ACCELERATED ALGORITHMS FOR STOCHASTIC MONOTONE INCLUSIONS WITH FIXED QUERIES

Anonymous authors
Paper under double-blind review

ABSTRACT

We study acceleration of stochastic first-order methods for monotone Lipschitz inclusions with a fixed number of oracle queries per iteration, measured by the expected squared residual. Existing methods either have suboptimal oracle complexity, require an increasing number of oracle queries per iteration, or both. Under same-sample oracle access and mean-square Lipschitz stochastic noise, we develop variance-reduced anchored forward-backward (VRAF), an anytime composite method that makes two oracle queries per iteration. VRAF attains the first log-free $O(1/\epsilon)$ complexity under these oracle assumptions and hence the optimal $O(\sigma^2/\epsilon)$ variance dependence. This, however, leaves open whether the optimal $O(1/\sqrt{\epsilon})$ deterministic term can also be attained. We establish an impossibility result by extending the lower bound of Foster et al. (2019) to stochastic oracles permitting repeated queries to sampled stochastic operators, showing that $O(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ oracle complexity is unattainable in general. We therefore develop recentered regularized stochastic extragradient (RRSEG), which attains the near-optimal oracle complexity previously achievable only with an increasing number of queries per iteration, while using a fixed number of queries.

1 INTRODUCTION

Finding equilibrium points and solving minimax problems are fundamental tasks in optimization and modern machine learning. Representative examples arise in generative adversarial networks (Goodfellow et al., 2014; Arjovsky et al., 2017), game theory (Rosen, 1965), multi-agent learning (Littman, 1994; Mazumdar et al., 2025), and recent approaches to LLM alignment (Munos et al., 2024; Xu et al., 2025).

In large-scale learning problems, the underlying operator is typically accessible only through a stochastic oracle, motivating extensive efforts to develop efficient stochastic methods. Yet our understanding remains limited even for basic monotone Lipschitz inclusions. For the expected squared residual, a standard optimality measure in this setting, the oracle complexity has a lower bound of $\Omega(\sigma^2/\epsilon)$ for achieving accuracy ϵ (Foster et al., 2019), where σ^2 bounds the stochastic oracle variance. However, no stochastic method is known to attain the optimal dependence on σ^2. E-Halpern and S-Dual-OHM have $O(\sigma^2\epsilon^{-3/2})$ stochastic terms, while RAIN improves this to the near-optimal $\widetilde{O}(\sigma^2/\epsilon)$ term (Cai et al., 2022; Yoon & Loizou, 2026; Chen & Luo, 2024). These methods, however, require an increasing number of oracle queries per iteration as iterations proceed or as ϵ decreases. This motivates our first question: *Can the optimal $O(\sigma^2/\epsilon)$ stochastic term be attained with a fixed number of oracle queries per iteration?*

We answer this question affirmatively with variance-reduced anchored forward-backward (VRAF), an anytime fixed-query method for monotone Lipschitz inclusions. VRAF makes exactly two oracle queries per iteration, both using the same sampled stochastic operator. It achieves $O(((L^2+L_\Delta^2)D^2+\sigma^2)/\epsilon)$ oracle complexity, where L is the Lipschitz constant of the operator, L_Δ is the mean-square Lipschitz constant of the stochastic noise, and D is the initial distance. To our knowledge, this is the first log-free $O(1/\epsilon)$ oracle complexity for the expected squared residual of monotone Lipschitz inclusions under same-sample access and Lipschitz stochastic noise, even among methods allowing an increasing number of oracle queries 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.

A worst-case example shows that the $O(L^2D^2/\epsilon)$ deterministic term of VRAF is order-tight. This leads to our second question: *Can a different method retain the optimal dependence on σ^2 while also attaining the optimal $O(LD/\sqrt{\epsilon})$ deterministic term?* Such simultaneous optimality is possible in stochastic convex minimization and, for gap functions, in stochastic variational inequalities (Lan, 2012; Juditsky et al., 2011). For the expected squared residual, the corresponding target is $O(LD/\sqrt{\epsilon}+\sigma^2/\epsilon)$. We establish an impossibility result by extending the worst-case construction of Foster et al. (2019) to stochastic oracles that allow repeated queries to previously sampled stochastic operators, showing that this complexity is unattainable even for two-dimensional monotone Lipschitz equations satisfying $L_\Delta \le L$.

This impossibility motivates our second method, recentered regularized stochastic extragradient (RRSEG). RRSEG uses two independent oracle queries per iteration and achieves $\widetilde{O}(LD/\sqrt{\epsilon}+\sigma^2/\epsilon)$ oracle complexity, with an optimal deterministic term and a near-optimal stochastic term. More explicitly, RRSEG has an $O(\sigma^2 \log^3 N/N)$ stochastic term, matching the logarithmic power in the near-optimal bound of RAIN, while using a fixed, rather than increasing, number of oracle queries per iteration. Corollary 5.2 rules out every logarithmic power $p < 1$ for fixed-query methods, leaving the range $1 \le p < 3$ open.

2 Problem Setting

Let $\mathcal{H}$ be a real Hilbert space. We consider the monotone Lipschitz inclusion

$$0 \in F(z_\star) + A(z_\star), \tag{1}$$

where $z_\star \in \mathcal{H}$. When $A = 0$, equation 1 reduces to $F(z_\star) = 0$, which we call a monotone Lipschitz equation.

Assumption 2.1. *For some $L > 0$, the operator $F : \mathcal{H} \to \mathcal{H}$ is monotone and L-Lipschitz continuous, i.e., for all $x, y \in \mathcal{H}$,*

$$\langle F(x) - F(y), x - y\rangle \ge 0, \qquad \|F(x) - F(y)\| \le L\|x - y\|,$$

and $A : \mathcal{H} \rightrightarrows \mathcal{H}$ is maximally monotone. For a given initial point z_0, there exists a solution $z_\star$ of equation 1 such that $\|z_0 - z_\star\| \le D$ for some $D > 0$.

We access F only through its unbiased and bounded stochastic oracle F_ξ and assume exact access to the resolvent $J_{\alpha A} := (I + \alpha A)^{-1}$ for every $\alpha > 0$.

Assumption 2.2 (Unbiased and bounded stochastic oracle)**.** *For some $\sigma \ge 0$ and every $z \in \mathcal{H}$,*

$$F_\xi(z) = F(z) + \delta_\xi(z), \qquad \mathbb{E}[\delta_\xi(z)] = 0, \qquad \mathbb{E}\|\delta_\xi(z)\|^2 \le \sigma^2.$$

We also use the following additional oracle conditions where stated.

Assumption 2.3 (Same-sample oracle access)**.** *For each sampled stochastic operator F_ξ, the algorithm may evaluate F_ξ at multiple query points.*

Assumption 2.4 (Lipschitz stochastic noise)**.** *There exists $L_\Delta \ge 0$ such that, for all $x, y \in \mathcal{H}$,*

$$\mathbb{E}\|\delta_\xi(x) - \delta_\xi(y)\|^2 \le L_\Delta^2\|x - y\|^2.$$

Same-sample oracle access and mean-square Lipschitz continuity of the stochastic oracle are commonly used in stochastic variance reduction Cai et al. (2022); Pethick et al. (2023). Together with the Lipschitz continuity of F, Assumption 2.4 is equivalent up to constants to the latter condition.[1] Assumptions 2.1, 2.2, and 2.4 therefore imply

$$\mathbb{E}\|F_\xi(x) - F_\xi(y)\|^2 \le \overline{L}^2\|x - y\|^2,$$

where $\overline{L}^2 := L^2 + L_\Delta^2$.

Fresh samples are independent. The access in Assumption 2.3 is available when ξ identifies, for example, a retained data sample or a replayable random seed, but not for a fresh-sample-only black-box oracle. Each evaluation $F_\xi(z)$ counts as one oracle query; resolvent evaluations are not counted as oracle queries.

[1] Under Assumption 2.2, $\mathbb{E}\|F_\xi(x) - F_\xi(y)\|^2 = \|F(x) - F(y)\|^2 + \mathbb{E}\|\delta_\xi(x) - \delta_\xi(y)\|^2$.

We call a method *fixed-query* if there exists a constant r, depending only on the method, such that every iteration uses at most r oracle queries, uniformly over all admissible problem instances, iteration indices, the total number of iterations (horizon), and the target accuracy.

Define the residual $\mathcal{R}(z) := \operatorname{dist}(0, F(z) + A(z))$. We use the expected squared residual as our optimality measure (Cai et al., 2022; Pethick et al., 2023):

$$\mathbb{E}[\mathcal{R}(z)^2].$$

In the deterministic setting, the corresponding criterion is the squared tangent residual used by Cai et al. (2024). When $A = 0$, $\mathcal{R}(z) = \|F(z)\|$. The oracle complexity is the number of stochastic oracle queries required to produce an output $\widehat{z}$ satisfying

$$\mathbb{E}[\mathcal{R}(\widehat{z})^2] \le \epsilon.$$

3 Related Work

Deterministic methods. For monotone Lipschitz equations, extragradient and Popov's method have $O(L^2D^2/\epsilon)$ oracle complexity for squared residual accuracy ϵ (Korpelevich, 1976; Popov, 1980; Gorbunov et al., 2022a;b). Extragradient-type methods extend to monotone Lipschitz inclusions through forward-backward-forward splitting (Tseng, 2000). Anchoring-only methods attain the same order, and P-AGD extends the anchored forward-backward update to monotone Lipschitz inclusions (Ryu et al., 2020; Surina et al., 2026; Cai & Zheng, 2026). In these anchoring-only methods, the stepsize is diminishing. By contrast, combining anchoring with an extragradient or Popov-style optimistic correction can maintain a nonvanishing stepsize and attain the optimal $O(LD/\sqrt{\epsilon})$ complexity; representative examples include EAG, FEG, and their extensions (Yoon & Ryu, 2021; Lee & Kim, 2021; Cai et al., 2024; Alcala et al., 2023; Tran-Dinh & Luo, 2021; Cai & Zheng, 2023). This distinction is relevant here because VRAF builds on an anchoring-only forward-backward structure.

Stochastic methods. Stochastic extensions of accelerated anchored methods, such as S-FEG and moving-anchor EAG-V, preserve the accelerated deterministic term under prescribed variance decay but yield a suboptimal stochastic term (Lee & Kim, 2021; Alcala et al., 2025). More generally, no existing stochastic method for the expected squared residual attains the optimal $O(\sigma^2/\epsilon)$ term in the oracle complexity. E-Halpern has $O(\epsilon^{-3/2})$ oracle complexity but uses an increasing number of oracle queries per iteration (Cai et al., 2022). S-Dual-OHM achieves the same order with a number of oracle queries per iteration that increases as ϵ decreases, but requires cocoercivity in expectation (Yoon & Loizou, 2026). RAIN attains the near-optimal $\widetilde{O}(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ oracle complexity for monotone Lipschitz problems, but is not fixed-query (Chen & Luo, 2024). Among fixed-query methods, BC-(P)SEG+, stochastic GOMA, and a single-call stochastic Halpern method have $O(\epsilon^{-2})$ oracle complexity (Pethick et al., 2023; Sohrabi et al., 2026; Kim & Alacaoglu, 2026). VRAF attains the optimal $O(\sigma^2/\epsilon)$ term with a fixed number of oracle queries per iteration. Table 1 summarizes the main comparisons.

Deterministic and stochastic optimality. In stochastic convex minimization, the optimal deterministic and stochastic terms can be attained together for function-value accuracy. For the expected squared residual considered here, however, the lower-bound construction of Foster et al. (2019) yields an additional logarithmic factor. Section 5 extends this construction to stochastic oracles that allow repeated queries to previously sampled stochastic operators, showing that $O(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ oracle complexity is unattainable in general.

Recursive regularization. To approach the accelerated deterministic rate while retaining a near-optimal stochastic rate, RAIN employs recursive regularization. This technique was introduced for stochastic convex optimization by Allen-Zhu (2018) and was later adapted to stochastic minimax optimization by Chen & Luo (2024), who proposed RAIN, and to differentially private stochastic variational inequalities by Bassily et al. (2024). RAIN solves each regularized subproblem with a stochastic subroutine, which leads to an increasing number of oracle queries per iteration. RAIN also considers a single-loop variant in its numerical experiments (Algorithm 9 in Chen & Luo (2024)), but does not provide a convergence-rate guarantee for this variant. Building on this single-loop structure,

we develop RRSEG with horizon-dependent parameters and prove an $O(L^2D^2/N^2 + \sigma^2 \log^3 N/N)$ bound on the expected squared residual using two oracle queries per iteration.

4 VRAF: Optimal Variance Dependence for Composite Inclusions

We now develop variance-reduced anchored forward-backward (VRAF), an anytime fixed-query method for monotone Lipschitz inclusions that attains the optimal $O(\sigma^2/\epsilon)$ dependence on the variance parameter.

As discussed in Section 3, deterministic acceleration of anchored methods relies on an extragradient or optimistic correction that permits a nonvanishing stepsize. In existing stochastic extensions, the resulting stochastic error terms accumulate, and prescribed variance decay is used to control this accumulation while preserving the accelerated deterministic rate (Lee & Kim, 2021; Alcala et al., 2025). To avoid this additional source of error, VRAF instead uses an anchoring-only forward-backward update with a diminishing stepsize (Ryu et al., 2020; Surina et al., 2026; Cai & Zheng, 2026), at the cost of an $O(1/N)$ deterministic term.

VRAF combines this update with a STORM-type variance-reduced recursive estimator (Cutkosky & Orabona, 2019). Under Assumption 2.3, the variance reduction evaluates the same sampled stochastic operator at z_k and z_{k+1}; Assumption 2.4 then gives $\mathbb{E}\|\delta_{\xi_{k+1}}(z_{k+1}) - \delta_{\xi_{k+1}}(z_k)\|^2 \le L_\Delta^2\|z_{k+1} - z_k\|^2$. This gives the following recursion.

For $\alpha_k > 0$ and $\beta_k, \gamma_{k+1} \in (0, 1]$, initialize $v_0 = F_{\xi_0}(z_0)$ and iterate

$$\begin{aligned} z_{k+1} &= J_{\alpha_k A}\big((1-\beta_k)z_k + \beta_k z_0 - \alpha_k v_k\big), \\ v_{k+1} &= F_{\xi_{k+1}}(z_{k+1}) + (1-\gamma_{k+1})\big(v_k - F_{\xi_{k+1}}(z_k)\big). \end{aligned} \tag{VRAF}$$

Theorem 4.1. *Under Assumptions 2.1–2.4, VRAF with* $\alpha_0 = \frac{7}{12\overline{L}}$, $\alpha_{k+1} = \frac{2(k+3)}{2k+7}\alpha_k$, $\beta_k = \frac{3}{k+3}$, *and* $\gamma_{k+1} = \frac{4k+13}{4(k+3)(k+4)}$ *satisfies, for every integer* $N \ge 1$,

$$\mathbb{E}[\mathcal{R}(z_N)^2] \le 35\frac{\overline{L}^2D^2 + \sigma^2}{N+2}.$$

Table 1: Oracle complexity for $\mathbb{E}[\mathcal{R}(z)^2] \le \epsilon$. The Composite column indicates whether the guarantee applies to $A \neq 0$; $\triangle$ denotes constrained problems only. In the Output column, Terminal denotes a nonrandomized terminal output; for RRSEG, this is the corrected output $\widehat{z}_N$ in Theorem 6.1. The notation $\widetilde{O}$ hides logarithmic factors. *Appendix A.1 derives the displayed expected-squared-residual complexity from the proof of Chen & Luo (2024), whose theorem is stated for the expected residual. †The prescribed variance decay is translated into oracle complexity using independent increasing batches under Assumption 2.2. VRAF requires Assumption 2.3 and 2.4; RRSEG requires neither.

Method	Fixed-query	Output	Oracle complexity	Composite
Moving-anchor EAG-V		Terminal	$O\left(\frac{LD}{\sqrt{\epsilon}} + \frac{\sigma^2L^3D^3}{\epsilon^{5/2}}\right)$ †	
S-FEG		Terminal	$O\left(\frac{LD}{\sqrt{\epsilon}} + \frac{\sigma^2L^2D^2}{\epsilon^2}\right)$ †	
BC-(P)SEG+	✓	Randomized	$\widetilde{O}\left(\frac{\overline{L}^4D^4+\sigma^4}{\epsilon^2}\right)$	✓
Stochastic GOMA	✓	Terminal	$O\left(\frac{L^4D^4+\sigma^4}{\epsilon^2}\right)$	
E-Halpern		Terminal	$O\left(\frac{\overline{L}^3D^3+\sigma^2\overline{L}D}{\epsilon^{3/2}}\right)$	$\triangle$
RAIN		Randomized	$\widetilde{O}\left(\frac{LD}{\sqrt{\epsilon}} + \frac{\sigma^2}{\epsilon}\right)$ *	
VRAF (ours)	✓	Terminal	$O\left(\frac{\overline{L}^2D^2+\sigma^2}{\epsilon}\right)$	✓
RRSEG (ours)	✓	Terminal	$\widetilde{O}\left(\frac{LD}{\sqrt{\epsilon}} + \frac{\sigma^2}{\epsilon}\right)$	✓

Remark 4.2. The VRAF stepsize satisfies $\alpha_k = \Theta\big(\frac{1}{\overline{L}\sqrt{k}}\big)$, while $\beta_k = \Theta\big(\frac{1}{k}\big)$. Indeed, since $\sqrt{\frac{k+5/2}{k+7/2}} \le \frac{\alpha_{k+1}}{\alpha_k} = \frac{k+3}{k+7/2} \le \sqrt{\frac{k+3}{k+4}}$, multiplying from $j = 0, \dots, k-1$ gives $\frac{7}{12\overline{L}}\sqrt{\frac{5}{2k+5}} \le \alpha_k \le \frac{7}{12\overline{L}}\sqrt{\frac{3}{k+3}}$.

