\|\mathbf{z}_0 - \mathbf{z}^\star\|^2 + \frac{2\sigma^2 + 25B^2\|\mathbf{z}_0 - \mathbf{z}^\star\|^2}{H^2(a-\frac{1}{2})\sqrt{a}} < \infty.$$

Combining this bound with Theorem 2.1 gives the assertion. ∎

Corollary C.1. *Under the setup of Corollary 2.2, suppose that* $\mathbf{z} = (\mathbf{x}, \mathbf{y})$, $\mathcal{U} = \mathcal{U}_\mathbf{x} \times \mathcal{U}_\mathbf{y}$, *and*

$$G(\mathbf{z}) = \begin{pmatrix} \nabla_\mathbf{x} f(\mathbf{x}, \mathbf{y}) \\ -\nabla_\mathbf{y} f(\mathbf{x}, \mathbf{y}) \end{pmatrix},$$

where f *is convex in* $\mathbf{x}$ *and concave in* $\mathbf{y}$. *Then, we have*

$$\mathbb{E}\left[\max_{\mathbf{y}'\in\mathcal{U}_\mathbf{y}} f(\mathbf{x}_t, \mathbf{y}') - \min_{\mathbf{x}'\in\mathcal{U}_\mathbf{x}} f(\mathbf{x}', \mathbf{y}_t)\right] = O\left(\frac{1}{t^{1/4}}\right).$$

Proof of Corollary C.1. For every $(\mathbf{x}', \mathbf{y}') \in \mathcal{U}_\mathbf{x} \times \mathcal{U}_\mathbf{y}$, convexity in $\mathbf{x}$ and concavity in $\mathbf{y}$ give

$$\begin{aligned} f(\mathbf{x}, \mathbf{y}') - f(\mathbf{x}', \mathbf{y}) &= f(\mathbf{x}, \mathbf{y}') - f(\mathbf{x}, \mathbf{y}) + f(\mathbf{x}, \mathbf{y}) - f(\mathbf{x}', \mathbf{y}) \\ &\le \langle \nabla_\mathbf{y} f(\mathbf{x}, \mathbf{y}), \mathbf{y}' - \mathbf{y}\rangle + \langle \nabla_\mathbf{x} f(\mathbf{x}, \mathbf{y}), \mathbf{x} - \mathbf{x}'\rangle \\ &= \left\langle G(\mathbf{z}), \mathbf{z} - \begin{pmatrix} \mathbf{x}' \\ \mathbf{y}' \end{pmatrix} \right\rangle. \end{aligned}$$

Taking the maximum over $(\mathbf{x}', \mathbf{y}') \in \mathcal{U}_\mathbf{x} \times \mathcal{U}_\mathbf{y}$ gives

$$\max_{\mathbf{y}'\in\mathcal{U}_\mathbf{y}} f(\mathbf{x}, \mathbf{y}') - \min_{\mathbf{x}'\in\mathcal{U}_\mathbf{x}} f(\mathbf{x}', \mathbf{y}) \le \mathrm{gap}(\mathbf{z}).$$

Taking expectations at $\mathbf{z} = \mathbf{z}_t$ and applying Corollary 2.2 proves the result. ∎

D Technical Results

The following lemma is generalizing a useful idea from (Cai & Zheng, 2026) and is used many times in the proofs.

Lemma D.1. *Let* $G\colon \mathbb{R}^d \to \mathbb{R}^d$ *be monotone and* L*-Lipschitz, and* $\alpha, \gamma \ge 0$. *Then, for any* $\mathbf{u}, \mathbf{s} \in \mathbb{R}^d$ *we have*

$$\|\alpha(\mathbf{u} - \mathbf{s}) - \gamma(G(\mathbf{u}) - G(\mathbf{s}))\| \le \sqrt{\alpha^2 + \gamma^2 L^2}\|\mathbf{u} - \mathbf{s}\|.$$

Proof. By expanding the square of the left-hand side, we obtain

$$\begin{aligned} \|\alpha(\mathbf{u} - \mathbf{s}) - \gamma(G(\mathbf{u}) - G(\mathbf{s}))\|^2 = \alpha^2\|\mathbf{u} - \mathbf{s}\|^2 + \gamma^2\|G(\mathbf{u}) - G(\mathbf{s})\|^2 \\ - 2\alpha\gamma\langle \mathbf{u} - \mathbf{s}, G(\mathbf{u}) - G(\mathbf{s})\rangle. \end{aligned} \tag{D.1}$$

Due to Lipschitzness of G, we have

$$\gamma^2\|G(\mathbf{u}) - G(\mathbf{s})\|^2 \le \gamma^2 L^2\|\mathbf{u} - \mathbf{s}\|^2.$$

By monotonicity of G, we have

$$-2\alpha\gamma\langle \mathbf{u} - \mathbf{s}, G(\mathbf{u}) - G(\mathbf{s})\rangle \le 0.$$

Combining the last two estimates in (D.1) and taking the square root of both sides gives the assertion. ∎

D.1 Results on Optimality Measures

The following standard result proves that the rate of convergence that we prove for the gradient mapping norm directly translates to the same rate on the restricted gap, up to an additional multiplicative constant.