Each VRAF iteration uses two oracle queries. Hence, Theorem 4.1 gives $O((\overline{L}^2 D^2 + \sigma^2)/\epsilon)$ oracle complexity to achieve $\mathbb{E}[\mathcal{R}(z)^2] \le \epsilon$. The parameters depend only on k and $\overline{L}$, so VRAF is anytime. We provide a proof sketch below; the complete proof is given in Appendix B.

Proof sketch. For $k \ge 1$, let $h_k \in A(z_k)$ be the element selected by the resolvent step of VRAF, so that

$$z_k + \alpha_{k-1} h_k = (1-\beta_{k-1}) z_{k-1} + \beta_{k-1} z_0 - \alpha_{k-1} v_{k-1}.$$

Define the potential

$$\begin{aligned} V_k :=& \frac{1}{2}\mathbb{E}\left\| z_k - z_0 + \frac{\alpha_k}{\beta_k}\big(v_k - F(z_k)\big) + \frac{\alpha_{k-1}}{\beta_{k-1}}\big(F(z_k) + h_k\big)\right\|^2 \\ &+ \frac{1}{2}\mathbb{E}\left\| \frac{\alpha_k}{\beta_k}\big(v_k - F(z_k)\big) - (z_0 - z_\star)\right\|^2. \end{aligned}$$

The proof establishes that, for $k \ge 1$,

$$\begin{aligned} \mathbb{E}[\mathcal{R}(z_k)^2] &\le 4\left(\frac{\beta_{k-1}}{\alpha_{k-1}}\right)^2 V_k, \\ V_{k+1} &\le \left(1 - \frac{2}{5(k+3)}\right) V_k + \frac{2}{5(k+3)}\frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right), \\ V_1 &\le \frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right). \end{aligned}$$

The last two inequalities imply by induction that $V_k \le \frac{9}{10}\big(D^2 + \frac{\sigma^2}{\overline{L}^2}\big)$ for all $k \ge 1$. Combining this bound with the first inequality gives

$$\mathbb{E}[\mathcal{R}(z_k)^2] \le 4\left(\frac{\beta_{k-1}}{\alpha_{k-1}}\right)^2 \frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right).$$

Finally, the parameter choices satisfy $\big(\frac{\beta_{k-1}}{\alpha_{k-1}}\big)^2 \le \frac{175\overline{L}^2}{18(k+2)}$. Setting $k = N$ proves the theorem. □

The $O(L^2 D^2/N)$ deterministic term in Theorem 4.1 is order-tight for the VRAF. This already holds for the one-dimensional linear operator $F(z) = Lz$ with $A = 0$; see Appendix B.4. Hence, the VRAF cannot attain the optimal $O(L^2 D^2/N^2)$ deterministic rate.

5 Impossibility of achieving optimal deterministic and stochastic rates simultaneously

Having established the optimal $O(\sigma^2/\epsilon)$ variance dependence with VRAF, we next ask whether it can be attained together with the optimal $O(LD/\sqrt{\epsilon})$ deterministic rate. We answer this question negatively with the following lower bound theorem.

Theorem 5.1 (Impossibility of simultaneous optimality)**.** *Under Assumptions 2.1, 2.2, and 2.4, even when restricted to two-dimensional monotone Lipschitz equations with $A = 0$ and $L_\Delta \le L$, it is impossible to uniformly guarantee $\mathbb{E}[\mathcal{R}(\widehat{z})^2] \le \epsilon$ using $O\big(\frac{LD}{\sqrt{\epsilon}} + \frac{\sigma^2}{\epsilon}\big)$ oracle queries. This impossibility holds even when previously sampled stochastic operators may be queried repeatedly.*

We provide a proof sketch below; the complete proof is given in Appendix C.

Proof sketch. Let F_Z be a rescaled version of the worst-case family of Foster et al. (2019). Appendix C extends this family so that $L = L_\Delta = 3$, $D = 1$, and $\sigma^2 = \frac{1}{M^2}$ for $M \geq 16$. At $\epsilon_M = \frac{1}{256M^4}$, by the lower bound construction of Foster et al. (2019) and this extension, achieving $\mathbb{E}[\mathcal{R}(\widehat{z})^2] \leq \epsilon_M$ requires $\Omega(M^2 \log M)$ oracle queries, while $LD/\sqrt{\epsilon_M} + \sigma^2/\epsilon_M = 304M^2$. □

For fixed-query methods, the $\Omega(M^2 \log M)$ lower bound readily translates into a corresponding lower bound on the iteration count N as below.

Corollary 5.2 (Necessity of a logarithmic factor for fixed-query methods)**.** *For every $0 \leq p < 1$, no fixed-query method can uniformly guarantee, for all integers $N \geq 2$,*

$$\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2] = O\left(\frac{L^2D^2}{N^2} + \frac{\sigma^2 \log^p N}{N}\right).$$

The proof is given in Appendix C. Corollary 5.2 leaves the range $1 \leq p < 3$ open: it rules out every logarithmic power $p < 1$, whereas the near-optimal bounds of both RAIN and RRSEG have stochastic logarithmic power $p = 3$; RRSEG achieves this dependence with a fixed number of oracle queries per iteration.

6 RRSEG: Optimal deterministic and near-optimal stochastic rates

The impossibility result shows that $O(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ oracle complexity is unattainable in general. We therefore ask whether a fixed-query method can retain the optimal deterministic rate while losing only logarithmic factors in the stochastic rate.

RAIN achieved the near-optimal rate using recursive regularization (Chen & Luo, 2024). Let $H_k(z) := F(z) + a_k(z - c_k)$ and let r_k be the unique zero of $A + H_k$, with $c_0 = z_0$. Since $\mathcal{R}(r_0) \leq a_0 D$, choosing a small a_0 gives a small residual for the original inclusion but also a small strong-monotonicity constant.[2] If r_k were available, adding $\lambda_k(z - r_k)$ to H_k would increase the strong-monotonicity constant while preserving r_k as the zero of the regularized inclusion.

However, RAIN implements this recursive regularization by solving each regularized subproblem with a stochastic subroutine (Chen & Luo, 2024). Chen & Luo (2024) also consider a single-loop variant with one stochastic extragradient step per update, but without providing a convergence rate analysis. Building on this single-loop structure, RRSEG applies one stochastic forward-backward-forward step to each regularized inclusion and then recenters the added regularization at the resulting iterate. When $A = 0$, this step reduces to stochastic extragradient. With horizon-dependent parameters, RRSEG retains the optimal deterministic rate while incurring only logarithmic factors in the stochastic rate. Unlike VRAF, its analysis uses only Assumptions 2.1 and 2.2; in particular, it requires neither Assumption 2.3 nor Assumption 2.4.

Since r_k is unavailable, RRSEG replaces it with the one-step approximation z_{k+1}:

$$H_{k+1}(z) = H_k(z) + \lambda_k(z - z_{k+1}).$$

This gives the following single-loop recursion.

For $a_0 > 0$, $c_0 = z_0$, $\eta_k > 0$, and $\lambda_k \geq 0$, draw independent samples ξ_k and ζ_k and iterate

$$\begin{aligned} y_k &= J_{\eta_k A}\big(z_k - \eta_k\big(F_{\xi_k}(z_k) + a_k(z_k - c_k)\big)\big), \\ z_{k+1} &= y_k - \eta_k\big(F_{\zeta_k}(y_k) + a_k(y_k - c_k) - F_{\xi_k}(z_k) - a_k(z_k - c_k)\big), \\ a_{k+1} &= a_k + \lambda_k, \\ c_{k+1} &= \frac{a_k c_k + \lambda_k z_{k+1}}{a_{k+1}}. \end{aligned} \tag{RRSEG}$$

For $a_0 \in (0, L)$ and $0 < \gamma \leq \frac{1}{3}$, define $Q_{N,\gamma}(a_0) := \frac{\frac{\gamma N}{3} - (L/a_0 - 1)}{\log(L/a_0)}$. The theorem below gives the exact finite-horizon bound. In particular, the concrete parameter choices in Corollary 6.2 yields the simplified bound $O\big(\frac{L^2D^2}{N^2} + \frac{\sigma^2 \log^3 N}{N}\big)$.

[2] An operator H is μ-strongly monotone if $\langle H(x) - H(y), x - y\rangle \geq \mu\|x - y\|^2$ for all x, y.

Theorem 6.1 (RRSEG). *Under Assumptions 2.1 and 2.2, fix an integer $N \geq 2$, $a_0 \in (0, L)$, and $0 < \gamma \leq \frac{1}{3}$ such that $Q_{N,\gamma}(a_0) \geq 1$. Choose the RRSEG parameters as $\eta_k = \frac{1}{3(L+Q_{N,\gamma}(a_0)a_k)}$, $a_{k+1} = \min\{L, \frac{a_k}{1-\gamma\eta_k a_k}\}$, and $\lambda_k := a_{k+1} - a_k$. Then $a_N = L$. Let $m := \lfloor \min\{Q_{N,\gamma}(a_0), N\} \rfloor$, $\overline{F}_N := \frac{1}{m}\sum_{k=N-m}^{N-1} F_{\xi_k}(z_k)$, and define the output $\widehat{z}_N := J_{\frac{1}{L}A}\big(z_N - \frac{1}{L}\overline{F}_N\big)$. Then RRSEG satisfies*

$$\begin{aligned}\big(\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2]\big)^{1/2} \leq& \sqrt{2}LD\left[2\frac{a_0}{L} + \left(\sqrt{2} - 1 + \frac{77}{32}\right)\left(\frac{a_0}{L}\right)^2\right] + \frac{\sqrt{2}\sigma}{\sqrt{m}} \\ &+ \sigma\sqrt{\frac{5}{2(2-3\gamma)(Q_{N,\gamma}(a_0)+1)}}\left(\log\frac{L}{a_0} + \sqrt{2} + \frac{9}{4}\right). \end{aligned} \tag{2}$$

The average $\overline{F}_N$ uses oracle evaluations already generated by RRSEG, so computing $\widehat{z}_N$ requires no additional stochastic oracle queries.

We provide a proof sketch below; the complete proof is given in Appendix D.

Proof sketch. For each k, let r_k be the zero of $A + H_k$. The proof establishes the following bounds on the distances from the iterates to these zeros:

$$\begin{aligned}\mathbb{E}\|z_k - r_k\|^2 &\leq \left(\frac{a_0}{a_k}\right)^4 D^2 + \frac{5\sigma^2}{4(2-3\gamma)(Q_{N,\gamma}(a_0)+1)a_k^2}, \\ \mathbb{E}\|z_{k+1} - r_k\|^2 &\leq \left(\frac{a_0}{a_{k+1}}\right)^4 D^2 + \frac{5\sigma^2}{4(2-3\gamma)(Q_{N,\gamma}(a_0)+1)a_{k+1}^2}.\end{aligned}$$

With $p_N := J_{\frac{1}{L}A}\left(z_N - \frac{1}{L}F(z_N)\right)$, the recentering relation and firm nonexpansiveness of the resolvent further give

$$L\left(\mathbb{E}\|z_N - p_N\|^2\right)^{1/2} \leq a_0 D + \sum_{k=0}^{N-1}\lambda_k\left(\mathbb{E}\|z_{k+1} - r_k\|^2\right)^{1/2} + \sqrt{2}L\left(\mathbb{E}\|z_N - r_N\|^2\right)^{1/2}.$$

The choice of $Q_{N,\gamma}(a_0)$ gives $a_N = L$. Then, using $\sqrt{u+v} \leq \sqrt{u} + \sqrt{v}$, the above inequalities give

$$\begin{aligned}&L\left(\mathbb{E}\|z_N - p_N\|^2\right)^{1/2} \\ &\leq \left(a_0 + a_0^2\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}^2} + \sqrt{2}\frac{a_0^2}{L}\right)D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q_{N,\gamma}(a_0)+1)}}\left(\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}} + \sqrt{2}\right) \\ &\leq \left(2a_0 + (\sqrt{2}-1)\frac{a_0^2}{L}\right)D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q_{N,\gamma}(a_0)+1)}}\left(\log\frac{L}{a_0} + \sqrt{2}\right), \end{aligned} \tag{3}$$

where the last inequality follows from

$$\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}^2} \leq \int_{a_0}^{L}\frac{da}{a^2} = \frac{1}{a_0} - \frac{1}{L}, \qquad \sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}} \leq \int_{a_0}^{L}\frac{da}{a} = \log\frac{L}{a_0}.$$

The proof also gives

$$\left(\mathbb{E}\|\overline{F}_N - F(z_N)\|^2\right)^{1/2} \leq \frac{77}{32}\frac{a_0^2}{L}D + \frac{\sigma}{\sqrt{m}} + \frac{9}{4}\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q_{N,\gamma}(a_0)+1)}}.$$

Finally,

$$\left(\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2]\right)^{1/2} \leq \sqrt{2}L\left(\mathbb{E}\|z_N - p_N\|^2\right)^{1/2} + \sqrt{2}\left(\mathbb{E}\|\overline{F}_N - F(z_N)\|^2\right)^{1/2}.$$

Substituting the preceding bounds gives equation 2; the complete calculation is given in Appendix D. □

Corollary 6.2 (RRSEG with concrete parameter choices)**.** *Under the assumptions of Theorem 6.1, let $N \geq 19$ and set $\gamma = \frac{1}{3}$, $a_0 = \frac{18L}{N}$, $Q_N = \frac{N/18+1}{\log(N/18)}$, and $m_N = \lfloor \min\{Q_N, N\} \rfloor$. Then $Q_N > 1$, and RRSEG satisfies*

$$\left(\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2]\right)^{1/2} \leq \frac{36\sqrt{2}LD}{N} + \frac{324\sqrt{2}LD}{N^2}\left(\sqrt{2} - 1 + \frac{77}{32}\right) + \frac{\sqrt{2}\sigma}{\sqrt{m_N}} + \sigma\sqrt{\frac{5}{2(Q_N+1)}}\left(\log\frac{N}{18} + \sqrt{2} + \frac{9}{4}\right).$$

Consequently, $\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2] = O\big(\frac{L^2D^2}{N^2} + \frac{\sigma^2 \log^3 N}{N}\big)$.

Proof. For $N \geq 19$, the displayed choice satisfies $a_0 \in (0, L)$ and $Q_N > 1$; Appendix D gives the elementary verification of the latter inequality. Substituting $\gamma = \frac{1}{3}$ and $a_0 = \frac{18L}{N}$ into Theorem 6.1 gives the upper bound and the stated order. □

Corollary 6.2 is consistent with Theorem 5.1, since the stochastic term is optimal only up to logarithmic factors. Repeating RRSEG from the same initial point with doubling horizons yields an anytime method without changing the asymptotic rate; see Appendix F.1.

When $A = 0$, $p_N = z_N - \frac{1}{L}F(z_N)$ and hence $L(z_N - p_N) = F(z_N)$. Therefore, equation 3 directly bounds $(\mathbb{E}\|F(z_N)\|^2)^{1/2}$, so the last iterate z_N attains the same oracle-complexity order without the final output correction.

Appendix D.8 illustrates the horizon-dependent schedule used in Corollary 6.2 and in the experiments.

7 Experiments

Figure 1 compares VRAF (Section 4), RRSEG (Section 6), BC-(P)SEG+ (Pethick et al., 2023), stochastic GOMA (Sohrabi et al., 2026), RAIN-SL (Chen & Luo, 2024), and E-Halpern (Cai et al., 2022). Problem 3 has $A \neq 0$ but is not a constrained problem, so only VRAF, RRSEG, and BC-PSEG+ are evaluated. Each uses at most 10^7 oracle queries. VRAF and RRSEG use their theoretical parameters without empirical tuning. VRAF (practical) replaces α_k by $4\alpha_k$; it is heuristic, and Appendix E.3 compares multipliers $1, 2, 4$, and 8. RAIN-SL is the tuned Algorithm 9 of Chen & Luo (2024). E-Halpern is calibrated through its target error ϵ. Appendix E gives the remaining details.

A single RRSEG run is horizon-dependent, and E-Halpern uses a prescribed target accuracy. Their solid curves show the corresponding intermediate iterates. For RRSEG, stars show terminal outputs from separate horizon-matched runs; on Problem 3, they show the corrected outputs $\widehat{z}_N$. The RRSEG dashed curve is horizon-matched, using equation 63 on Problems 1 and 2 and Corollary 6.2 on Problem 3. For RAIN-SL, NaN marks the first plotted query count at which at least one of the 50 runs becomes nonfinite.

Problems 1-3 are monotone, with $A \neq 0$ in Problem 3. For rock-paper-scissors (RPS), we parameterize the mixed strategies by softmax logits. The resulting operator is nonmonotone, so we include this problem to examine the behavior of the methods outside the monotone setting. Appendix E.4 gives the exact problems.

The RRSEG curves decrease most rapidly near their terminal horizons. RAIN-SL decreases rapidly at early query counts, but at least one of its 50 runs becomes nonfinite on every problem. For RPS, softmax saturation can make the residual in the logit variables small without making the saddle-point gap small. Appendix E.5 therefore also reports the mixed-strategy saddle-point gap, which decreases together with the squared residual for RRSEG and both VRAF variants.

8 Discussion

The main open question left by Theorem 5.1 is the tradeoff between the deterministic and stochastic terms. In particular, it remains open whether the exact $O(\sigma^2/N)$ stochastic term can be attained while improving the deterministic term $O(1/N)$, and whether the logarithmic power in the

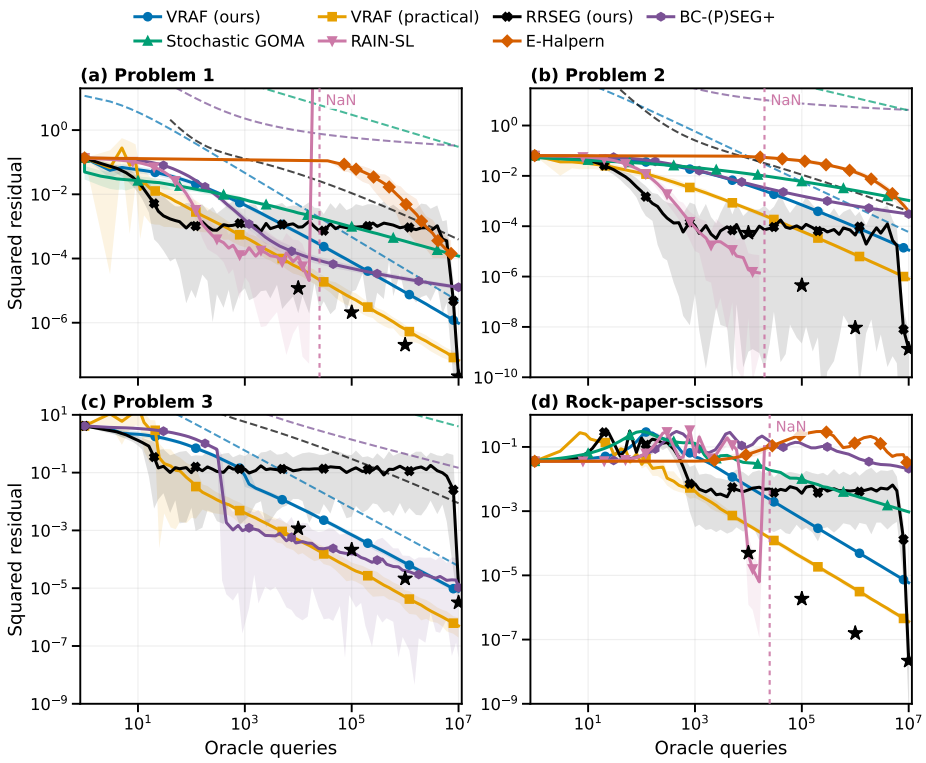

Figure 1: Squared residual versus the number of oracle queries. Solid curves show the mean over 50 independent runs, and shaded regions show the pointwise 2.5%–97.5% quantiles. Dashed curves show the theoretical bounds described in the text. RRSEG solid curves show intermediate outputs from a single horizon-dependent run, whereas stars show terminal outputs from separate horizon-matched runs.

$O(\sigma^2 \log^3 N/N)$ stochastic term can be reduced toward the lower-bound threshold while preserving the optimal $O(1/N^2)$ deterministic term under fixed-query access.

Our guarantees also rely on monotonicity. Although deterministic acceleration extends to negative comonotonicity (Lee & Kim, 2021; Gorbunov et al., 2023), corresponding stochastic guarantees are unknown; the RPS experiment therefore lies outside our theory.

Finally, residual convergence does not imply pointwise convergence. Appendix F gives a counterexample for VRAF, and the RRSEG guarantees do not address pointwise convergence.

9 Conclusion

We studied fixed-query acceleration for stochastic monotone Lipschitz inclusions under the expected squared residual criterion. VRAF is the first method to attain the optimal $O(\sigma^2/\epsilon)$ variance dependence, even among methods allowing an increasing number of oracle queries per iteration, while itself using only a fixed number of queries. We also established a lower bound showing that the optimal deterministic and stochastic rates cannot, in general, be attained simultaneously, even when $L_\Delta \le L$ and repeated queries to previously sampled stochastic operators are allowed. RRSEG attains the near-optimal $\widetilde{O}(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ oracle complexity previously achieved with increasing oracle queries, while using only a fixed number of oracle queries per iteration.

AI USE STATEMENT

In this work, we used generative AI tools to assist with algorithmic exploration, mathematical analysis and proof development, numerical experiments, literature search, code development, translation, and manuscript editing. The research objectives and major methodological choices were determined by the authors. All AI-assisted work was reviewed and independently checked by the authors, who take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

The appendix contains complete proofs of the VRAF, RRSEG, and pointwise results, together with the lower-bound transfer. For the experiments, the appendix specifies the exact problems, schedules, tuning procedures, and evaluation protocol. The supplementary material provides the complete experimental code and the VRAF proof-verification code.

REFERENCES

Ahmet Alacaoglu and Yura Malitsky. Stochastic variance reduction for variational inequality methods. In *Conference on Learning Theory*, pp. 778–816. PMLR, 2022.

James Alcala, Yat Tin Chow, and Mahesh Sunkula. Stochastic Moving Anchor Algorithms and a Popov's Scheme with Moving Anchor. *arXiv preprint arXiv:2506.07290*, 2025.

James K Alcala, Yat Tin Chow, and Mahesh Sunkula. Moving anchor extragradient methods for smooth structured minimax problems. *arXiv preprint arXiv:2308.12359*, 2023.

Zeyuan Allen-Zhu. How to make the gradients small stochastically: Even faster convex and nonconvex sgd. *Advances in Neural Information Processing Systems*, 31, 2018.

Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*, pp. 214–223. Pmlr, 2017.

Raef Bassily, Cristóbal Guzmán, and Michael Menart. Private algorithms for stochastic saddle points and variational inequalities: Beyond euclidean geometry. *Advances in Neural Information Processing Systems*, 37:128603–128635, 2024.

Radu I. Boţ and Enis Chenchene. Extragradient Method with Flexible Anchoring: Strong Convergence and Fast Residual Decay. *SIAM Journal on Optimization*, 36(3):1420–1445, 2026.

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, 2022.

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

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. Accelerated Algorithms for Constrained Nonconvex-Nonconcave Min-Max Optimization and Comonotone Inclusion. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 5312–5347. PMLR, 2024.

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.

Zihao Chen. A Unifying View of Anchoring via Operator-Side Tikhonov Regularization. *arXiv preprint arXiv:2605.30905*, 2026.

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

Jelena Diakonikolas. Halpern iteration for near-optimal and parameter-free monotone inclusion and strong solutions to variational inequalities. In *Conference on Learning Theory*, pp. 1428–1451. PMLR, 2020.

Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient Methods for Structured Nonconvex-Nonconcave Min-Max Optimization. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, pp. 2746–2754. PMLR, 2021.

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, 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, 2020.

Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.

Eduard Gorbunov, Nicolas Loizou, and Gauthier Gidel. Extragradient method: O (1/k) last-iterate convergence for monotone variational inequalities and connections with cocoercivity. In *International Conference on Artificial Intelligence and Statistics*, pp. 366–402. PMLR, 2022a.

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, 2022b.

Eduard Gorbunov, Adrien Taylor, Samuel Horváth, and Gauthier Gidel. Convergence of proximal point and extragradient-based methods beyond monotonicity: The case of negative comonotonicity. In *International Conference on Machine Learning*, pp. 11614–11641. PMLR, 2023.

Anatoli Juditsky, Arkadi Nemirovski, and Claire Tauvel. Solving Variational Inequalities with Stochastic Mirror-Prox Algorithm. *Stochastic Systems*, 1(1):17–58, 2011.

Donghwan Kim. Accelerated proximal point method for maximally monotone operators. *Mathematical Programming*, 190(1-2):57–87, 2021.

Jun-Hyun Kim and Ahmet Alacaoglu. Improving the Last-Iterate Guarantees of Anytime Algorithms for Stochastic Monotone Variational Inequalities. *arXiv preprint arXiv:2609.15257*, 2026.

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

Guanghui Lan. An optimal method for stochastic composite optimization. *Mathematical Programming*, 133(1):365–397, 2012.

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.

Sucheol Lee and Donghwan Kim. Semi-Anchored Multi-Step Gradient Descent Ascent Method for Structured Nonconvex-Nonconcave Composite Minimax Problems. *arXiv preprint arXiv:2105.15042*, 2022.

Chris Junchi Li, Wenlong Mou, Martin Wainwright, and Michael Jordan. Root-sgd: Sharp nonasymptotics and asymptotic efficiency in a single algorithm. In *Conference on Learning Theory*, pp. 909–981. PMLR, 2022.

Zhize Li, Hongyan Bao, Xiangliang Zhang, and Peter Richtarik. PAGE: A Simple and Optimal Probabilistic Gradient Estimator for Nonconvex Optimization. In *Proceedings of the 38th International Conference on Machine Learning*, pp. 6286–6295. PMLR, 2021.

Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine Learning Proceedings 1994*, pp. 157–163. Elsevier, 1994.

Eric Mazumdar, Kishan Panaganti, and Laixi Shi. Tractable multi-agent reinforcement learning through behavioral economics. In *International Conference on Learning Representations*, volume 2025, pp. 18058–18072, 2025.

Remi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegel, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J. Mankowitz, Doina Precup, and Bilal Piot. Nash Learning from Human Feedback. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 36743–36768. PMLR, 2024.

Yuyuan Ouyang and Yangyang Xu. Lower complexity bounds of first-order methods for convex-concave bilinear saddle-point problems. *Mathematical Programming*, 185(1-2):1–35, 2021.

Jisun Park and Ernest K. Ryu. Exact optimal accelerated complexity for fixed-point iterations. In *International Conference on Machine Learning*, pp. 17420–17457. PMLR, 2022.

Thomas Pethick, Puya Latafat, Panos Patrinos, Olivier Fercoq, and Volkan Cevher. Escaping limit cycles: Global convergence for constrained nonconvex-nonconcave minimax problems. In *International Conference on Learning Representations*, 2022.

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

L. D. 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.

J. Ben Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica: Journal of the Econometric Society*, pp. 520–534, 1965.

Ernest K. Ryu, Kun Yuan, and Wotao Yin. ODE Analysis of Stochastic Gradient Methods with Optimism and Anchoring for Minimax Problems. *arXiv preprint arXiv:1905.10899*, 2020.

Motahareh Sohrabi, Jianxin You, Simon Lacoste-Julien, Eduard Gorbunov, and Gauthier Gidel. Accelerated and Stable Convergence with Anchored Generalized Optimistic Method. In *Forty-Third International Conference on Machine Learning*, 2026.

Chaobing Song, Zhengyuan Zhou, Yichao Zhou, Yong Jiang, and Yi Ma. Optimistic Dual Extrapolation for Coherent Non-monotone Variational Inequalities. *Advances in Neural Information Processing Systems*, 33:14303–14314, 2020.

Jaewook Suh, Jisun Park, and Ernest Ryu. Continuous-time analysis of anchor acceleration. *Advances in Neural Information Processing Systems*, 36:32782–32866, 2023.

Anja Surina, Arun Suggala, George Tsoukalas, Anton Kovsharov, Sergey Shirobokov, Francisco J. R. 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 Yang Luo. Halpern-Type Accelerated and Splitting Algorithms For Monotone Inclusions. *arXiv preprint arXiv:2110.08150*, 2021.

Paul Tseng. A Modified Forward-Backward Splitting Method for Maximal Monotone Mappings. *SIAM Journal on Control and Optimization*, 38(2):431–446, 2000.

Zaiyan Xu, Sushil Vemuri, Kishan Panaganti, Dileep Kalathil, Rahul Jain, and Deepak Ramachandran. Robust LLM alignment via distributionally robust direct preference optimization. In *The Thirty-Ninth Annual Conference on Neural Information Processing Systems*, 2025.

TaeHo Yoon and Nicolas Loizou. Direct Acceleration of Stochastic Root-Finding Without Variance Reduction and Regularization. *arXiv preprint arXiv:2608.12043*, 2026.

TaeHo Yoon and Ernest K. Ryu. Accelerated algorithms for smooth convex-concave minimax problems with O (1/kˆ 2) rate on squared gradient norm. In *International Conference on Machine Learning*, pp. 12098–12109. PMLR, 2021.

TaeHo Yoon and Ernest K. Ryu. Accelerated Minimax Algorithms Flock Together. *SIAM Journal on Optimization*, 35(1):180–209, 2025.

A Further related work

Gap lower bounds. First-order lower complexity bounds are known for convex-concave bilinear saddle-point problems under saddle-point criteria (Ouyang & Xu, 2021). For smooth convex-concave saddle-point problems, Golowich et al. (2020) prove an $O(N^{-1/2})$ last-iterate rate for extragradient and matching lower bounds for stationary linear iterative methods, for criteria including the restricted primal-dual gap and the Hamiltonian. These criteria differ from the tangent residual used in this paper, so these results do not imply the lower bound in Section 5.

Beyond monotonicity. Several first-order methods retain convergence guarantees under assumptions weaker than monotonicity. EG+ is analyzed for structured nonconvex-nonconcave problems under a weak Minty-type condition (Diakonikolas et al., 2021), and Pethick et al. (2022) extend extragradient-type methods to constrained and regularized weak-Minty problems. FEG attains accelerated residual convergence under negative comonotonicity (Lee & Kim, 2021), while proximal-point, extragradient, and optimistic methods have also been analyzed under negative comonotonicity (Gorbunov et al., 2023). Optimistic dual extrapolation and semi-anchored schemes provide further guarantees for coherent nonmonotone variational inequalities and structured nonconvex-nonconcave composite minimax problems (Song et al., 2020; Lee & Kim, 2022).

Implicit and strong-convergence viewpoints. Anchoring is closely connected to Halpern iteration and proximal-point methods. Halpern-based methods give near-optimal guarantees for monotone inclusions (Diakonikolas, 2020), while accelerated proximal-point and Halpern-type splitting methods give fast residual guarantees for maximally monotone operators and monotone inclusions (Kim, 2021; Tran-Dinh & Luo, 2021). For nonexpansive and contractive fixed-point problems, Park & Ryu (2022) derive exact optimal accelerated methods. Flexible anchoring has also been used to obtain strong convergence together with fast residual decay (Boţ & Chenchene, 2026), and Yoon & Ryu (2025) identify common trajectory structure among several accelerated minimax and fixed-point methods.

Continuous-time and structural interpretations of anchoring. Continuous-time models have been used to explain optimism and anchoring in minimax dynamics (Ryu et al., 2020) and to characterize the rate produced by general anchor schedules (Suh et al., 2023). Flexible anchored extragradient can be derived by discretizing a Tikhonov-regularized monotone flow (Boţ & Chenchene, 2026), while operator-side Tikhonov regularization gives a common construction for Halpern, forward, extragradient, and Popov-type methods (Chen, 2026).

Variance-reduction mechanisms. Recursive stochastic estimators have a broad literature outside monotone problems. STORM and PAGE are representative recursive estimators for stochastic nonconvex optimization (Cutkosky & Orabona, 2019; Li et al., 2021), and ROOT-SGD uses recursive averaging under a stochastic Lipschitz condition on the gradient noise (Li et al., 2022). For variational inequalities and monotone inclusions, variance reduction has been developed directly for extragradient and splitting schemes (Alacaoglu & Malitsky, 2022) and for Halpern iteration (Cai et al., 2022).

A.1 RAIN under the expected squared residual

Table 1 uses the criterion $\mathbb{E}\|F(z)\|^2 \le \epsilon$, while the RAIN theorems of Chen & Luo (2024) are stated for $\mathbb{E}\|F(z)\|$. The required second-moment rate follows from the same proof.

Proposition A.1. *For a monotone L-Lipschitz equation with $\|z_0 - z_\star\| \le D$, RAIN attains*

$$\mathbb{E}\|F(z_S)\|^2 \le \epsilon$$

with $\widetilde{O}(LD/\sqrt{\epsilon}+\sigma^2/\epsilon)$ oracle queries. In particular, the stochastic contribution in the oracle bound is $O\left(\frac{\sigma^2}{\epsilon}\log^3\left(1+\frac{LD}{\sqrt{\epsilon}}\right)\right)$.