Corollary D.2. *Under Assumption 1.1, let $\mathcal{U}$ be a compact set. Let $r = \delta_Z$, where Z is nonempty, closed, and convex, but not necessarily bounded. Then, for every $\mathbf{z} \in Z$ and any $\rho > 0$, we have*

$$\mathrm{gap}_{\mathcal{U}}(\mathbf{z}) \leq \left(\max_{\mathbf{u} \in \mathcal{B}} \|\mathbf{z}_0 - \mathbf{u}\| + \rho \|G(\mathbf{z}_0)\| + (1 + \rho L)\|\mathbf{z} - \mathbf{z}_0\| \right) \|\mathcal{G}_\rho(\mathbf{z})\| + \rho \|\mathcal{G}_\rho(\mathbf{z})\|^2. \quad \text{(D.2)}$$

Proof. Fix $\mathbf{z} \in C$ and let $\mathbf{s} = \mathrm{proj}_Z(\mathbf{z} - \rho G(\mathbf{z}))$. By (1.8), we have $\frac{\mathbf{z} - \mathbf{s} - \rho G(\mathbf{z})}{\rho} \in \partial \delta_Z(\mathbf{s})$, and since $\mathcal{G}_\rho(\mathbf{z}) = \frac{\mathbf{z} - \mathbf{s}}{\rho}$, we have $\mathcal{G}_\rho(\mathbf{z}) - G(\mathbf{z}) \in \partial \delta_Z(\mathbf{s})$. As a result, convexity of δ_Z gives, for every $\mathbf{u} \in \mathcal{B}$, that

$$0 \leq \langle \mathcal{G}_\rho(\mathbf{z}) - G(\mathbf{z}), \mathbf{s} - \mathbf{u} \rangle \iff \langle G(\mathbf{z}), \mathbf{s} - \mathbf{u} \rangle \leq \langle \mathcal{G}_\rho(\mathbf{z}), \mathbf{s} - \mathbf{u} \rangle. \quad \text{(D.3)}$$

Using $\mathbf{z} - \mathbf{s} = \rho \mathcal{G}_\rho(\mathbf{z})$, we obtain

$$\begin{aligned} \langle G(\mathbf{z}), \mathbf{z} - \mathbf{u} \rangle &= \rho \langle G(\mathbf{z}), \mathcal{G}_\rho(\mathbf{z}) \rangle + \langle G(\mathbf{z}), \mathbf{s} - \mathbf{u} \rangle \\ &\leq \rho \langle G(\mathbf{z}), \mathcal{G}_\rho(\mathbf{z}) \rangle + \langle \mathcal{G}_\rho(\mathbf{z}), \mathbf{s} - \mathbf{u} \rangle \\ &\leq (\|\mathbf{s} - \mathbf{u}\| + \rho \|G(\mathbf{z})\|) \, \|\mathcal{G}_\rho(\mathbf{z})\|, \end{aligned} \quad \text{(D.4)}$$

where the first inequality is by (D.3) and the second by Cauchy-Schwarz.

Since $\mathbf{s} = \mathrm{proj}_C(\mathbf{z} - \rho G(\mathbf{z})) \in C$, for every $\mathbf{u} \in \mathcal{U}$, we have

$$\begin{aligned} \|\mathbf{s} - \mathbf{u}\| &\leq \|\mathbf{s} - \mathbf{z}\| + \|\mathbf{z} - \mathbf{z}_0\| + \|\mathbf{z}_0 - \mathbf{u}\| \\ &= \rho \|\mathcal{G}_\rho(\mathbf{z})\| + \|\mathbf{z} - \mathbf{z}_0\| + \|\mathbf{z}_0 - \mathbf{u}\|, \end{aligned}$$

where the first identity is because of $\mathbf{z} - \mathbf{s} = \rho \mathcal{G}_\rho(\mathbf{z})$. Lipschitzness of G gives $\|G(\mathbf{z})\| \leq \|G(\mathbf{z}_0)\| + L\|\mathbf{z} - \mathbf{z}_0\|$. Substituting these inequalities into (D.4) proves (D.2). ■

The following result shows that the natural residual (gradient mapping norm) is upper bounded by a residual-type quantity. Moreover, we prove that the natural residual is Lipschitz, which is another standard property.

Lemma D.3. *Under Assumption 1.1, for $\mathbf{z} \in \mathrm{dom}\, \partial r$ and any $\mathbf{s} \in \partial r(\mathbf{z})$, any $\rho > 0$ and for all $\mathbf{x}, \mathbf{y} \in \mathrm{dom}\, r$, we have*

$$\|\mathcal{G}_\rho(\mathbf{z})\| \leq \|G(\mathbf{z}) + \mathbf{s}\|, \quad \text{(D.5)}$$

$$\|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\| \leq \sqrt{\rho^{-2} + L^2} \|\mathbf{x} - \mathbf{y}\|. \quad \text{(D.6)}$$

Proof. If $\mathbf{s} \in \partial r(\mathbf{z})$, then $\mathbf{z} = \mathrm{prox}_{\rho r}(\mathbf{z} + \rho \mathbf{s})$ since the latter is equivalent to $\mathbf{z} + \rho \partial r(\mathbf{z}) \ni \mathbf{z} + \rho \mathbf{s}$. Using this identity with nonexpansiveness gives

$$\rho \|\mathcal{G}_\rho(\mathbf{z})\| = \|\mathbf{z} - \mathrm{prox}_{\rho r}(\mathbf{z} - \rho G(\mathbf{z}))\| = \|\mathrm{prox}_{\rho r}(\mathbf{z} + \rho \mathbf{s}) - \mathrm{prox}_{\rho r}(\mathbf{z} - \rho G(\mathbf{z}))\| \leq \rho \|G(\mathbf{z}) + \mathbf{s}\|.$$

We divide by ρ to obtain (D.5).