Proof. First consider the strongly monotone subproblem used by RAIN. Lemma 3.1 and equation (17) in the proof of Theorem 4.1 of Chen & Luo (2024) give, for a target $\delta > 0$,

$$\|G(z_S)\| \le 16\sum_{s=0}^{S-1}\lambda_s\|e_s\|, \qquad 256\lambda_s^2 S^2\mathbb{E}\|e_s\|^2 \le \delta^2.$$

Hence Cauchy-Schwarz gives $\mathbb{E}\|G(z_S)\|^2 \le \delta^2$. For a general monotone equation, RAIN applies this result to $G(z) = F(z)+\lambda(z-z_0)$. Lemma 3.2 of Chen & Luo (2024) gives $\|F(z)\| \le 2\|G(z)\|+\lambda D$. Choosing the internal tolerance and λD as constant multiples of $\sqrt{\epsilon}$ therefore yields $\mathbb{E}\|F(z_S)\|^2 \le \epsilon$, with only constant changes to the parameters. Substituting $\delta = \Theta(\sqrt{\epsilon})$ and $\lambda = \Theta(\sqrt{\epsilon}/D)$ into the oracle count in Theorems 4.1-4.2 of Chen & Luo (2024) gives the stated order. □

B Complete VRAF Proof

B.1 Potential and residual extraction

Let $\mathcal{F}_k := \sigma(\xi_0,\dots,\xi_k)$ and $\mathbb{E}_k[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_k]$. For $k \ge 1$, let $h_k \in A(z_k)$ be the element selected by the resolvent at iteration $k-1$, so that

$$z_k + \alpha_{k-1}h_k = (1-\beta_{k-1})z_{k-1} + \beta_{k-1}z_0 - \alpha_{k-1}v_{k-1}.$$

Define

$$T_k := \frac{\alpha_k}{\beta_k}, \qquad u_k := T_k\big(v_k - F(z_k)\big) \quad (k \ge 0),$$
$$p_k := T_{k-1}\big(F(z_k)+h_k\big) \quad (k \ge 1), \qquad d := z_0 - z_\star,$$

and, for $k \ge 1$,

$$V_k := \frac{1}{2}\mathbb{E}\|z_k - z_0 + u_k + p_k\|^2 + \frac{1}{2}\mathbb{E}\|u_k - d\|^2.$$

Since $p_k/T_{k-1} = F(z_k)+h_k \in (F+A)(z_k)$ and $0 \in (F+A)(z_\star)$, monotonicity gives

$$\|z_k - z_\star + p_k\|^2 - \|p_k\|^2 = \|z_k - z_\star\|^2 + 2\langle p_k, z_k - z_\star\rangle \ge 0.$$

Hence

$$\begin{aligned}\mathbb{E}[\mathcal{R}(z_k)^2] &\le \frac{1}{T_{k-1}^2}\mathbb{E}\|p_k\|^2 \\ &\le \frac{1}{T_{k-1}^2}\mathbb{E}\|z_k - z_\star + p_k\|^2 \\ &= \frac{1}{T_{k-1}^2}\mathbb{E}\big\|(z_k - z_0 + u_k + p_k) - (u_k - d)\big\|^2 \\ &\le \frac{4V_k}{T_{k-1}^2}. \end{aligned} \tag{4}$$

We first prove the uniform bound

$$V_k \le \frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right), \qquad k \ge 1. \tag{5}$$

B.2 One-step bound

The exact algebraic identities and coefficient-positivity checks in this proof are independently verified by the symbolic proof checker provided in the supplementary material; all checks use exact arithmetic.

Proposition B.1. *For every $k \geq 1$,*

$$V_{k+1} \leq \left(1 - \frac{2}{5(k+3)}\right) V_k + \frac{2}{5(k+3)} \frac{9}{10} \left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right). \tag{6}$$

Proof. Fix $k \geq 1$ and set $t := k + 3 \geq 4$. Since $\beta_k = 3/t$ and $T_k = \alpha_k/\beta_k$, the schedule gives

$$\begin{aligned}
\frac{T_{k-1}}{T_k} &= \frac{\alpha_{k-1}}{\alpha_k} \frac{\beta_k}{\beta_{k-1}} = \frac{2t-1}{2(t-1)} \frac{t-1}{t} = \frac{2t-1}{2t}, \\
\frac{T_{k+1}}{T_k} &= \frac{\alpha_{k+1}}{\alpha_k} \frac{\beta_k}{\beta_{k+1}} = \frac{2t}{2t+1} \frac{t+1}{t} = \frac{2(t+1)}{2t+1}, \\
(1-\gamma_{k+1}) \frac{T_{k+1}}{T_k} &= \left(1 - \frac{4t+1}{4t(t+1)}\right) \frac{2(t+1)}{2t+1} = \frac{2t-1}{2t}.
\end{aligned}$$

Set

$$q := \frac{2t-1}{2t}, \qquad \rho := 1 - \frac{2}{5t}.$$

We will also use

$$\overline{L}^2 \alpha_k^2 \leq \frac{1}{k+3}, \qquad k \geq 1. \tag{7}$$

Indeed, $\alpha_1 = 1/(2\overline{L})$, so the claim holds for $k = 1$. If it holds at k, then

$$\overline{L}^2 \alpha_{k+1}^2 \leq \frac{1}{t} \left(\frac{2t}{2t+1}\right)^2 = \frac{4t}{(2t+1)^2} < \frac{1}{t+1}.$$

Set

$$s := \frac{z_k - z_{k+1}}{\beta_k}, \qquad g := T_k\big(F(z_{k+1}) - F(z_k)\big), \qquad w := z_k - z_0 + u_k + p_k.$$

The resolvent relation at iteration k gives

$$\begin{aligned}
z_{k+1} + \alpha_k h_{k+1} &= (1-\beta_k) z_k + \beta_k z_0 - \alpha_k v_k, \\
s &= z_k - z_0 + T_k(v_k + h_{k+1}) \\
&= z_k - z_0 + u_k + p_{k+1} - g.
\end{aligned}$$

Subtracting $F(z_{k+1})$ from the definition of v_{k+1} and multiplying by T_{k+1} gives

$$\begin{aligned}
u_{k+1} =& q u_k + q T_k\big(\delta_{\xi_{k+1}}(z_{k+1}) - \delta_{\xi_{k+1}}(z_k)\big) \\
&+ (T_{k+1} - q T_k) \delta_{\xi_{k+1}}(z_{k+1}) \\
=& q u_k + \omega_{k+1},
\end{aligned} \tag{8}$$

where

$$\omega_{k+1} := q T_k\big(\delta_{\xi_{k+1}}(z_{k+1}) - \delta_{\xi_{k+1}}(z_k)\big) + (T_{k+1} - q T_k) \delta_{\xi_{k+1}}(z_{k+1}).$$

Both z_k and z_{k+1} are $\mathcal{F}_k$-measurable, so unbiasedness gives

$$\mathbb{E}_k \omega_{k+1} = 0.$$

Since $\mathbb{E} u_0 = T_0 \mathbb{E} \delta_{\xi_0}(z_0) = 0$, equation 8 also gives

$$\mathbb{E} u_k = 0, \qquad \mathbb{E}\langle u_k, d \rangle = 0, \qquad k \geq 0. \tag{9}$$

Young's inequality with $\varepsilon = 9/20$, together with Assumptions 2.2 and 2.4, gives

$$\begin{aligned}
\mathbb{E}_k \|\omega_{k+1}\|^2 &\leq \frac{29}{20} q^2 T_k^2 L_\Delta^2 \|z_{k+1} - z_k\|^2 + \frac{29}{9} (T_{k+1} - q T_k)^2 \sigma^2 \\
&= \frac{29}{20} q^2 \alpha_k^2 L_\Delta^2 \|s\|^2 + \frac{29}{9} (T_{k+1} - q T_k)^2 \sigma^2.
\end{aligned} \tag{10}$$

Using $z_{k+1} = z_k - \beta_k s$, equation 8, and $z_k - z_0 + p_{k+1} = s - u_k + g$,

$$z_{k+1} - z_0 + u_{k+1} + p_{k+1} = (1-\beta_k)s + (q-1)u_k + g + \omega_{k+1},$$
$$u_{k+1} - d = qu_k - d + \omega_{k+1}.$$

Substituting these identities into V_{k+1}, using $\mathbb{E}_k\omega_{k+1} = 0$, and subtracting ρV_k gives

$$\begin{aligned} V_{k+1} - \rho V_k =& \frac{1}{2}\mathbb{E}\|(1-\beta_k)s + (q-1)u_k + g\|^2 - \frac{\rho}{2}\mathbb{E}\|w\|^2 \\ &+ \frac{q^2-\rho}{2}\mathbb{E}\|u_k\|^2 + \frac{1-\rho}{2}\|d\|^2 + \mathbb{E}\|\omega_{k+1}\|^2, \end{aligned} \tag{11}$$

where equation 9 removes the cross term with d.

Monotonicity of $F + A$ between z_k and z_{k+1} gives

$$\begin{aligned} 0 &\le \langle F(z_k) + h_k - F(z_{k+1}) - h_{k+1}, z_k - z_{k+1}\rangle \\ &= \frac{\beta_k}{qT_k}\langle p_k - qp_{k+1}, s\rangle. \end{aligned}$$

Since $\beta_k/(qT_k) > 0$ and $p_k - qp_{k+1} = w - s - g + (1-q)p_{k+1}$,

$$0 \le \langle w, s\rangle + (1-q)\langle p_{k+1}, s\rangle - \|s\|^2 - \langle g, s\rangle. \tag{12}$$

Monotonicity between z_{k+1} and $z_\star$ gives

$$\begin{aligned} 0 &\le T_k\langle F(z_{k+1}) + h_{k+1}, z_{k+1} - z_\star\rangle \\ &= \langle p_{k+1}, d - u_k - p_{k+1} + g + (1-\beta_k)s\rangle. \end{aligned} \tag{13}$$

Finally, Lipschitz continuity of F gives

$$0 \le \alpha_k^2 L^2\|s\|^2 - \|g\|^2. \tag{14}$$

For $t \ge 4$, the coefficients $1 - 8/(5t)$, $1/(5t)$, and $29q^2/20$ are nonnegative. Add these multiples of equation 12, equation 13, and equation 14, respectively, to the right-hand side of equation 11. Then equation 10 gives

$$\begin{aligned} V_{k+1} - \rho V_k \le& \mathbb{E}\left[\frac{1}{2}\left\|\left(1-\frac{3}{t}\right)s + (q-1)u_k + g\right\|^2 - \frac{\rho}{2}\|w\|^2 + \frac{q^2-\rho}{2}\|u_k\|^2\right. \\ &+ \left(1-\frac{8}{5t}\right)\left(\langle w, s\rangle + (1-q)\langle p_{k+1}, s\rangle - \|s\|^2 - \langle g, s\rangle\right) \\ &+ \frac{1}{5t}\langle p_{k+1}, d - u_k - p_{k+1} + g + (1-3/t)s\rangle \\ &\left.+ \frac{29}{20}q^2\alpha_k^2\overline{L}^2\|s\|^2 - \frac{29}{20}q^2\|g\|^2\right] \\ &+ \frac{1-\rho}{2}D^2 + \frac{29}{9}(T_{k+1} - qT_k)^2\sigma^2. \end{aligned}$$

Using equation 7 and

$$\left(1-\frac{8}{5t}\right)(1-q) + \frac{1-3/t}{5t} = \frac{7(t-2)}{10t^2},$$

this becomes

$$\begin{aligned} V_{k+1} - \rho V_k \le& \mathbb{E}\left[\frac{1}{2}\left\|\left(1-\frac{3}{t}\right)s + (q-1)u_k + g\right\|^2 + \frac{q^2-\rho}{2}\|u_k\|^2\right. \\ &- \frac{\rho}{2}\|w\|^2 + \left(1-\frac{8}{5t}\right)\langle w, s\rangle \\ &- \frac{1}{5t}\|p_{k+1}\|^2 + \frac{1}{5t}\left\langle p_{k+1}, d - u_k + g + \frac{7(t-2)}{2t}s\right\rangle \\ &\left.+ \left(\frac{29q^2}{20t} - 1 + \frac{8}{5t}\right)\|s\|^2 - \left(1-\frac{8}{5t}\right)\langle g, s\rangle - \frac{29}{20}q^2\|g\|^2\right] \\ &+ \frac{1-\rho}{2}D^2 + \frac{29}{9}(T_{k+1} - qT_k)^2\sigma^2. \end{aligned}$$

The two terms containing w satisfy

$$-\frac{\rho}{2}\|w\|^2+\left(1-\frac{8}{5t}\right)\langle w,s\rangle=-\frac{\rho}{2}\left\|w-\frac{1-8/(5t)}{\rho}s\right\|^2+\frac{(1-8/(5t))^2}{2\rho}\|s\|^2,$$

and the two terms containing p_{k+1} satisfy

$$\begin{aligned}&-\frac{1}{5t}\|p_{k+1}\|^2+\frac{1}{5t}\left\langle p_{k+1},d-u_k+g+\frac{7(t-2)}{2t}s\right\rangle\\&=-\frac{1}{5t}\left\|p_{k+1}-\frac{1}{2}(d-u_k+g)-\frac{7(t-2)}{4t}s\right\|^2+\frac{1}{20t}\left\|d-u_k+g+\frac{7(t-2)}{2t}s\right\|^2.\end{aligned}$$

After discarding these two nonpositive squared norms, define

$$\begin{aligned}\widehat{Q}_t(u,s,g,d):=&\frac{1}{2}\left\|\left(1-\frac{3}{t}\right)s+(q-1)u+g\right\|^2\\&+\left[\frac{(1-8/(5t))^2}{2\rho}-1+\frac{8}{5t}+\frac{29q^2}{20t}\right]\|s\|^2-\left(1-\frac{8}{5t}\right)\langle g,s\rangle\\&+\frac{1}{20t}\left\|d-u+g+\frac{7(t-2)}{2t}s\right\|^2-\frac{29}{20}q^2\|g\|^2+\frac{q^2-\rho}{2}\|u\|^2.\end{aligned}$$

Then

$$V_{k+1}-\rho V_k\le\mathbb{E}\widehat{Q}_t(u_k,s,g,d)+\frac{1-\rho}{2}D^2+\frac{29}{9}(T_{k+1}-qT_k)^2\sigma^2. \tag{15}$$

Substituting $q=(2t-1)/(2t)$ and $\rho=1-2/(5t)$ gives

$$\begin{aligned}&\frac{2}{5}(1-\rho)\|d\|^2-\frac{3}{10t}\langle u,d\rangle-\widehat{Q}_t(u,s,g,d)\\&=\frac{t-1}{4t^2}\|u\|^2+\frac{295t^3-646t^2-1029t+450}{80t^3(5t-2)}\|s\|^2+\frac{76t^2-120t+29}{80t^2}\|g\|^2+\frac{11}{100t}\|d\|^2\\&\quad+\frac{17t-44}{20t^2}\langle u,s\rangle+\frac{3}{5t}\langle u,g\rangle-\frac{1}{5t}\langle u,d\rangle+\frac{7(3t+2)}{20t^2}\langle s,g\rangle\\&\quad-\frac{7(t-2)}{20t^2}\langle s,d\rangle-\frac{1}{10t}\langle g,d\rangle.\end{aligned}$$

Completing the terms involving d gives

$$\begin{aligned}&\frac{11}{100t}\left\|d-\frac{10}{11}u-\frac{5}{11}g-\frac{35(t-2)}{22t}s\right\|^2\\&\quad+\frac{7t-11}{44t^2}\|u\|^2+\frac{2020t^3-1716t^2-18179t+6910}{880t^3(5t-2)}\|s\|^2+\frac{836t^2-1340t+319}{880t^2}\|g\|^2\\&\quad+\frac{117t-344}{220t^2}\langle u,s\rangle+\frac{28}{55t}\langle u,g\rangle+\frac{7(7t+8)}{55t^2}\langle s,g\rangle.\end{aligned}$$

Completing the terms involving u gives

$$\begin{aligned}&\frac{11}{100t}\left\|d-\frac{10}{11}u-\frac{5}{11}g-\frac{35(t-2)}{22t}s\right\|^2\\&\quad+\frac{7t-11}{44t^2}\left\|u+\frac{(117t-344)s+112tg}{70t-110}\right\|^2\\&\quad+\frac{205t^4+23518t^3-117687t^2+134397t-34550}{400t^3(5t-2)(7t-11)}\|s\|^2\\&\quad+\frac{7(t^2+53t-40)}{25t^2(7t-11)}\langle s,g\rangle+\frac{2660t^3-9584t^2+7715t-1595}{400t^2(7t-11)}\|g\|^2.\end{aligned}$$

For $c > 0$,

$$a\|s\|^2 + b\langle s, g\rangle + c\|g\|^2 = c\left\|g + \frac{b}{2c}s\right\|^2 + \left(a - \frac{b^2}{4c}\right)\|s\|^2.$$

Set $m := t - 4 \geq 0$. For the last three terms above, the coefficient of $\|g\|^2$ is

$$\frac{2660m^3 + 22336m^2 + 58723m + 46161}{400(m+4)^2(7m+17)} > 0.$$

After completing the square in g, the remaining coefficient of $\|s\|^2$ is

$$\frac{\begin{pmatrix}15580m^6 + 2129188m^5 + 26216605m^4 + 125305406m^3 \\ + 267420385m^2 + 214441225m + 2604726\end{pmatrix}}{80(m+4)^3(5m+18)(2660m^3 + 22336m^2 + 58723m + 46161)} > 0.$$

Therefore

$$\widehat{Q}_t(u, s, g, d) \leq \frac{2}{5}(1-\rho)\|d\|^2 - \frac{3}{10t}\langle u, d\rangle. \tag{16}$$

Taking expectation in equation 16 and using equation 9,

$$\mathbb{E}\widehat{Q}_t(u_k, s, g, d) \leq \frac{2}{5}(1-\rho)D^2. \tag{17}$$

Finally, $qT_k = (1-\gamma_{k+1})T_{k+1}$, so $T_{k+1} - qT_k = \gamma_{k+1}T_{k+1}$. By equation 7 at $k+1$ and $\beta_{k+1} = 3/(t+1)$,

$$\overline{L}^2 T_{k+1}^2 = \overline{L}^2 \alpha_{k+1}^2 \frac{(t+1)^2}{9} \leq \frac{t+1}{9}.$$

Hence

$$\frac{29}{9}\overline{L}^2(T_{k+1} - qT_k)^2 \leq \frac{29(4t+1)^2}{1296t^2(t+1)},$$
$$\frac{9}{25t} - \frac{29(4t+1)^2}{1296t^2(t+1)} = \frac{64t^2 + 5864t - 725}{32400t^2(t+1)} > 0.$$

Since $9/(25t) = (9/10)(1-\rho)$,

$$\frac{29}{9}(T_{k+1} - qT_k)^2\sigma^2 \leq \frac{9}{10}(1-\rho)\frac{\sigma^2}{\overline{L}^2}. \tag{18}$$

Combining equation 15, equation 17, and equation 18 gives

$$\begin{aligned} V_{k+1} - \rho V_k &\leq \left(\frac{1}{2} + \frac{2}{5}\right)(1-\rho)D^2 + \frac{9}{10}(1-\rho)\frac{\sigma^2}{\overline{L}^2} \\ &= \frac{9}{10}(1-\rho)\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right). \end{aligned}$$