For the second assertion, the definition of the proximal operator gives $\mathcal{G}_\rho(\mathbf{x}) - G(\mathbf{x}) \in \partial r(\mathbf{x} - \rho \mathcal{G}_\rho(\mathbf{x}))$ since $\rho \mathcal{G}_\rho(\mathbf{x}) = \mathbf{x} - \mathrm{prox}_{\rho r}(\mathbf{x} - \rho G(\mathbf{x}))$, and the analogous inclusion for $\mathbf{y}$. Monotonicity of ∂r implies

$$\langle \mathcal{G}_\rho(\mathbf{x}) - G(\mathbf{x}) - [\mathcal{G}_\rho(\mathbf{y}) - G(\mathbf{y})], \mathbf{x} - \rho \mathcal{G}_\rho(\mathbf{x}) - [\mathbf{y} - \rho \mathcal{G}_\rho(\mathbf{y})] \rangle \geq 0.$$

Expanding the inner product gives

$$\rho \|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\|^2 \leq \langle \mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y}), \mathbf{x} - \mathbf{y} + \rho(G(\mathbf{x}) - G(\mathbf{y})) \rangle - \langle G(\mathbf{x}) - G(\mathbf{y}), \mathbf{x} - \mathbf{y} \rangle. \quad \text{(D.7)}$$

Young's inequality gives

$$\langle \mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y}), \mathbf{x} - \mathbf{y} + \rho(G(\mathbf{x}) - G(\mathbf{y}))\rangle \le \frac{\rho}{2}\|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\|^2 + \frac{1}{2\rho}\left(\|\mathbf{x}-\mathbf{y}\|^2 + 2\rho\langle \mathbf{x}-\mathbf{y}, G(\mathbf{x}) - G(\mathbf{y})\rangle + \rho^2\|G(\mathbf{x}) - G(\mathbf{y})\|^2\right).$$

Plugging to (D.7) yields

$$\rho\|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\|^2 \le \frac{\rho}{2}\|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\|^2 + \frac{1}{2\rho}\|\mathbf{x}-\mathbf{y}\|^2 + \frac{\rho}{2}\|G(\mathbf{x}) - G(\mathbf{y})\|^2.$$

After rearranging and using Lipschitz continuity of G, we deduce

$$\|\mathcal{G}_\rho(\mathbf{x}) - \mathcal{G}_\rho(\mathbf{y})\|^2 \le \rho^{-2}\|\mathbf{x}-\mathbf{y}\|^2 + \|G(\mathbf{x}) - G(\mathbf{y})\|^2 \le (\rho^{-2} + L^2)\|\mathbf{x}-\mathbf{y}\|^2.$$

Taking square root of both sides proves (D.6). ■

D.2 Numerical Consequences of Parameters

The following fact collects some tedious, yet important estimations we used, which follow from the choice of our parameters.

Fact D.4. *Suppose that $a \ge 2$, $H^2 = L^2 + 2B^2$, and*

$$\beta_t = \frac{a}{t+a}, \qquad \eta_t = \frac{1}{H(t+a)^{3/4}}.$$

Then, for every $t \ge 1$,

$$\sqrt{(1-\beta_t)^2 + \eta_t^2 L^2} \le 1 - \frac{2}{3}\beta_t. \tag{D.8}$$

Moreover, for every $t \ge 0$, we have,

$$\left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)^2 + \eta_{t+1}^2 L^2 \le \left(1 - \frac{3}{2(t+a+1)}\right)^2, \tag{D.9}$$

and

$$\left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right| \le \frac{a}{4(t+a)(t+a+1)}. \tag{D.10}$$

Finally, for every $t \ge 1$,

$$(1-\beta_t)^2 + \eta_t^2(L^2 + 2B^2) \le 1 - \frac{a}{t+a}. \tag{D.11}$$

Proof. For the first estimate, we have

$$\begin{aligned}\left(1 - \frac{2}{3}\beta_t\right)^2 - (1-\beta_t)^2 - \eta_t^2 L^2 &= \frac{2}{3}\beta_t - \frac{5}{9}\beta_t^2 - \eta_t^2 L^2 \\ &= \frac{1}{t+a}\left(\frac{2a}{3} - \frac{5a^2}{9(t+a)} - \frac{L^2}{H^2\sqrt{t+a}}\right) \\ &\ge \frac{1}{t+a}\left(\frac{2a}{3} - \frac{5a^2}{9(t+a)} - \frac{1}{\sqrt{t+a}}\right) \\ &\ge \frac{1}{t+a}\left(\frac{a(a+6)}{9(a+1)} - \frac{1}{\sqrt{a+1}}\right) \ge 0.\end{aligned}$$

Here, the first inequality uses $H \ge L$, and the second uses $t \ge 1$. The last inequality follows from $a(a+6) \ge 9\sqrt{a+1}$ for $a \ge 2$. This proves (D.8).