Since $1 - \rho = 2/[5(k+3)]$, this is equation 6. □

B.3 Initialization and conclusion

At $k = 0$,

$$\beta_0 = 1, \qquad T_0 = \alpha_0 = \frac{7}{12\overline{L}}, \qquad \alpha_1 = \frac{1}{2\overline{L}},$$
$$\beta_1 = \frac{3}{4}, \qquad T_1 = \frac{2}{3\overline{L}}, \qquad \gamma_1 = \frac{13}{48}.$$

Set

$$s_0 := z_0 - z_1, \qquad g_0 := T_0\big(F(z_1) - F(z_0)\big), \qquad \ell := \frac{L^2}{\overline{L}^2} \in [0,1].$$

Since $\beta_0 = 1$, the resolvent relation gives

$$\begin{aligned} s_0 &= T_0(v_0 + h_1) = u_0 + T_0\big(F(z_0) + h_1\big) = u_0 + p_1 - g_0, \\ p_1 &= s_0 - u_0 + g_0. \end{aligned}$$

Subtracting $F(z_1)$ from the definition of v_1 and multiplying by T_1 gives

$$\begin{aligned} u_1 =& \frac{5}{6}u_0 + \frac{35}{72\overline{L}}\big(\delta_{\xi_1}(z_1) - \delta_{\xi_1}(z_0)\big) + \frac{13}{72\overline{L}}\delta_{\xi_1}(z_1) \\ =& \frac{5}{6}u_0 + \omega_1, \end{aligned}$$

where

$$\omega_1 := \frac{35}{72\overline{L}}\big(\delta_{\xi_1}(z_1) - \delta_{\xi_1}(z_0)\big) + \frac{13}{72\overline{L}}\delta_{\xi_1}(z_1), \qquad \mathbb{E}_0\omega_1 = 0.$$

Young's inequality with $\varepsilon = 3/10$ gives

$$\mathbb{E}_0\|\omega_1\|^2 \le \frac{3185}{10368}(1-\ell)\|s_0\|^2 + \frac{2197}{15552}\frac{\sigma^2}{\overline{L}^2}. \tag{19}$$

Since $z_1 - z_0 + p_1 = -u_0 + g_0$,

$$\begin{aligned} z_1 - z_0 + u_1 + p_1 &= g_0 - \frac{1}{6}u_0 + \omega_1, \\ u_1 - d &= \frac{5}{6}u_0 - d + \omega_1. \end{aligned}$$

Using $\mathbb{E}_0\omega_1 = 0$ and equation 19 in the definition of V_1,

$$V_1 \le \mathbb{E}\left[\frac{1}{2}\left\|g_0 - \frac{1}{6}u_0\right\|^2 + \frac{1}{2}\left\|\frac{5}{6}u_0 - d\right\|^2 + \frac{3185}{10368}(1-\ell)\|s_0\|^2\right] + \frac{2197}{15552}\frac{\sigma^2}{\overline{L}^2}. \tag{20}$$

Monotonicity, Lipschitz continuity, and unbiasedness give

$$\langle s_0 - u_0 + g_0, d - s_0\rangle \ge 0, \tag{21}$$

$$\frac{49\ell}{144}\|s_0\|^2 - \|g_0\|^2 \ge 0, \tag{22}$$

$$\mathbb{E}\langle u_0, d\rangle = 0. \tag{23}$$

For any $g, u, d, s \in \mathcal{H}$,

$$\begin{aligned} &\frac{9}{10}\|d\|^2 + \frac{21}{10}\|u\|^2 - \frac{1}{2}\left\|g - \frac{1}{6}u\right\|^2 - \frac{1}{2}\left\|\frac{5}{6}u - d\right\|^2 - \frac{3185}{10368}(1-\ell)\|s\|^2 \\ &\quad - \frac{7}{10}\langle s - u + g, d - s\rangle - \left(\frac{49\ell}{144}\|s\|^2 - \|g\|^2\right) - \frac{4}{5}\langle u, d\rangle \\ &= \frac{1}{2}\left\|g + \frac{1}{6}u - \frac{7}{10}d + \frac{7}{10}s\right\|^2 + \frac{69}{40}\left\|u + \frac{17}{69}d - \frac{49}{207}s\right\|^2 \\ &\quad + \frac{347}{6900}\left\|d - \frac{91}{1041}s\right\|^2 + \frac{4417}{249840}\|s\|^2 + (1-\ell)\frac{343}{10368}\|s\|^2. \end{aligned} \tag{24}$$

The right-hand side is nonnegative. Applying equation 24 to (g_0, u_0, d, s_0), taking expectation, and using equation 21–equation 23 gives

$$\begin{aligned} &\mathbb{E}\left[\frac{1}{2}\left\|g_0 - \frac{1}{6}u_0\right\|^2 + \frac{1}{2}\left\|\frac{5}{6}u_0 - d\right\|^2 + \frac{3185}{10368}(1-\ell)\|s_0\|^2\right] \\ &\qquad \le \frac{9}{10}\|d\|^2 + \frac{21}{10}\mathbb{E}\|u_0\|^2. \end{aligned} \tag{25}$$

Moreover,

$$\mathbb{E}\|u_0\|^2 = T_0^2\mathbb{E}\|\delta_{\xi_0}(z_0)\|^2 \le \frac{49}{144}\frac{\sigma^2}{\overline{L}^2}.$$

Combining this bound with equation 20 and equation 25,

$$\begin{aligned} V_1 &\le \frac{9}{10}\|d\|^2 + \frac{66551}{77760}\frac{\sigma^2}{\overline{L}^2} \\ &\le \frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right). \end{aligned} \tag{26}$$

By equation 6, equation 26 implies equation 5 by induction.

It remains to lower-bound T_k. The schedule gives

$$\begin{aligned} \frac{T_{k+1}}{T_k} &= \frac{2(k+4)}{2k+7}, \\ \frac{T_{k+1}^2/(k+4)}{T_k^2/(k+3)} &= 1 - \frac{1}{(2k+7)^2}. \end{aligned}$$

Since $\overline{L}^2 T_0^2/3 = 49/432$,

$$\begin{aligned} \frac{\overline{L}^2 T_k^2}{k+3} &= \frac{49}{432}\prod_{j=0}^{k-1}\left(1 - \frac{1}{(2j+7)^2}\right) \\ &\ge \frac{49}{432}\prod_{j=0}^{\infty}\left(1 - \frac{1}{(2j+7)^2}\right) \\ &= \frac{49}{432}\frac{75\pi}{256} = \frac{1225\pi}{36864} > \frac{18}{175}, \end{aligned}$$

where the equality uses the Wallis product and the last inequality follows from $\pi > 31/10$. Thus

$$T_k^2 \ge \frac{1225\pi}{36864\overline{L}^2}(k+3) > \frac{18}{175\overline{L}^2}(k+3), \qquad k \ge 0. \tag{27}$$

Finally, equation 4, equation 5, and equation 27 give, for every $N \ge 1$,

$$\begin{aligned} \mathbb{E}[\mathcal{R}(z_N)^2] &\le \frac{4V_N}{T_{N-1}^2} \\ &\le 4 \cdot \frac{9}{10}\left(D^2 + \frac{\sigma^2}{\overline{L}^2}\right)\frac{175\overline{L}^2}{18(N+2)} \\ &= 35\frac{\overline{L}^2 D^2 + \sigma^2}{N+2}. \end{aligned}$$

This proves Theorem 4.1.

B.4 Tightness of the deterministic rate

The $O(1/N)$ deterministic term in Theorem 4.1 is tight for the VRAF schedule. The following argument applies more generally to the equation VRAF whenever $L\alpha_k = \Theta(k^{-1/2})$ and $\beta_k = \Theta(k^{-1})$.

Lemma B.2. *Consider the deterministic iteration*

$$z_{k+1} = z_k - \alpha_k F(z_k) + \beta_k(z_0 - z_k), \tag{28}$$

where $\alpha_k, \beta_k > 0$, $L\alpha_k = \Theta(k^{-1/2})$, and $\beta_k = \Theta(k^{-1})$. For the one-dimensional monotone L-Lipschitz operator $F(z) = Lz$ and $z_0 = D > 0$,

$$\|F(z_k)\|^2 = \Omega\left(\frac{L^2 D^2}{k}\right).$$

Proof. For $F(z) = Lz$,

$$z_{k+1} = (1 - L\alpha_k - \beta_k)z_k + D\beta_k. \tag{29}$$

For all sufficiently large k, $0 \le 1 - L\alpha_k - \beta_k \le 1$, and there are constants $a, b > 0$ such that

$$L\alpha_k \le \frac{a}{\sqrt{k+1}}, \qquad \beta_k \ge \frac{b}{k+1}.$$

Moreover, $z_k > 0$ from some index onward. Whenever $z_k \le 0$, equation 29 gives

$$z_{k+1} - z_k = -(L\alpha_k + \beta_k)z_k + D\beta_k \ge D\beta_k.$$

Since $\sum_k \beta_k = \infty$, the iterates cannot remain nonpositive; once $z_k > 0$, the nonnegative coefficient in equation 29 keeps all subsequent iterates positive.

Choose an index K from which these properties hold and choose $c > 0$ sufficiently small that

$$\frac{cD}{\sqrt{K+1}} \le z_K, \qquad c \le \frac{1}{2}, \qquad ac < \frac{b}{2}.$$

Set $w_k = cD/\sqrt{k+1}$. For every $k \ge K$,

$$\begin{aligned}(1 - L\alpha_k - \beta_k)w_k + D\beta_k - w_{k+1} &= (w_k - w_{k+1}) + \beta_k(D - w_k) - L\alpha_k w_k \\ &\ge \frac{bD}{2(k+1)} - \frac{acD}{k+1} > 0.\end{aligned}$$

Hence equation 29 and induction give $z_k \ge w_k = cD/\sqrt{k+1}$ for every $k \ge K$. Therefore

$$\|F(z_k)\|^2 = L^2 z_k^2 \ge \frac{c^2 L^2 D^2}{k+1},$$

which proves the claim. □

In the deterministic case, VRAF has $v_k = F(z_k)$ and therefore reduces to equation 28 when $A = 0$. Moreover, Remark 4.2 gives $L\alpha_k = \Theta(k^{-1/2})$, and $\beta_k = 3/(k+3) = \Theta(k^{-1})$. Lemma B.2 therefore shows that the deterministic $O(L^2D^2/N)$ term in Theorem 4.1 is order-tight for the VRAF schedule.

C Proofs for Impossibility of Simultaneous Optimality

We prove the lower bound for adaptive randomized algorithms. At each oracle query, the algorithm may choose the query point and may query either a fresh stochastic operator or any previously sampled stochastic operator, based on its internal randomness and all previous observations.

C.1 Hard family

We use a rescaled version of the one-dimensional noisy binary search family in Foster et al. (2019, Theorem 11). Fix $M \ge 16$ and a hidden index $j \in \{0, \dots, M-1\}$. Let $Z = (Z_0, \dots, Z_M) \in \{-1, 1\}^{M+1}$ have independent coordinates with

$$\mathbb{P}_j(Z_i = 1) = \begin{cases} \frac{1}{2} - \frac{1}{8M}, & i \le j, \\ \frac{1}{2} + \frac{1}{8M}, & i > j. \end{cases}$$

Define $F_Z : \mathbb{R} \to \mathbb{R}$ by

$$F_Z(x) = \begin{cases} \frac{Z_0}{M}, & x < 0, \\ \frac{(i+1-Mx)Z_i + (Mx-i)Z_{i+1}}{M}, & x \in [i/M, (i+1)/M], \quad i = 0, \dots, M-1, \\ \frac{Z_M}{M}, & x > 1. \end{cases}$$

Let $\mathbb{E}_j$ denote expectation under $\mathbb{P}_j$ and define $F_j := \mathbb{E}_j[F_Z]$. For the monotone inclusion, set $A = 0$ and $z_0 = 0$.

Lemma C.1. *For every $M \geq 16$ and $j \in \{0, \ldots, M-1\}$, F_j is monotone and $1/(2M)$-Lipschitz, with unique zero $x_j^\star = (j + \frac{1}{2})/M$. The oracle F_Z satisfies Assumptions 2.2 and 2.4. Hence Assumptions 2.1–2.4 hold with*

$$D = 1, \qquad L = \frac{1}{2M}, \qquad \sigma^2 = \frac{1}{M^2}, \qquad L_\Delta = 3. \tag{30}$$

Proof. The coordinate means are

$$\mathbb{E}_j[Z_i] = \begin{cases} -\frac{1}{4M}, & i \leq j, \\ \frac{1}{4M}, & i > j. \end{cases}$$

Therefore

$$F_j(x) = \begin{cases} -\frac{1}{4M^2}, & x \leq \frac{j}{M}, \\ \frac{x - (j+\frac{1}{2})/M}{2M}, & \frac{j}{M} \leq x \leq \frac{j+1}{M}, \\ \frac{1}{4M^2}, & x \geq \frac{j+1}{M}. \end{cases}$$

Thus F_j is monotone and $1/(2M)$-Lipschitz and vanishes only at $x_j^\star$. Since $x_j^\star \in (0,1)$ and $z_0 = 0$, $|z_0 - x_j^\star| \leq 1$.

Unbiasedness follows from $F_j = \mathbb{E}_j[F_Z]$. Moreover, $|F_Z(x)| \leq 1/M$ for every x, so

$$\mathbb{E}_j |F_Z(x) - F_j(x)|^2 = \mathrm{Var}_{\mathbb{P}_j}(F_Z(x)) \leq \frac{1}{M^2}.$$

Finally, adjacent grid values of F_Z differ by at most $2/M$ over an interval of length $1/M$, so every F_Z is 2-Lipschitz. Since F_j is $1/(2M)$-Lipschitz,

$$\begin{aligned} |[F_Z(x) - F_j(x)] - [F_Z(y) - F_j(y)]| &\leq \left(2 + \frac{1}{2M}\right) |x - y| \\ &\leq 3|x - y|. \end{aligned}$$

This proves equation 30. □

C.2 Reduction to noisy binary search

Lemma C.2 (Adaptive noisy binary search; cf. Foster et al. (2019, proof of Theorem 11))**.** *Let J be uniform on $\{0, \ldots, M-1\}$ and, conditional on $J = j$, let the rows $Z^{(s)} = (Z_0^{(s)}, \ldots, Z_M^{(s)})$, $s \geq 1$, be independent with independent coordinates satisfying*

$$\mathbb{P}(Z_i^{(s)} = 1 \mid J = j) = \begin{cases} \frac{1}{2} - \frac{1}{8M}, & i \leq j, \\ \frac{1}{2} + \frac{1}{8M}, & i > j. \end{cases}$$

An adaptive randomized algorithm may query arbitrary entries (s, i), including multiple entries from the same row and repeated entries. If it identifies J with probability at least $3/4$, then it requires $\Omega(M^2 \log M)$ matrix-entry queries.

Proof. Let $a := 1/(8M)$ and let $\mathcal{T}_q$ be the transcript after q queries, including the algorithm's internal randomness. The query location at step t is determined by $\mathcal{T}_{t-1}$. A repeated query to an already observed entry gives no additional information. For a new entry, conditional on $\mathcal{T}_{t-1}$ its distribution is either $\mathrm{Ber}(\frac{1}{2} - a)$ or $\mathrm{Ber}(\frac{1}{2} + a)$, according to the queried coordinate and J. Hence

$$\begin{aligned} I(J; Y_t \mid \mathcal{T}_{t-1}) &\leq \max\left\{ D_{\mathrm{KL}}\left(\mathrm{Ber}(\tfrac{1}{2} + a) \,\middle\|\, \mathrm{Ber}(\tfrac{1}{2} - a)\right), D_{\mathrm{KL}}\left(\mathrm{Ber}(\tfrac{1}{2} - a) \,\middle\|\, \mathrm{Ber}(\tfrac{1}{2} + a)\right) \right\} \\ &= 2a \log \frac{1 + 2a}{1 - 2a} \leq 16a^2 = \frac{1}{4M^2}. \end{aligned}$$

By the chain rule, $I(J; \mathcal{T}_q) \leq q/(4M^2)$. If an estimator $\widehat{J}$ based on $\mathcal{T}_q$ has error probability at most $1/4$, Fano's inequality gives

$$I(J; \mathcal{T}_q) \geq \frac{3}{4} \log M - \log 2 \geq \frac{1}{2} \log M,$$

where the last inequality uses $M \geq 16$. Therefore $q \geq 2M^2 \log M$, proving the claimed order. □

Each query to a sampled stochastic operator can be simulated using at most two matrix-entry queries. If $x \in [i/M, (i+1)/M]$, then

$$F_{Z^{(s)}}(x) = \frac{(i+1-Mx)Z_i^{(s)} + (Mx-i)Z_{i+1}^{(s)}}{M},$$

so only $Z_i^{(s)}$ and $Z_{i+1}^{(s)}$ are needed; outside $[0,1]$, only one boundary entry is needed. Repeatedly querying the same sampled stochastic operator corresponds to querying further entries of the same matrix row. Hence Q stochastic-oracle queries reveal at most $2Q$ matrix entries, and Lemma C.2 implies that identification with probability at least $3/4$ requires

$$Q = \Omega(M^2 \log M). \tag{31}$$

C.3 Residual accuracy implies identification

Define

$$\tau_M := \frac{1}{8M^2}, \qquad I_j := \left[\frac{j+\frac{1}{4}}{M}, \frac{j+\frac{3}{4}}{M}\right].$$

Since $A = 0$, $\mathcal{R}(x) = |F_j(x)|$.