We next prove the estimates for $t \ge 0$. Since

$$\frac{\eta_{t+1}}{\eta_t} = \frac{(t+a)^{3/4}}{(t+a+1)^{3/4}} \ge \frac{t+a}{t+a+1} \ge \beta_{t+1} = \frac{a}{t+a+1},$$

we have $\eta_{t+1}/\eta_t - \beta_{t+1} \geq 0$. Write $s = t + a + 1 \geq 3$. By concavity of $x^{3/4}$, we have $\left(1 - \frac{1}{s}\right)^{3/4} \leq 1 - \frac{3}{4s}$. Therefore, by definition of η_t and β_{t+1}, we have

$$0 \leq \frac{\eta_{t+1}}{\eta_t} - \beta_{t+1} = \left(1 - \frac{1}{s}\right)^{3/4} - \frac{a}{s} \leq 1 - \frac{3}{4s} - \frac{a}{s} \leq 1 - \frac{11}{4s},$$

where the last inequality uses $a \geq 2$. Moreover, $\eta_{t+1}^2 L^2 = \frac{L^2}{H^2} s^{-3/2} \leq s^{-3/2}$ because $H \geq L$. Therefore,

$$\left(\frac{\eta_{t+1}}{\eta_t} - \beta_{t+1}\right)^2 + \eta_{t+1}^2 L^2 \leq \left(1 - \frac{11}{4s}\right)^2 + s^{-3/2} \leq \left(1 - \frac{3}{2s}\right)^2.$$

The last inequality follows from $\frac{5}{2} - \frac{85}{16s} - s^{-1/2} \geq \frac{35}{48} - \frac{1}{\sqrt{3}} > 0$ because $s \geq 3$. This proves (D.9).

Similarly, using the definitions of β_t and η_t, we have

$$\begin{aligned}
\left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right| &= \left|\frac{a}{t+a+1} - \left(\frac{t+a}{t+a+1}\right)^{3/4} \frac{a}{t+a}\right| \\
&= \frac{a}{t+a+1}\left|1 - \left(\frac{t+a+1}{t+a}\right)^{1/4}\right| \\
&= \frac{a}{t+a+1}\left[\left(1 + \frac{1}{t+a}\right)^{1/4} - 1\right].
\end{aligned}$$

The last equality follows because $\left(1 + \frac{1}{t+a}\right)^{1/4} \geq 1$. Then, concavity of $x^{1/4}$ yields $\left(1 + \frac{1}{t+a}\right)^{1/4} \leq 1 + \frac{1}{4(t+a)}$. Therefore, we have

$$\left|\beta_{t+1} - \frac{\eta_{t+1}}{\eta_t}\beta_t\right| \leq \frac{a}{4(t+a)(t+a+1)},$$

which proves (D.10).

Finally, for $t \geq 1$,

$$\begin{aligned}
(t+a)\left[1 - (1-\beta_t)^2 - \eta_t^2(L^2 + 2B^2)\right] &= 2a - \frac{a^2}{t+a} - \frac{L^2 + 2B^2}{H^2\sqrt{t+a}} \\
&\geq 2a - \frac{a^2}{t+a} - \frac{1}{\sqrt{t+a}} \geq a. \qquad \text{(D.12)}
\end{aligned}$$

Here, the first inequality uses $H^2 = L^2 + 2B^2$, while the last inequality follows from

$$a - \frac{a^2}{t+a} - \frac{1}{\sqrt{t+a}} \geq \frac{a}{a+1} - \frac{1}{\sqrt{a+1}} \geq 0$$

for $t \geq 1$ and $a \geq 2$. Dividing (D.12) by $t + a$ proves (D.11). ■

Fact D.5. *Suppose* $H = \max\left\{2L,\ 4\sqrt{1024\Lambda^2 + 4 + 16B^2}\right\}$ *and let* $\Gamma = 1024\Lambda^2 + 4 + 16B^2$, *with*

$$\alpha_t = \beta_t = \frac{1}{t+1}, \quad \textit{and} \quad \eta_t = \frac{1}{H\sqrt{t+1}}.$$

Then for every $t \geq 1$, *following five estimates hold:*

1.

$$\left(1 - \frac{1}{t+1}\right)^2 \frac{\Gamma C}{t} + \frac{512\Lambda^2 C}{(t+1)^2} + \frac{2(C + 4B^2 C)}{(t+1)^2} = \frac{\Gamma C}{t+1} - \frac{\Gamma C}{2(t+1)^2}. \qquad \text{(D.13)}$$

2.

$$\left(1-\beta_t+\frac{\eta_t^2L^2}{1-\beta_t}\right)\le\left(1-\frac{\beta_t}{2}\right). \tag{D.14}$$

3.

$$\begin{aligned}&40(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)^2C+\frac{10(t+1)\eta_t^2\alpha_t^2\Gamma C}{t}\\&+\eta_t^2\left(\frac{512\Lambda^2C}{(t+1)^2}+2\alpha_t^2\left(\sigma^2+4B^2C\right)\right)\le\left(40+\frac{41\Gamma}{2H^2}\right)\frac{C}{(t+1)^3}.\end{aligned} \tag{D.15}$$

4.

$$40+\frac{41\Gamma}{2H^2}\le\frac{1321}{32}\le 64. \tag{D.16}$$

5.

$$\begin{aligned}&\frac{256C}{(t+1)^2}\left(1+\frac{1}{4(t+1)}\right)\left(\left(1-\frac{3}{2(t+1)}\right)^2+\frac{1}{4(t+1)}\right)+\frac{64C}{(t+1)^3}\\&\le\frac{256C}{(t+2)^2}.\end{aligned} \tag{D.17}$$

Proof. We start from the first statement, because $1-\frac{1}{t+1}=\frac{t}{t+1}$, the left-hand side satisfies

$$\begin{aligned}&\left(1-\frac{1}{t+1}\right)^2\frac{\Gamma C}{t}+\frac{512\Lambda^2C}{(t+1)^2}+\frac{2(C+4B^2C)}{(t+1)^2}\\&=\frac{t\Gamma C}{(t+1)^2}+\frac{\Gamma C}{2(t+1)^2}\\&=\frac{(t+1/2)\Gamma C}{(t+1)^2}\\&=\frac{\Gamma C}{t+1}-\frac{\Gamma C}{2(t+1)^2},\end{aligned}$$

which proves (D.13).