Lemma C.3. *For every $M \geq 16$, any adaptive randomized algorithm satisfying*

$$\sup_{0 \leq j < M} \mathbb{E}_j |F_j(\widehat{X})|^2 \leq \frac{1}{256M^4} \tag{32}$$

for the family in Subsection C.1 requires $\Omega(M^2 \log M)$ oracle queries.

Proof. The expression for F_j in the proof of Lemma C.1 gives

$$|F_j(x)| \leq \tau_M \quad \Longleftrightarrow \quad x \in I_j. \tag{33}$$

The intervals $I_0, \ldots, I_{M-1}$ are pairwise disjoint. By Markov's inequality and equation 32,

$$\begin{aligned}
\mathbb{P}_j(\widehat{X} \notin I_j) &= \mathbb{P}_j\big(|F_j(\widehat{X})| > \tau_M\big) \\
&\leq \frac{\mathbb{E}_j |F_j(\widehat{X})|^2}{\tau_M^2} \\
&\leq \frac{1}{4}.
\end{aligned}$$

Define $\widehat{J} = \ell$ when $\widehat{X} \in I_\ell$, and assign any value when $\widehat{X}$ lies outside all I_ℓ. Then

$$\mathbb{P}_j(\widehat{J} = j) \geq \mathbb{P}_j(\widehat{X} \in I_j) \geq \frac{3}{4}.$$

Thus $\widehat{J}$ identifies j with probability at least $3/4$, and equation 31 gives the result. □

C.4 Proof of Theorem 5.1

Proof. Fix $M \geq 16$ and define on $\mathbb{R}^2$ $\widetilde{F}_Z(x,u) := (F_Z(x), 3u)$ and $\widetilde{F}_j(x,u) := (F_j(x), 3u)$, with $\widetilde{A} = 0$ and $\widetilde{z}_0 = (0,0)$. By Lemma C.1,

$$\begin{aligned}
\langle \widetilde{F}_j(x,u) - \widetilde{F}_j(y,v), (x-y, u-v)\rangle &= (F_j(x) - F_j(y))(x-y) + 3|u-v|^2 \geq 0, \\
\|\widetilde{F}_j(x,u) - \widetilde{F}_j(y,v)\|^2 &\leq \frac{1}{4M^2}|x-y|^2 + 9|u-v|^2 \leq 9\|(x,u)-(y,v)\|^2.
\end{aligned}$$

Thus $\widetilde{F}_j$ is monotone, and taking $x = y$ and $u \neq v$ shows that its Lipschitz modulus is exactly $L = 3$. Its unique zero is $(x_j^\star, 0)$, whose distance from $\widetilde{z}_0$ is at most 1.

The lifted noise is $\widetilde{\delta}_Z(x,u) = (F_Z(x) - F_j(x), 0)$, so $\mathbb{E}_j\|\widetilde{\delta}_Z(x,u)\|^2 \leq 1/M^2$ and $\mathbb{E}_j\|\widetilde{\delta}_Z(x,u) - \widetilde{\delta}_Z(y,v)\|^2 \leq 9\|(x,u)-(y,v)\|^2$. Hence the lifted family satisfies the assumptions with $D = 1$, $L = 3$, $\sigma^2 = 1/M^2$, and $L_\Delta = 3$, and $\mathcal{R}_j(x,u)^2 = |F_j(x)|^2 + 9u^2$.

Each query to $\widetilde{F}_Z(x,u)$ is simulated by one query to $F_Z(x)$ followed by the deterministic coordinate $3u$, including repeated queries to a previously sampled operator. Therefore, at $\epsilon_M := 1/(256M^4)$, lifted residual accuracy implies the scalar criterion in Lemma C.3, and hence requires $\Omega(M^2 \log M)$ oracle queries. But $LD/\sqrt{\epsilon_M} + \sigma^2/\epsilon_M = 48M^2 + 256M^2 = 304M^2$, so a uniform $O(LD/\sqrt{\epsilon} + \sigma^2/\epsilon)$ bound gives only $O(M^2)$ queries, a contradiction for sufficiently large M. □

C.5 Proof of Corollary 5.2

Proof. Fix $0 \le p < 1$ and suppose that a fixed-query method satisfies, uniformly over the problem class, for some constant $C > 0$ and every integer $N \ge 2$,

$$\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2] \le C\left(\frac{L^2D^2}{N^2} + \frac{\sigma^2 \log^p N}{N}\right). \tag{34}$$

By the definition of a fixed-query method, there is a constant r such that each iteration uses at most r oracle queries. Apply equation 34 to the lifted hard family in the proof of Theorem 5.1, for which $L = 3$, $D = 1$, $\sigma^2 = M^{-2}$, and $\epsilon_M = (256M^4)^{-1}$. Choose

$$N_M = \left\lceil KM^2(\log M)^p \right\rceil,$$

where K is a sufficiently large constant depending only on C and p. Since $\log N_M = O(\log M)$,

$$C\frac{L^2D^2}{N_M^2} = O\left(\frac{1}{M^4(\log M)^{2p}}\right),$$
$$C\frac{\sigma^2 \log^p N_M}{N_M} = O\left(\frac{1}{KM^4}\right).$$

Thus, after increasing K if necessary, equation 34 gives $\mathbb{E}[\mathcal{R}(\widehat{z}_{N_M})^2] \le \epsilon_M$ for all sufficiently large M. The method then uses at most

$$rN_M = O\big(M^2(\log M)^p\big) = o(M^2 \log M)$$

oracle queries because $p < 1$. This contradicts the $\Omega(M^2 \log M)$ lower bound in equation 31 and Lemma C.3. □

D Complete RRSEG proof

D.1 Preliminary bounds

For each k, define $H_k(z) := F(z) + a_k(z - c_k)$ and let r_k be the unique zero of $A + H_k$. The operator H_k is a_k-strongly monotone and $(L + a_k)$-Lipschitz. We first derive bounds that do not depend on the particular horizon-dependent schedule.

D.1.1 One-step stochastic forward-backward-forward bound

Lemma D.1 (One-step stochastic forward-backward-forward contraction). *Let $B : \mathcal{H} \to \mathcal{H}$ be μ-strongly monotone and M-Lipschitz, let A be maximally monotone, and let r satisfy $0 \in B(r) + A(r)$. Assume $2\mu \le M$ and consider*

$$y = J_{\eta A}\big(x - \eta(B(x) + \varepsilon)\big),$$
$$x^+ = y - \eta\big(B(y) + \zeta - B(x) - \varepsilon\big),$$

where $\mathbb{E}\varepsilon = 0$, $\mathbb{E}\|\varepsilon\|^2 \le \sigma^2$, and, conditional on ε,

$$\mathbb{E}[\zeta \mid \varepsilon] = 0, \qquad \mathbb{E}[\|\zeta\|^2 \mid \varepsilon] \le \sigma^2.$$

If $\eta M \le 1/3$, then

$$\mathbb{E}\|x^+ - r\|^2 \le \left(1 - \frac{4}{3}\eta\mu\right)\|x - r\|^2 + \frac{5}{2}\eta^2\sigma^2. \tag{35}$$

Proof. The resolvent step gives

$$h_y := \frac{x-y}{\eta} - B(x) - \varepsilon \in A(y).$$

Choose $h_r = -B(r) \in A(r)$. Then

$$x^+ = x - \eta\big(h_y + B(y) + \zeta\big).$$

Conditioning on ε and using the conditional mean of ζ,

$$\begin{aligned}\mathbb{E}[\|x^+ - r\|^2 \mid \varepsilon] =& \|x-r\|^2 - 2\eta\langle h_y + B(y), x - r\rangle \\ &+ \eta^2\|h_y + B(y)\|^2 + \eta^2\mathbb{E}[\|\zeta\|^2 \mid \varepsilon] \\ \le& \|x-r\|^2 - 2\eta\mu\|y-r\|^2 - 2\eta\langle h_y + B(y), x - y\rangle \\ &+ \eta^2\|h_y + B(y)\|^2 + \eta^2\sigma^2,\end{aligned}$$

where the inequality uses monotonicity of A and strong monotonicity of B. Moreover,

$$\begin{aligned}&-2\eta\langle h_y + B(y), x - y\rangle + \eta^2\|h_y + B(y)\|^2 \\ &\quad = -\|x-y\|^2 + \eta^2\|\varepsilon + B(x) - B(y)\|^2.\end{aligned}$$

Let $t := \eta M$. Young's inequality and Lipschitz continuity of B give

$$\begin{aligned}\eta^2\|\varepsilon + B(x) - B(y)\|^2 &\le \frac{\eta^2}{1-t}\|\varepsilon\|^2 + \frac{\eta^2}{t}\|B(x) - B(y)\|^2 \\ &\le \frac{\eta^2}{1-t}\|\varepsilon\|^2 + t\|x-y\|^2.\end{aligned}$$

Taking expectation over ε yields

$$\begin{aligned}\mathbb{E}\|x^+ - r\|^2 \le& \|x-r\|^2 - 2\eta\mu\mathbb{E}\|y-r\|^2 - (1-t)\mathbb{E}\|x-y\|^2 \\ &+ \frac{2-t}{1-t}\eta^2\sigma^2.\end{aligned} \tag{36}$$

Write $s := \eta\mu$. Since $2\mu \le M$, we have $2s \le t$. Together with $t \le 1/3$, this gives $1 - t - 4s \ge 1 - 3t \ge 0$. Hence

$$\begin{aligned}&2s\|y-r\|^2 + (1-t)\|x-y\|^2 - \frac{4}{3}s\|x-r\|^2 \\ &\quad = (1 - t + 2s)\left\|x - y - \frac{2s}{1-t+2s}(x-r)\right\|^2 \\ &\qquad + \frac{2s(1-t-4s)}{3(1-t+2s)}\|x-r\|^2 \ge 0.\end{aligned}$$

Also,

$$\frac{5}{2} - \frac{2-t}{1-t} = \frac{1-3t}{2(1-t)} \ge 0.$$

Substituting the last two inequalities into equation 36 proves equation 35. □

For the schedules considered below, $a_k \le L$, so $2a_k \le L + a_k$. Applying Lemma D.1 conditionally with

$$B(z) = H_k(z), \qquad \mu = a_k, \qquad M = L + a_k,$$

gives, whenever $\eta_k(L + a_k) \le 1/3$,

$$\mathbb{E}\|z_{k+1} - r_k\|^2 \le \left(1 - \frac{4}{3}\eta_k a_k\right)\mathbb{E}\|z_k - r_k\|^2 + \frac{5}{2}\eta_k^2\sigma^2. \tag{37}$$

D.1.2 EFFECT OF RECENTERING

For each k, choose $h_k \in A(r_k)$ such that

$$F(r_k) + h_k + a_k(r_k - c_k) = 0. \tag{38}$$

The recentering update gives

$$a_{k+1}(z - c_{k+1}) = a_k(z - c_k) + \lambda_k(z - z_{k+1}). \tag{39}$$

Subtracting equation 38 at k and $k+1$ and using equation 39,

$$\begin{aligned}\lambda_k(z_{k+1} - r_k) =& F(r_{k+1}) + h_{k+1} - F(r_k) - h_k \\ &+ a_{k+1}(r_{k+1} - r_k).\end{aligned} \tag{40}$$

Lemma D.2 (Bounds from recentering)**.** *For every RRSEG update,*

$$\|z_{k+1} - r_{k+1}\| \le \|z_{k+1} - r_k\|, \tag{41}$$

$$\|r_{k+1} - r_k\| \le \frac{\lambda_k}{a_{k+1}}\|z_{k+1} - r_k\|. \tag{42}$$

Moreover, if

$$g_k := F(r_k) + h_k + a_0(r_k - z_0),$$

then $g_0 = 0$ *and*

$$\|g_{k+1} - g_k\| \le \lambda_k\|z_{k+1} - r_k\|. \tag{43}$$

Proof. The monotonicity of $F + A$ gives

$$\langle F(r_{k+1}) + h_{k+1} - F(r_k) - h_k, r_{k+1} - r_k\rangle \ge 0.$$

If $\lambda_k = 0$, then equation 40 and strong monotonicity give $r_{k+1} = r_k$, so all conclusions are immediate. Suppose $\lambda_k > 0$. Taking the inner product of equation 40 with $r_{k+1} - r_k$ gives

$$a_{k+1}\|r_{k+1} - r_k\|^2 \le \lambda_k\langle z_{k+1} - r_k, r_{k+1} - r_k\rangle,$$

which proves equation 42. Using the same inequality,

$$\begin{aligned}&\|z_{k+1} - r_k\|^2 - \|z_{k+1} - r_{k+1}\|^2 \\ &\quad= 2\langle z_{k+1} - r_k, r_{k+1} - r_k\rangle - \|r_{k+1} - r_k\|^2 \\ &\quad\ge \left(\frac{2a_{k+1}}{\lambda_k} - 1\right)\|r_{k+1} - r_k\|^2 \ge 0,\end{aligned}$$

which proves equation 41.

The identity equation 40 can also be written as

$$\lambda_k(z_{k+1} - r_k) = (g_{k+1} - g_k) + (a_{k+1} - a_0)(r_{k+1} - r_k).$$

Moreover,

$$\langle g_{k+1} - g_k, r_{k+1} - r_k\rangle \ge a_0\|r_{k+1} - r_k\|^2 \ge 0.$$

Therefore

$$\lambda_k^2\|z_{k+1} - r_k\|^2 \ge \|g_{k+1} - g_k\|^2,$$

which proves equation 43. □

D.1.3 SCHEDULE-FREE RESIDUAL BOUNDS

Proposition D.3 (Residual bound at the regularized zero)**.** *For every* $m \geq 1$*,*

$$\begin{aligned}\left(\mathbb{E}[\mathcal{R}(r_m)^2]\right)^{1/2} &\leq \left(\mathbb{E}\|F(r_m)+h_m\|^2\right)^{1/2} \\ &\leq a_0 D + \sum_{k=0}^{m-1} \lambda_k \left(\mathbb{E}\|z_{k+1}-r_k\|^2\right)^{1/2}. \end{aligned} \tag{44}$$

Proof. Let $z_\star$ solve the original inclusion and choose $h_\star \in A(z_\star)$ such that $F(z_\star)+h_\star = 0$. At $k=0$, equation 38 gives

$$F(r_0)+h_0 = a_0(z_0-r_0).$$

Monotonicity of $F+A$ gives

$$\begin{aligned}0 &\leq \langle F(r_0)+h_0-F(z_\star)-h_\star, r_0-z_\star\rangle \\ &= a_0\langle z_0-r_0, z_0-z_\star\rangle - a_0\|z_0-r_0\|^2,\end{aligned}$$

so

$$\|z_0-r_0\| \leq D. \tag{45}$$

By equation 43 and $g_0=0$,

$$\|g_m\| \leq \sum_{k=0}^{m-1} \lambda_k \|z_{k+1}-r_k\|. \tag{46}$$

Monotonicity of $F+A$ also gives

$$\langle g_m, r_m-r_0\rangle \geq a_0\|r_m-r_0\|^2.$$

Hence

$$\|g_m - a_0(r_m-r_0)\|^2 \leq \|g_m\|^2.$$

Since

$$F(r_m)+h_m = g_m - a_0(r_m-r_0) + a_0(z_0-r_0),$$

we obtain

$$\|F(r_m)+h_m\| \leq a_0 D + \|g_m\|.$$

Combining this inequality with equation 46, taking L_2 norms, and applying Minkowski's inequality proves equation 44. □

The final composite output will be obtained from the forward-backward displacement.

Lemma D.4 (Forward-backward displacement)**.** *Let*

$$p_m := J_{\frac{1}{L}A}\left(z_m - \frac{1}{L}F(z_m)\right).$$

Then

$$\begin{aligned}L\left(\mathbb{E}\|z_m-p_m\|^2\right)^{1/2} \leq& \sqrt{2}L\left(\mathbb{E}\|z_m-r_m\|^2\right)^{1/2} + \left(\mathbb{E}\|F(r_m)+h_m\|^2\right)^{1/2} \\ \leq& \sqrt{2}L\left(\mathbb{E}\|z_m-r_m\|^2\right)^{1/2} + a_0 D + \sum_{k=0}^{m-1}\lambda_k\left(\mathbb{E}\|z_{k+1}-r_k\|^2\right)^{1/2}. \end{aligned} \tag{47}$$

Proof. Write

$$d := z_m - r_m, \qquad f := \frac{F(z_m)-F(r_m)}{L}, \qquad q := \frac{F(r_m)+h_m}{L}.$$

Since $h_m \in A(r_m)$,

$$r_m = J_{\frac{1}{L}A}\left(r_m - \frac{1}{L}F(r_m) + q\right).$$

Let $s := p_m - r_m$. Firm nonexpansiveness of the resolvent gives

$$\|s\|^2 \le \langle s, d - f - q\rangle,$$

which is equivalent to

$$\left\|s - \frac{1}{2}(d - f - q)\right\| \le \frac{1}{2}\|d - f - q\|.$$

Therefore

$$\begin{aligned}\|z_m - p_m\| &= \|d - s\| \\ &\le \frac{1}{2}\big(\|d + f + q\| + \|d - f - q\|\big) \\ &\le \sqrt{\|d\|^2 + \|f + q\|^2} \\ &\le \sqrt{\|d\|^2 + \big(\|d\| + \|q\|\big)^2} \\ &\le \sqrt{2}\|d\| + \|q\|.\end{aligned}$$

Here the third line follows from the parallelogram identity, the fourth uses the L-Lipschitz continuity of F, and the last uses $\sqrt{a^2 + (a+b)^2} \le \sqrt{2}a + b$ for $a, b \ge 0$. Taking L_2 norms gives the first inequality in equation 47; the second follows from Proposition D.3. □

D.2 Tracking bound and schedule

The schedule-free bounds contain the sum $\sum_k \lambda_k (\mathbb{E}\|z_{k+1} - r_k\|^2)^{1/2}$. For the deterministic part, consider the envelope $(a_0/a_k)^p D^2$. When $a_N = L$,

$$\begin{aligned}a_0^{p/2} \sum_{k=0}^{N-1} \frac{\lambda_k}{a_{k+1}^{p/2}} &\le a_0^{p/2} \int_{a_0}^{L} \frac{da}{a^{p/2}} \\ &= \frac{2L}{p-2}\left[\frac{a_0}{L} - \left(\frac{a_0}{L}\right)^{p/2}\right]\end{aligned}$$

for $p > 2$, so any $p > 2$ keeps this contribution of order $a_0 D$. For the stochastic part, an envelope proportional to σ^2/a_k^2 makes the same sum proportional to $\sum_k \lambda_k/a_{k+1}$ and hence logarithmic. This motivates the two envelopes used below.