Then we start the second statement. Using the schedule of β_t and η_t, we have $\beta_t\le 1/2$ and $\eta_t^2L^2\le\beta_t/4$. Using $\frac{1}{4(1-\beta_t)}\le\frac{1}{2}$, it suffices to show,

$$\left(1-\beta_t+\frac{\eta_t^2L^2}{1-\beta_t}\right)\le 1-\beta_t+\frac{\beta_t}{4(1-\beta_t)}\le\left(1-\frac{\beta_t}{2}\right),$$

for every $t\ge 1$. This finishes proof of (D.14).

We next move to the third statement. We estimate first term of the left-hand side, applying the definition of η_t and β_t gives,

$$\begin{aligned}40(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)^2C&=40(t+1)\left(\frac{1}{\sqrt{t(t+1)}}-\frac{1}{t+1}\right)^2C\\&=40(t+1)\left(\frac{1}{\sqrt{t}(t+1)(\sqrt{t+1}+\sqrt{t})}\right)^2C.\end{aligned}$$

Then, because $\sqrt{t(t+1)}\ge 1$, we have

$$\frac{1}{\sqrt{t}(t+1)(\sqrt{t+1}+\sqrt{t})}\le\frac{1}{(t+1)^2}.$$

Therefore, we have

$$40(t+1)\left(\frac{\eta_t}{\eta_{t-1}}\beta_{t-1}-\beta_t\right)^2 C \le \frac{40C}{(t+1)^3}.$$

For the second term, using the definition of η_t and α_t,

$$\frac{10(t+1)\eta_t^2\alpha_t^2\Gamma C}{t} = \frac{10\Gamma C}{H^2 t(t+1)^2} \le \frac{20\Gamma}{H^2}\frac{C}{(t+1)^3},$$

where we used $1/t \le 2/(1+t)$ for the last inequality.

For the last term, substituting the η_t and α_t yield

$$\begin{aligned}\eta_t^2\left(\frac{512\Lambda^2 C}{(t+1)^2}+2\alpha_t^2(\sigma^2+4B^2C)\right) &= \frac{1}{H^2(t+1)}\left(\frac{512\Lambda^2 C}{(t+1)^2}+\frac{2(\sigma^2+4B^2C)}{(t+1)^2}\right)\\ &\le \frac{\Gamma}{2H^2}\frac{C}{(t+1)^3},\end{aligned}$$

where the last inequality follows from $\sigma^2 \le C$.

Combining the three inequalities gives us

$$\frac{40C}{(t+1)^3}+\frac{20\Gamma}{H^2}\frac{C}{(t+1)^3}+\frac{\Gamma}{2H^2}\frac{C}{(t+1)^3} = \left(40+\frac{41\Gamma}{2H^2}\right)\frac{C}{(t+1)^3}.$$

Which proves the claim.

For the fourth statement, the definition of H gives $H^2 \ge 16\Gamma$. Thus we have $\Gamma/H^2 \le 1/16$, and consequently

$$40+\frac{41\Gamma}{2H^2} \le 40+\frac{41}{32} = \frac{1321}{32} \le 64.$$

This proves the fourth statement.

It remains to prove the last statement. Expanding the factor gives,

$$\left(1+\frac{1}{4(t+1)}\right)\left(\left(1-\frac{3}{2(t+1)}\right)^2+\frac{1}{4(t+1)}\right) = \frac{16(t+1)^3-40(t+1)^2+25(t+1)+9}{16(t+1)^3}. \tag{D.18}$$

Subtracting the left-hand side of (D.17) from its right-hand side yields

$$\begin{aligned}&\frac{256C}{(t+2)^2}-\frac{256C}{(t+1)^2}\left(1+\frac{1}{4(t+1)}\right)\left(\left(1-\frac{3}{2(t+1)}\right)^2+\frac{1}{4(t+1)}\right)-\frac{64C}{(t+1)^3}\\ &\quad=\frac{256C}{(t+2)^2}-\frac{16C(16(t+1)^3-36(t+1)^2+25(t+1)+9)}{(t+1)^5}\\ &\quad=\frac{16C(4(t+1)^4+31(t+1)^3-23(t+1)^2-43(t+1)-9)}{(t+1)^5(t+2)^2},\end{aligned}$$

where the second equality uses (D.18), and the remaining equalities follow from direct calculation.

The numerator satisfies

$$\begin{aligned}&4(t+1)^4+31(t+1)^3-23(t+1)^2-43(t+1)-9\\ &\quad=(t-1)(4(t+1)^3+39(t+1)^2+55(t+1)+67)+125>0,\end{aligned}$$

for every $t \ge 1$. This completes the proof of the last statement. ■

E Details of Section 4 and Further Numerical Results

We provide additional details and experiments supporting the numerical results in Section 4. We first define the optimality measures used in our experiments and describe the experiment settings for two constrained problems considered in Section 4. We then consider the quadratic problem, which is referred to as the hard convex-concave problem by Abe et al. (2025) as an additional benchmark. Next, we give give details on the initialization and hyperparameter tuning. Finally, we report initialization sensitivity for our two Halpern method.

E.1 Optimality measures

Let us first restate the optimality measures used throughout our experiments.

Gap. We use the gap defined as (1.5),

$$\mathrm{gap}(\mathbf{z}) = \max_{\mathbf{u}\in Z} \{\langle G(\mathbf{z}), \mathbf{z}-\mathbf{u}\rangle\}.$$

For the simplex bilinear problem, this reduces to standard primal-dual gap used for matrix games, see for example (Carmon et al., 2019),

$$\mathrm{gap}(\mathbf{x},\mathbf{y}) = \max_j (A^\top \mathbf{x})_j - \min_i (A\mathbf{y})_i.$$

Gradient mapping. We report the gradient mapping norm

$$\mathcal{G}_{1/L}(\mathbf{z}) = L\big(\mathbf{z} - \mathrm{proj}_Z(\mathbf{z} - \frac{1}{L}G(\mathbf{z}))\big).$$

The projection is computed using Euclidean projection onto the simplex or clipping for the ℓ_∞ constraints.

Tangent residual. We measure stationarity using the tangent residual with $r = \delta_Z$,

$$\mathrm{res}(\mathbf{z}) = \min_{\mathbf{s}\in N_Z(\mathbf{z})} \|G(\mathbf{z}) + \mathbf{s}\|,$$

where the $N_Z(\mathbf{z})$ denotes the normal cone of Z at $\mathbf{z}$.

E.2 Experiment setup

We next provide further details on the two constrained problems considered in Section 4 and describe additional hard convex-concave problem from (Abe et al., 2025).

Matrix games. We consider the standard setup

$$\min_{\mathbf{x}\in\Delta^m} \max_{\mathbf{y}\in\Delta^n} \mathbf{x}^\top A\mathbf{y}.$$

Here, Δ^d denotes the d-dimensional simplex. We set $m = n = 50$ and construct A using the random payoff game of (Abe et al., 2025). Unlike their additive noise, rather we sample from the i-th row and j-th column of A, so that operator gives

$$\widetilde{G}(\mathbf{x},\mathbf{y};i,j) = \begin{pmatrix} nA_{:j}y_j \\ -mA_{i:}^\top x_i \end{pmatrix},$$

which is an unbiased estimator of $G(\mathbf{x},\mathbf{y}) = (A\mathbf{y}, -A^\top\mathbf{x})^\top$.

Robust least squares. We next test on the robust least squares problem studied by (Cai et al., 2022a), with additional ℓ_∞ constraints

$$\min_{\mathbf{x}\in\mathcal{X}} \max_{\mathbf{y}\in\mathcal{Y}} \frac{1}{2n}\|A\mathbf{x}-\mathbf{y}\|^2 - \frac{1.5}{2n}\|\mathbf{y}-\mathbf{b}\|^2,$$

where $\mathcal{X} = \{\mathbf{x} : \|\mathbf{x}\|_\infty \leq 1\}$ and $\mathcal{Y} = \{\mathbf{y} : \|\mathbf{y} - \mathbf{b}\|_\infty \leq 1/2\}$. Following (Cai et al., 2022a), we use the same real dataset with $n = 21263$ samples and $d = 81$ features. To adapt the constraint properly, we normalize each feature row and rescale the label vector by the root mean square values. We sample a data row uniformly and use

$$\widetilde{G}(\mathbf{x}, \mathbf{y}; i) = \begin{pmatrix} a_i(a_i^\top \mathbf{x} - y_i) \\ \mathbf{e}_i(a_i^\top \mathbf{x} + \frac{1}{2}y_i - \frac{3}{2}b_i) \end{pmatrix},$$

which is an unbiased estimator.

Hard convex-concave game. We further evaluate on the hard convex-concave game of (Abe et al., 2025). Following their setup, we consider

$$\min_{\mathbf{x}\in\mathcal{X}_\infty} \max_{\mathbf{y}\in\mathcal{X}_2} \frac{1}{2}\mathbf{x}^\top H\mathbf{x} - \mathbf{h}^\top \mathbf{x} - \langle A\mathbf{x} - \mathbf{b}, \mathbf{y}\rangle,$$

where $\mathcal{X}_1 = \mathcal{X}_2 = [-200, 200]^d$ with $d = 100$, $H = 2A^\top A$, and A, $\mathbf{b}$, and $\mathbf{h}$ are constructed as in (Abe et al., 2025). We use additive noise, so that operator is

$$\widetilde{G}(\mathbf{z}, \xi) = \begin{pmatrix} H\mathbf{x} - \mathbf{h} - A^\top \mathbf{y} \\ A\mathbf{x} - \mathbf{b} \end{pmatrix} + \begin{pmatrix} \xi_1 \\ \xi_2 \end{pmatrix}, \qquad \xi_1, \xi_2 \sim \mathcal{N}(0, 0.1^2 I).$$

Here, the known solution is $\mathbf{z}^\star = ((1, 2, , \ldots, d)^\top, -\frac{1}{2}\mathbf{1})$.

E.3 Additional optimality measures

We supplement the results in Section 4 with additional optimality measures for matrix games and robust least squares problems. We also report results for the hard convex-concave game.