Let $Q \ge 1$, $\gamma > 0$, and define

$$\begin{aligned}\eta_k &= \frac{1}{3(L + Qa_k)}, \\ a_{k+1} &= \min\left\{L, \frac{a_k}{1 - \gamma\eta_k a_k}\right\}, \qquad \lambda_k = a_{k+1} - a_k. \end{aligned} \tag{48}$$

Then

$$\eta_k(L + a_k) \le \frac{1}{3}, \qquad \eta_k a_k \le \frac{1}{3(Q+1)}. \tag{49}$$

Lemma D.5 (One-step envelope bounds)**.** *Let $p > 2$, $0 < \gamma < 2/3$, and $p\gamma \le 4/3$. Under equation 48,*

$$\left(1 - \frac{4}{3}\eta_k a_k\right)\left(\frac{a_0}{a_k}\right)^p \le \left(\frac{a_0}{a_{k+1}}\right)^p, \tag{50}$$

$$\left(1 - \frac{4}{3}\eta_k a_k\right)\frac{5}{4(2-3\gamma)(Q+1)a_k^2} + \frac{5}{2}\eta_k^2 \le \frac{5}{4(2-3\gamma)(Q+1)a_{k+1}^2}. \tag{51}$$

Proof. The update gives

$$\frac{a_k}{a_{k+1}} \geq 1 - \gamma\eta_k a_k.$$

Bernoulli's inequality and $p\gamma \leq 4/3$ therefore give

$$\left(\frac{a_k}{a_{k+1}}\right)^p \geq (1-\gamma\eta_k a_k)^p \geq 1 - p\gamma\eta_k a_k \geq 1 - \frac{4}{3}\eta_k a_k,$$

which proves equation 50.

For equation 51, let $u_k := 1 - a_k/a_{k+1}$. Then $0 \leq u_k \leq \gamma\eta_k a_k$ and

$$\begin{aligned}\frac{a_k^2}{a_{k+1}^2} - \left(1 - \frac{4}{3}\eta_k a_k\right) &= \frac{4}{3}\eta_k a_k - 2u_k + u_k^2 \\ &\geq \left(\frac{4}{3} - 2\gamma\right)\eta_k a_k.\end{aligned}$$

Hence

$$\begin{aligned}&\frac{5}{4(2-3\gamma)(Q+1)a_{k+1}^2} - \left(1 - \frac{4}{3}\eta_k a_k\right)\frac{5}{4(2-3\gamma)(Q+1)a_k^2} \\ &\quad\geq \frac{5\eta_k}{6(Q+1)a_k} \geq \frac{5}{2}\eta_k^2,\end{aligned}$$

where the last inequality follows from equation 49. □

Lemma D.6 (Tracking bound)**.** *Let $p > 2$, $0 < \gamma < 2/3$, and $p\gamma \leq 4/3$. Under equation 48,*

$$\mathbb{E}\|z_k - r_k\|^2 \leq \left(\frac{a_0}{a_k}\right)^p D^2 + \frac{5\sigma^2}{4(2-3\gamma)(Q+1)a_k^2}, \tag{52}$$

$$\mathbb{E}\|z_{k+1} - r_k\|^2 \leq \left(\frac{a_0}{a_{k+1}}\right)^p D^2 + \frac{5\sigma^2}{4(2-3\gamma)(Q+1)a_{k+1}^2}. \tag{53}$$

Proof. At $k = 0$, equation 45 gives $\|z_0 - r_0\| \leq D$. Assume equation 52 holds at iteration k. Using equation 37, equation 50, and equation 51,

$$\begin{aligned}\mathbb{E}\|z_{k+1} - r_k\|^2 &\leq \left(1 - \frac{4}{3}\eta_k a_k\right)\mathbb{E}\|z_k - r_k\|^2 + \frac{5}{2}\eta_k^2\sigma^2 \\ &\leq \left(\frac{a_0}{a_{k+1}}\right)^p D^2 + \frac{5\sigma^2}{4(2-3\gamma)(Q+1)a_{k+1}^2}.\end{aligned}$$

Thus equation 53 holds. Then equation 41 gives equation 52 at iteration $k+1$. □

D.3 Reaching $a_N = L$ and the core finite-horizon bound

Lemma D.7 (Condition for $a_N = L$)**.** *Suppose equation 48 is used for N iterations and*

$$\frac{3}{\gamma}\left(\frac{L}{a_0} - 1 + Q\log\frac{L}{a_0}\right) \leq N. \tag{54}$$

Then $a_N = L$.

Proof. Assume $a_N < L$. Then no clipping occurs and

$$\frac{1}{a_{k+1}} = \frac{1}{a_k} - \gamma\eta_k, \qquad \frac{a_{k+1}}{a_k} = \frac{1}{1-\gamma\eta_k a_k}.$$

Using $-\log(1-x) \geq x$,

$$\begin{aligned}\frac{3}{\gamma}\left[L\left(\frac{1}{a_k}-\frac{1}{a_{k+1}}\right)+Q\log\frac{a_{k+1}}{a_k}\right]&=3L\eta_k+\frac{3Q}{\gamma}\left[-\log(1-\gamma\eta_k a_k)\right]\\&\geq 3\eta_k(L+Qa_k)=1.\end{aligned}$$

Summing over $k=0,\ldots,N-1$ gives

$$\frac{3}{\gamma}\left[L\left(\frac{1}{a_0}-\frac{1}{a_N}\right)+Q\log\frac{a_N}{a_0}\right]\geq N.$$

The expression in brackets is strictly increasing in a_N. Since $a_N < L$, this contradicts equation 54. □

Proposition D.8 (Core finite-horizon bound)**.** *Let $p>2$, $0<\gamma<2/3$, $p\gamma\leq 4/3$, $Q\geq 1$, and $0<a_0<L$. If equation 54 holds, define*

$$p_N := J_{\frac{1}{L}A}\left(z_N-\frac{1}{L}F(z_N)\right).$$

Then

$$\begin{aligned}L\left(\mathbb{E}\|z_N-p_N\|^2\right)^{1/2}\leq& LD\left[\frac{p}{p-2}\frac{a_0}{L}+\left(\sqrt{2}-\frac{2}{p-2}\right)\left(\frac{a_0}{L}\right)^{p/2}\right]\\&+\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\left(\log\frac{L}{a_0}+\sqrt{2}\right).\end{aligned}\tag{55}$$

Proof. Lemma D.7 gives $a_N = L$. Lemma D.6 therefore gives

$$L\left(\mathbb{E}\|z_N-r_N\|^2\right)^{1/2}\leq LD\left(\frac{a_0}{L}\right)^{p/2}+\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}.$$

Proposition D.3 and equation 53 give

$$\begin{aligned}\left(\mathbb{E}\|F(r_N)+h_N\|^2\right)^{1/2}\leq& a_0D+a_0^{p/2}D\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}^{p/2}}\\&+\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}}.\end{aligned}$$

Since both integrands are decreasing,

$$\begin{aligned}a_0^{p/2}\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}^{p/2}}&\leq\frac{2L}{p-2}\left[\frac{a_0}{L}-\left(\frac{a_0}{L}\right)^{p/2}\right],\\\sum_{k=0}^{N-1}\frac{\lambda_k}{a_{k+1}}&\leq\log\frac{L}{a_0}.\end{aligned}$$

Thus

$$\begin{aligned}\left(\mathbb{E}\|F(r_N)+h_N\|^2\right)^{1/2}\leq& LD\left[\frac{p}{p-2}\frac{a_0}{L}-\frac{2}{p-2}\left(\frac{a_0}{L}\right)^{p/2}\right]\\&+\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\log\frac{L}{a_0}.\end{aligned}$$

Substituting the last two bounds into Lemma D.4 proves equation 55. □

D.4 Choosing the quartic tracking envelope

The admissible parameters satisfy $p > 2$ and $p\gamma \le 4/3$. Solving equation 54 for the largest admissible Q gives

$$Q = \frac{\frac{\gamma N}{3} - \left(\frac{L}{a_0} - 1\right)}{\log\left(\frac{L}{a_0}\right)}.$$

To compare the leading coefficients, write

$$a_0 = \frac{3\rho L}{\gamma N}, \qquad \rho > 1,$$

and hold ρ fixed. Then

$$Q = \frac{\gamma(\rho-1)}{3\rho}\frac{N}{\log N}\big(1 + o(1)\big).$$

The leading deterministic coefficient in Proposition D.8 is

$$\frac{3\rho}{N}\frac{p}{\gamma(p-2)} + o(N^{-1}).$$

For fixed γ, $p/(p-2)$ decreases with p, so $p\gamma \le 4/3$ selects

$$p = \frac{4}{3\gamma}.$$

With this choice, the leading deterministic coefficient and the square of the leading stochastic coefficient both depend on γ through

$$\frac{1}{\gamma(2-3\gamma)}.$$

Indeed,

$$\frac{p}{\gamma(p-2)} = \frac{2}{\gamma(2-3\gamma)},$$

while

$$\frac{1}{(2-3\gamma)Q} = \frac{3\rho}{\gamma(2-3\gamma)(\rho-1)}\frac{\log N}{N}\big(1 + o(1)\big).$$

Since $\gamma(2-3\gamma)$ is maximized at $\gamma = 1/3$, the corresponding value is $p = 4$. We therefore use the quartic deterministic envelope and allow every $0 < \gamma \le 1/3$ in the finite-horizon theorem; Corollary 6.2 uses $\gamma = 1/3$.

For $p = 4$, Proposition D.8 becomes

$$L\left(\mathbb{E}\|z_N - p_N\|^2\right)^{1/2} \le \left(2a_0 + (\sqrt{2}-1)\frac{a_0^2}{L}\right) D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\left(\log\frac{L}{a_0} + \sqrt{2}\right). \tag{56}$$

D.5 Equation case

The equation case uses the same one-step, recentering, tracking, and traversal bounds. Only the final conversion is simpler.

Corollary D.9 (Equation-case last-iterate bound)**.** *Suppose $A = 0$, $p = 4$, $0 < \gamma \le 1/3$, and the conditions of Lemma D.7 hold. Then*

$$\left(\mathbb{E}\|F(z_N)\|^2\right)^{1/2} \le 2a_0 D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\left(\log\frac{L}{a_0} + 1\right). \tag{57}$$

Proof. When $A = 0$, $h_N = 0$. Proposition D.3, equation 53, and the two integral bounds used in the proof of Proposition D.8 give

$$\left(\mathbb{E}\|F(r_N)\|^2\right)^{1/2} \le \left(2a_0 - \frac{a_0^2}{L}\right) D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\log\frac{L}{a_0}.$$

Lipschitz continuity and equation 52 give

$$\left(\mathbb{E}\|F(z_N) - F(r_N)\|^2\right)^{1/2} \le \frac{a_0^2}{L}D + \sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}.$$

Adding these two inequalities proves equation 57. □

D.6 Composite output

Let

$$m := \lfloor \min\{Q, N\} \rfloor, \qquad \overline{F}_N := \frac{1}{m}\sum_{k=N-m}^{N-1} F_{\xi_k}(z_k),$$

and define

$$\widehat{z}_N := J_{\frac{1}{L}A}\left(z_N - \frac{1}{L}\overline{F}_N\right).$$

The average uses oracle evaluations already computed by RRSEG.

Lemma D.10 (Accuracy of the averaged oracle evaluations)**.** *Assume $p = 4$, $0 < \gamma \le 1/3$, and $a_N = L$. Then*

$$\begin{aligned}\left(\mathbb{E}\|\overline{F}_N - F(z_N)\|^2\right)^{1/2} \le& \frac{77}{32}\frac{a_0^2}{L}D + \frac{\sigma}{\sqrt{m}} \\ &+ \frac{9}{4}\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}.\end{aligned} \tag{58}$$

Proof. From the update of a_k and equation 49,

$$1 - \frac{a_k}{a_{k+1}} \le \gamma\eta_k a_k \le \frac{1}{9(Q+1)}.$$

Since $m \le Q$ and $a_N = L$,

$$\begin{aligned}\frac{a_{N-m}}{L} &= \prod_{k=N-m}^{N-1}\frac{a_k}{a_{k+1}} \\ &\ge 1 - \sum_{k=N-m}^{N-1}\left(1 - \frac{a_k}{a_{k+1}}\right) \\ &\ge 1 - \frac{m}{9(Q+1)} > \frac{8}{9}.\end{aligned}$$

Thus, for $N - m \le k \le N$,

$$\left(\mathbb{E}\|z_k - r_k\|^2\right)^{1/2} \le \frac{81}{64}\frac{a_0^2}{L^2}D + \frac{9}{8}\frac{\sigma}{L}\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}. \tag{59}$$

The same bound holds for $(\mathbb{E}\|z_{k+1} - r_k\|^2)^{1/2}$ when $N - m \le k \le N - 1$. At the terminal iterate, the sharper bound

$$\left(\mathbb{E}\|z_N - r_N\|^2\right)^{1/2} \le \frac{a_0^2}{L^2}D + \frac{\sigma}{L}\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}} \tag{60}$$

follows directly from equation 52 and $a_N = L$.

By equation 42, for every $N - m \le k \le N$,

$$\begin{aligned}\left(\mathbb{E}\|r_k - r_N\|^2\right)^{1/2} &\le \sum_{j=k}^{N-1} \frac{\lambda_j}{a_{j+1}} \left(\mathbb{E}\|z_{j+1} - r_j\|^2\right)^{1/2} \\ &< \frac{1}{9}\left[\frac{81}{64}\frac{a_0^2}{L^2}D + \frac{9}{8}\frac{\sigma}{L}\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}\right].\end{aligned}$$

Combining this inequality with equation 59 and equation 60,

$$\left(\mathbb{E}\|z_k - z_N\|^2\right)^{1/2} \le \frac{77}{32}\frac{a_0^2}{L^2}D + \frac{9}{4}\frac{\sigma}{L}\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}. \tag{61}$$

Indeed, the deterministic coefficient is $(1 + 1/9)(81/64) + 1 = 77/32$, and the stochastic coefficient is $(1 + 1/9)(9/8) + 1 = 9/4$.

Now decompose

$$\begin{aligned}\overline{F}_N - F(z_N) =& \frac{1}{m}\sum_{k=N-m}^{N-1}\left(F_{\xi_k}(z_k) - F(z_k)\right) \\ &+ \frac{1}{m}\sum_{k=N-m}^{N-1}\left(F(z_k) - F(z_N)\right).\end{aligned}$$

Because z_k is determined before ξ_k is drawn, the errors in the first sum form a martingale difference sequence, so

$$\mathbb{E}\left\|\frac{1}{m}\sum_{k=N-m}^{N-1}\left(F_{\xi_k}(z_k) - F(z_k)\right)\right\|^2 \le \frac{\sigma^2}{m}.$$

For the second sum, Lipschitz continuity and equation 61 give

$$\begin{aligned}&\left(\mathbb{E}\left\|\frac{1}{m}\sum_{k=N-m}^{N-1}\left(F(z_k) - F(z_N)\right)\right\|^2\right)^{1/2} \\ &\le \frac{77}{32}\frac{a_0^2}{L}D + \frac{9}{4}\sigma\sqrt{\frac{5}{4(2-3\gamma)(Q+1)}}.\end{aligned}$$

Minkowski's inequality proves equation 58. □

Lemma D.11 (Residual of the final resolvent point)**.** *Let*

$$p_N = J_{\frac{1}{L}A}\left(z_N - \frac{1}{L}F(z_N)\right).$$

Then

$$\mathcal{R}(\widehat{z}_N) \le \sqrt{2}L\|z_N - p_N\| + \sqrt{2}\|\overline{F}_N - F(z_N)\|. \tag{62}$$

Proof. The resolvent definitions give

$$\begin{aligned}h_p &:= L(z_N - p_N) - F(z_N) \in A(p_N), \\ h_{\widehat{z}} &:= L(z_N - \widehat{z}_N) - \overline{F}_N \in A(\widehat{z}_N).\end{aligned}$$

By monotonicity and Lipschitz continuity of F,

$$\begin{aligned}\|F(p_N) + h_p\|^2 &= \|L(z_N - p_N) - (F(z_N) - F(p_N))\|^2 \\ &\le 2L^2\|z_N - p_N\|^2.\end{aligned}$$

Let $e := \widehat{z}_N - p_N$ and $\delta := \overline{F}_N - F(z_N)$. Firm nonexpansiveness applied to the two resolvent inputs gives

$$\|e\|^2 \le -\frac{1}{L}\langle e, \delta\rangle.$$

Therefore

$$\left\|e + \frac{\delta}{L}\right\|^2 \le \frac{\|\delta\|^2}{L^2} - \|e\|^2.$$

Using Lipschitz continuity,

$$\begin{aligned}
&\|(F(\widehat{z}_N) + h_{\widehat{z}}) - (F(p_N) + h_p)\| \\
&\quad \le L\|e\| + L\left\|e + \frac{\delta}{L}\right\| \\
&\quad \le \sqrt{2}\|\delta\|.
\end{aligned}$$

Combining the last two bounds proves equation 62. □

D.7 Final parameter choices

Fix $0 < \gamma \le 1/3$ and $a_0 \in (0, L)$. For fixed a_0 and γ, increasing Q decreases the stochastic term but slows the increase of a_k. We therefore take the largest Q allowed by equation 54:

$$Q_{N,\gamma}(a_0) := \frac{\frac{\gamma N}{3} - \left(\frac{L}{a_0} - 1\right)}{\log\left(\frac{L}{a_0}\right)}.$$

With $Q = Q_{N,\gamma}(a_0)$, Lemma D.7 gives $a_N = L$. Combining equation 56, Lemma D.10, and equation 62 gives

$$\begin{aligned}
\left(\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2]\right)^{1/2} \le &\sqrt{2}LD\left[2\frac{a_0}{L} + \left(\sqrt{2} - 1 + \frac{77}{32}\right)\left(\frac{a_0}{L}\right)^2\right] + \frac{\sqrt{2}\sigma}{\sqrt{m}} \\
&+ \sigma\sqrt{\frac{5}{2(2-3\gamma)(Q+1)}}\left(\log\frac{L}{a_0} + \sqrt{2} + \frac{9}{4}\right),
\end{aligned}$$

where $m = \lfloor \min\{Q, N\} \rfloor$. This proves Theorem 6.1.