Matrix games. Figure 2 reports the gradient mapping norm and tangent residual, complementing the gap results in Section 4.

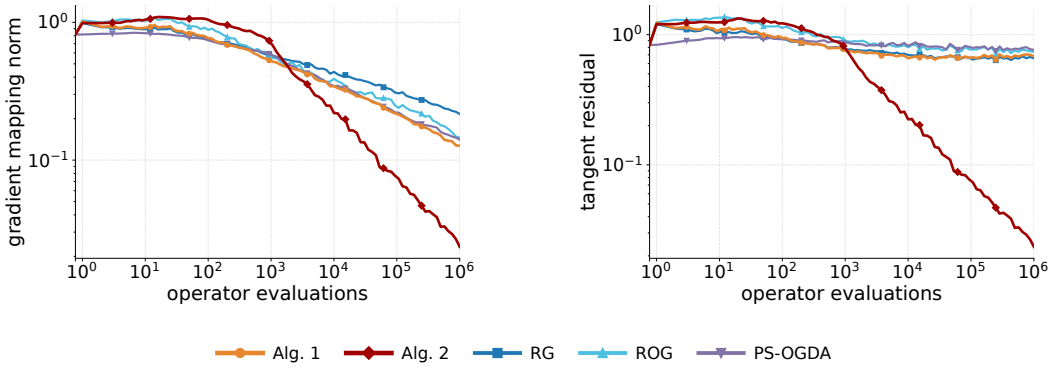

Figure 2: Matrix games. Left: gradient mapping norm. Right: tangent residual.

Robust least squares. Figure 3 reports the gap and tangent residual, complementing the gradient mapping norm results in Section 4.

Hard convex-concave game. Figure 4 reports the gap, gradient mapping, tangent residual, and distance to the solution $\mathbf{z}^\star$.

The additional optimality measures show a similar behavior to the results in Section 4. Across the tested problems, Algorithm 2 generally outperforms the baselines, while Algorithm 1 shows comparable or slightly better performance, depending on the problem.

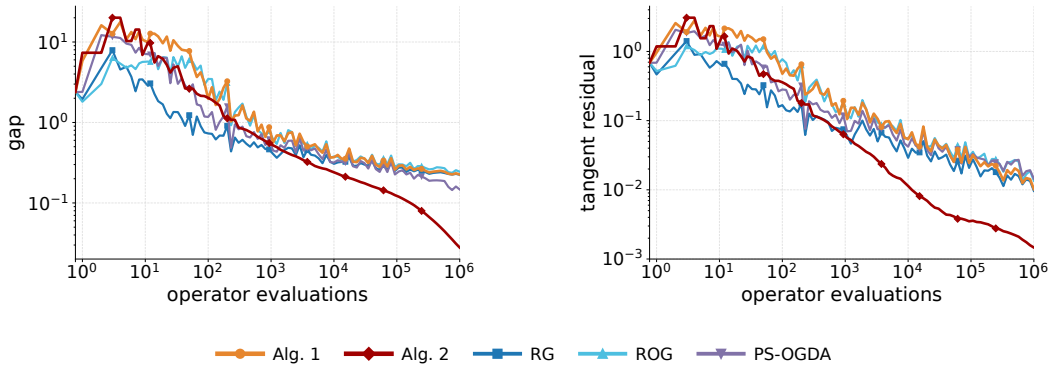

Figure 3: Robust least squares. Left: gap. Right: tangent residual.

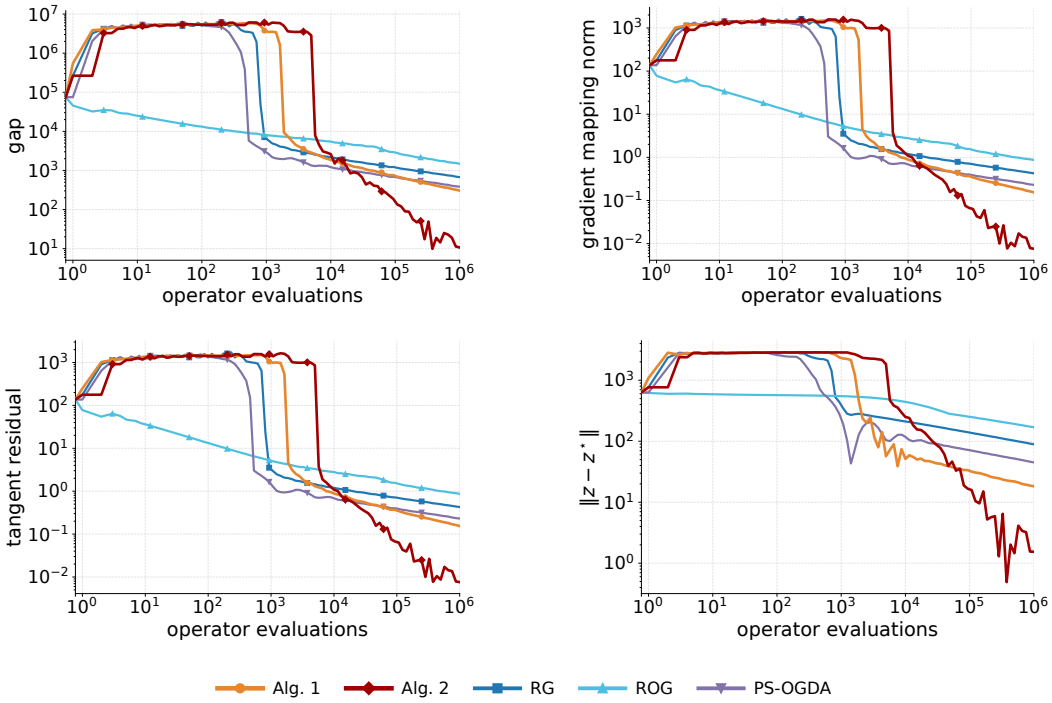

Figure 4: Hard convex-concave game. Top left: gap. Top right: gradient mapping norm. Bottom left: tangent residual. Bottom right: distance to the solution.

E.4 Initialization and hyperparameter choices

Initialization. All methods use the same initial point for each problem. For the matrix games, following Abe et al. (2025), we initialize $\mathbf{x}_0 = \frac{1}{m}\mathbf{1}$ and $\mathbf{y}_0 = \frac{1}{n}\mathbf{1}$. For robust least squares, we initialize $\mathbf{x}_0 = \mathbf{0}$ and $\mathbf{y}_0 = \mathbf{b}$. For the hard convex-concave game, we initialize $\mathbf{x}_0 = \mathbf{y}_0 = 200\mathbf{e}_{50}$.

Hyperparameter tuning. We follow the parameter schedules of each method and tune their coefficients for each problem. For each method, we consider 12 parameter choices for the matrix games and robust least squares problems, and 16 for the hard convex-concave game. We evaluate each choice over three validation runs, using 10^6 operator evaluations per run. We finally choose the coefficients that give the smallest average of the last-iterate gaps from the three validation runs. The search grids are listed in Table 3, and selected coefficients are reported in Table 4.

For η_0 and τ_0, we consider values of the 10^k or 5×10^k, within the ranges listed in Table 3. For RG and ROG, we consider $\gamma_0 \in \{0.1, 0.5, 1\}$, which keeps the anchoring weight in $(0, 1]$. For each stepsize η_0, we assign a value of γ_0 for RG and ROG, or τ_0 for PS-OGDA, from the following grid. We generate these parameter pairs for each method, then evaluate each combination in three validation runs.

Coefficient	**Matrix games**	**Robust least squares**	**Hard convex-concave**
η_0	$[10^{-4}, 50]$	$[10^{-4}, 50]$	$[10^{-4}, 5000]$
γ_0 (RG, ROG)	$\{0.1, 0.5, 1\}$	$\{0.1, 0.5, 1\}$	$\{0.1, 0.5, 1\}$
τ_0 (PS-OGDA)	$[10^{-3}, 10]$	$[5 \times 10^{-5}, 1]$	$[10^{-4}, 1]$
Number of config.	12	12	16

Table 3: Candidate coefficient grids.

Method	**Tuned coef.**	**Matrix games**	**Robust least squares**	**Hard convex-concave**
Alg. 1	η_0	0.1	50	500
Alg. 2	η_0	0.05	50	100
RG	(η_0, γ_0)	$(0.05, 1)$	$(5, 0.1)$	$(100, 1)$
ROG	(η_0, γ_0)	$(0.1, 0.1)$	$(1, 0.1)$	$(1, 0.5)$
PS-OGDA	(η_0, τ_0)	$(0.01, 1)$	$(10, 5 \times 10^{-5})$	$(50, 0.005)$

Table 4: Selected hyperparameter coefficients.

E.5 Sensitivity of initialization

We test the sensitivity of Algorithm 1 and 2 to initialization on the matrix games using ten random initial starting points. For each initial point, we draw from $\mathcal{U}(0, 1)$ and normalize each vector so that its coordinates sum to one. Both methods use the same ten initial points. We keep the payoff matrix and sampling seeds fixed, and use the parameters selected in the main experiment without retuning. Figure 5 reports the gap for ten individual runs.

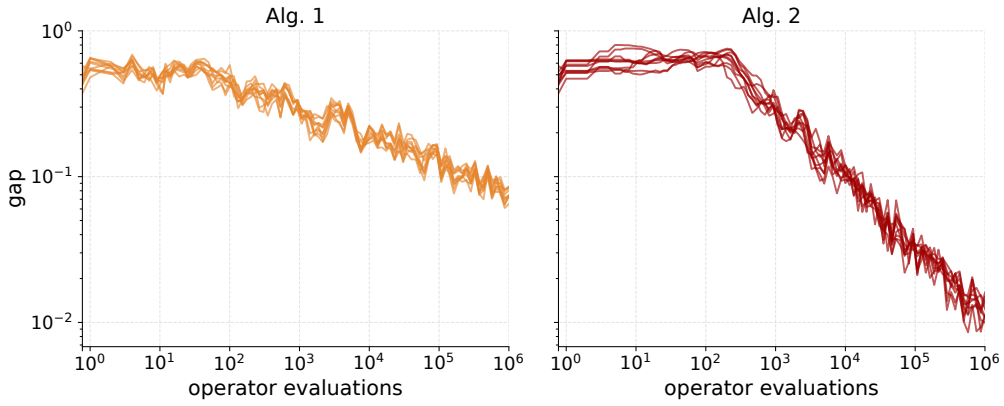

Figure 5: Left: Alg. 1. Right: Alg. 2

Both methods show similar convergence behavior across the ten initial points, suggesting that their performance is relatively insensitive to the tested random initializations.