For Corollary 6.2, choose

$$\gamma = \frac{1}{3}, \qquad a_0 = \frac{18L}{N}.$$

Then

$$Q_N = Q_{N,1/3}\left(\frac{18L}{N}\right) = \frac{\frac{N}{18} + 1}{\log\left(\frac{N}{18}\right)}.$$

For $N \ge 19$, $N/18 > 1$ and $0 < \log(N/18) < N/18 < N/18 + 1$, so $Q_N > 1$. Substituting these values into Theorem 6.1 gives

$$\begin{aligned}
\left(\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2]\right)^{1/2} \le &\frac{36\sqrt{2}LD}{N} + \frac{324\sqrt{2}LD}{N^2}\left(\sqrt{2} - 1 + \frac{77}{32}\right) + \frac{\sqrt{2}\sigma}{\sqrt{m_N}} \\
&+ \sigma\sqrt{\frac{5}{2(Q_N+1)}}\left(\log\frac{N}{18} + \sqrt{2} + \frac{9}{4}\right),
\end{aligned}$$

where

$$m_N = \lfloor \min\{Q_N, N\} \rfloor.$$

Since $Q_N = \Theta(N/\log N)$ and $m_N = \Theta(N/\log N)$,

$$\mathbb{E}[\mathcal{R}(\widehat{z}_N)^2] = O\left(\frac{L^2 D^2}{N^2} + \frac{\sigma^2 \log^3 N}{N}\right).$$

Finally, applying Corollary D.9 to the same concrete schedule gives the sharper equation-case last-iterate bound

$$\left(\mathbb{E}\|F(z_N)\|^2\right)^{1/2} \le \frac{36LD}{N} + \sigma\sqrt{\frac{5}{4(Q_N+1)}}\left(\log\frac{N}{18} + 1\right). \tag{63}$$

D.8 Shape of the concrete schedule

For the parameters in Corollary 6.2, a_k is nondecreasing and remains at most L. Before it reaches L,

$$a_{k+1} - a_k = \frac{\frac{1}{3}\eta_k a_k^2}{1 - \frac{1}{3}\eta_k a_k} = \frac{a_k^2}{9L + (9Q_N - 1)a_k}.$$

The right-hand side increases with a_k, so the regularization parameter grows more rapidly as it approaches L. Figure 2 illustrates this behavior for the horizon used in the experiments.

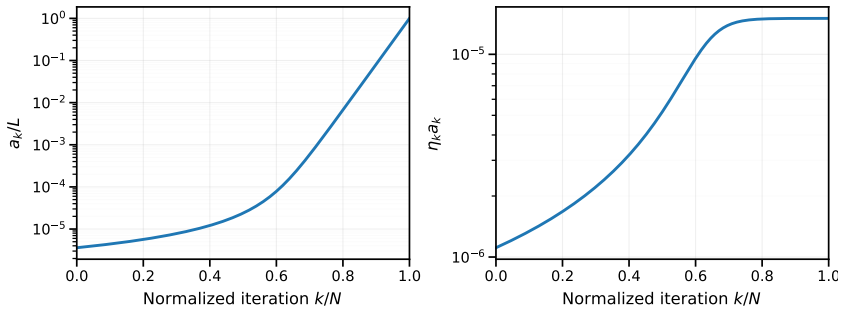

Figure 2: RRSEG parameters from Corollary 6.2 for $N = 5{,}000{,}000$, the horizon used for the solid RRSEG trajectories in the experiments. Left: a_k/L. Right: $\eta_k a_k$.

E Experimental details

E.1 Evaluation protocol and parameter choices

Problems 1 and 2 and rock-paper-scissors use $A = 0$, while Problem 3 uses a nonzero full-domain nonsmooth operator. Each $F_\xi(z)$ evaluation counts as one oracle query; every method uses 50 runs and at most 10^7 oracle queries per run.

RRSEG uses Corollary 6.2 with the single fixed horizon $N = 5{,}000{,}000$ on all four problems. Every solid RRSEG curve records the intermediate iterates z_k from that run. The stars in Figure 1 use separate RRSEG runs with horizon $N = Q/2$ at $Q \in \{10^4, 10^5, 10^6, 10^7\}$ oracle queries. Problems 1 and 2 and rock-paper-scissors use the terminal iterate z_N, while Problem 3 uses the corrected output $\widehat{z}_N$.

On Problems 1 and 2 and rock-paper-scissors, BC-SEG+ uses $\gamma = 1/(2L)$ and $\alpha_k = 1/[18(k/100 + 1)]$. On Problem 3, BC-PSEG+ uses $\gamma = 1/(2L)$ and $\alpha_k = 1/[18\sqrt{k/100 + 1}]$. Stochastic GOMA is run as in Theorem 4 of Sohrabi et al. (2026). Stochastic GOMA and RAIN-SL apply to equations, while the composite guarantee of E-Halpern is restricted to constrained problems; none is run on Problem 3.

For RAIN-SL, we use Algorithm 9 of Chen & Luo (2024), the single-loop experimental variant, and tune (η, λ, γ) over

$$\eta \in \{0.005, 0.01, 0.05, 0.1, 1, 5, 10\}, \qquad \lambda \in \{0.001, 0.01, 0.1, 1\}, \qquad \gamma \in \{0.001, 0.01, 0.1, 1\}.$$

All candidates use the same tuning seed within each problem, disjoint from the evaluation seeds. No candidate remains finite after 10^6 or 10^5 oracle queries, so we select the finite triple with the smallest squared residual after $12{,}000$ oracle queries. The selected triple is $(0.05, 0.001, 0.001)$ for Problem 1 and rock-paper-scissors and $(0.1, 0.001, 0.001)$ for Problem 2.

For E-Halpern, we use Algorithm 2 of Cai et al. (2022) and select ϵ using 50 independent calibration runs. Among 100 logarithmically spaced values from $0.8\mathcal{R}(z_0)$ to $\max\{10^{-4}\mathcal{R}(z_0), 10^{-6}\}$, we choose the smallest ϵ whose prescribed schedule uses at most 10^7 oracle queries on average. The selected values are 0.00413019, 0.00922484, and 0.02694452 for Problems 1, 2, and rock-paper-scissors, respectively. Each evaluation run stops before a batch would make the total exceed 10^7 oracle queries.

E.2 Plot construction

For E-Halpern, the query count of a completed iterate can differ across runs; at each checkpoint we use the last completed iterate within that budget and plot it at the mean actual query count.

All fixed-query solid curves use the same logarithmic checkpoints, include the common initial point, and are obtained from one trajectory.

The dashed curves use Theorem 4.1 for VRAF and Theorem 4 of Sohrabi et al. (2026) for stochastic GOMA. Following the experiments of Pethick et al. (2023), the BC-SEG+ solid curves show z_k, while the BC-PSEG+ solid curve shows the resolvent point $\bar{z}_k$. Their dashed curves use Theorems 6.1 and 7.1, respectively, which bound the corresponding α_k-weighted randomized outputs. For BC-PSEG+ on Problem 3, with $\alpha_0 = 1/18$, $\gamma = 1/2$, $L_{\widehat{F}} = \sqrt{1.01}$, and $\rho = 0$, the condition of Theorem 7.1 holds with $\eta \simeq 8.8995$ and $\mu \simeq 0.3133$. For RRSEG, Problems 1 and 2 use the sharper $A = 0$ last-iterate bound equation 63, while Problem 3 uses Corollary 6.2. For Problems 1 and 2, at each horizon $N \geq 19$ the sharper bound is

$$\left(\mathbb{E}\|F(z_N)\|^2\right)^{1/2} \leq \frac{36LD}{N} + \sigma\sqrt{\frac{5}{4(Q_N+1)}}\left(\log\frac{N}{18}+1\right), \qquad Q_N = \frac{N/18+1}{\log(N/18)}.$$

The bounds are instantiated separately at each horizon and squared for the squared-residual axis. These dashed RRSEG points are horizon-by-horizon theoretical guarantees; every empirical RRSEG solid curve comes from the single run described above. The displayed BC-SEG+ and stochastic GOMA assumptions hold on Problems 1 and 2 with $L_{\widehat{F}} = 1.2$ and $\kappa = 1$, respectively.

E.3 VRAF stepsize sensitivity

To assess VRAF (practical), we multiply the theorem-prescribed α_k by a common $c \in \{1, 2, 4, 8\}$ and leave all other coefficients unchanged. The same c is used on every problem, and the four curves share the same 50 seeds within each problem. Thus this is a sensitivity diagnostic rather than per-problem tuning; $c = 1$ is VRAF and $c = 4$ is VRAF (practical). Since the VRAF schedule is initialized by $\alpha_0 = 7/(12\bar{L})$ and the subsequent recurrence is multiplicative, multiplying all α_k by c is equivalent to using $\bar{L}/c$ in the stepsize schedule. The experiment therefore also measures robustness to multiplicative misspecification of the Lipschitz scale used by VRAF.

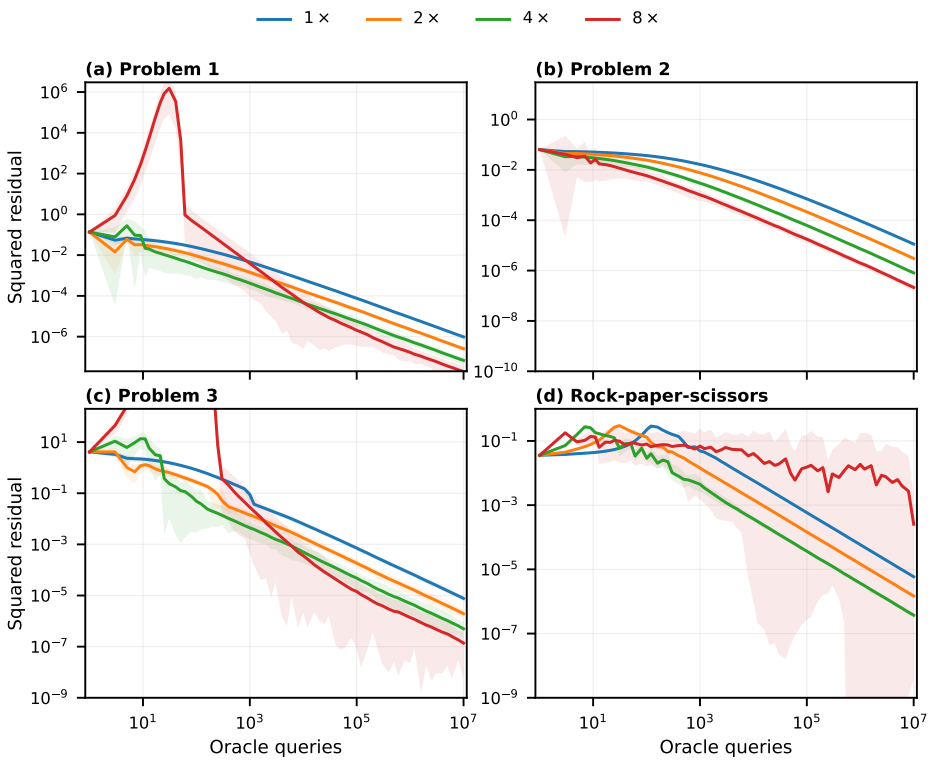

Figure 3: VRAF sensitivity to a common multiplier on α_k. Curves show the mean over 50 runs with pointwise 2.5%–97.5% quantiles.

E.4 Exact test problems and stochastic oracles

For Problems 1 and 2, write $z = (z^{(1)}, z^{(2)})$ and let $\phi(t) = t - \arctan t$. Both use $A = 0$, are monotone and 1-Lipschitz, and have $L_\Delta = 0.2$ and $\bar{L} = \sqrt{1.04}$.

Problem 1. We use $z_\star = (0.4, 0.6)$, $z_0 = (0.95, 0.05)$, and

$$F(z) = \frac{2}{3}\left(\phi\big(z^{(1)} - 0.4\big) + \frac{1}{4}\phi\big(z^{(1)} + z^{(2)} - 1\big),\ z^{(2)} - 0.6 + \frac{1}{4}\phi\big(z^{(1)} + z^{(2)} - 1\big)\right).$$

Here $\nabla F(z_\star) = \mathrm{diag}(0, 2/3)$, so the first coordinate is locally flat. The stochastic oracle is

$$\begin{aligned} F_\xi(z) = F(z) &+ 0.1\big(\rho_1 \tanh(z^{(1)} - 0.4), \rho_2 \tanh(z^{(2)} - 0.6)\big) \\ &+ 0.05\rho_3 \tanh\big(z^{(1)} + z^{(2)} - 1\big)(1,1) + 0.08(\rho_4, \rho_5), \end{aligned}$$

where $\rho_1, \ldots, \rho_5$ are independent Rademacher variables; we use $\sigma^2 = 0.0378$.

Problem 2. We use

$$F(z) = \frac{1}{5}\left(\phi(z^{(1)}) + 2\phi(z^{(1)} + z^{(2)}),\ \phi(z^{(2)}) + 2\phi(z^{(1)} + z^{(2)})\right),$$
$$z_\star = (0, 0), \qquad z_0 = (2, -2).$$

Here $\nabla F(z_\star) = 0$. The stochastic oracle is

$$\begin{aligned} F_\xi(z) = F(z) &+ 0.1\big(\rho_1 \tanh z^{(1)}, \rho_2 \tanh z^{(2)}\big) \\ &+ 0.05\rho_3 \tanh(z^{(1)} + z^{(2)})(1,1) + 0.08(\rho_4, \rho_5), \end{aligned}$$

where $\rho_1, \ldots, \rho_5$ are independent Rademacher variables; we use $\sigma^2 = 0.0378$.

Problem 3. Let $g(z) = \frac{1}{3}|z^{(1)}|$, $A = \partial g$, and $\phi(t) = t - \arctan t$. We use

$$F(z) = \frac{1}{3}\left(1 + \phi(z^{(1)}) + 2z^{(2)},\ z^{(2)} - 2z^{(1)}\right),$$
$$z_\star = (0,0), \qquad\qquad z_0 = (2,-2).$$

Since $F(z_\star) = (1/3, 0)$ and $A(z_\star) = [-1/3, 1/3] \times \{0\}$, $z_\star$ solves the inclusion. Moreover,

$$\frac{\nabla F(z) + \nabla F(z)^\top}{2} = \frac{1}{3}\begin{pmatrix}\phi'(z^{(1)}) & 0\\ 0 & 1\end{pmatrix} \succeq 0,$$

Since $0 \le \phi'(t) \le 1$, F is monotone and $\|\nabla F(z)\| \le (1+2)/3 = 1$. The inclusion is not locally strongly monotone: for $t > 0$, the common selection $(1/3, 0) \in A(t,0) \cap A(0,0)$ gives the monotonicity inner product $t\phi(t)/3 = o(t^2)$. The resolvent is

$$J_{\alpha A}(x,y) = (\operatorname{sign}(x)\max\{|x| - \alpha/3, 0\}, y),$$

so A has full domain and the RRSEG intermediate iterates have finite residuals. The stochastic oracle is

$$F_\xi(z) = F(z) + 0.1\big(\rho_1 \tanh z^{(1)}, \rho_2 \tanh z^{(2)}\big) + 0.35(\rho_3, \rho_4).$$

Here $\rho_1, \dots, \rho_4$ are independent Rademacher variables, $L_\Delta = 0.1$, $\bar{L} = \sqrt{1.01}$, and $\sigma^2 = 0.265$. The residual is

$$\mathcal{R}(z)^2 = \operatorname{dist}\Big(-F_1(z), \tfrac{1}{3}\partial|z^{(1)}|\Big)^2 + F_2(z)^2.$$

Rock-paper-scissors. Let

$$M_{\mathrm{RPS}} = \begin{pmatrix}0 & -1 & 1\\ 1 & 0 & -1\\ -1 & 1 & 0\end{pmatrix}, \qquad p = \operatorname{softmax}(x), \qquad q = \operatorname{softmax}(y),$$

and define $F(x,y) = (\nabla_x \Psi(x,y), -\nabla_y \Psi(x,y))$ for $\Psi(x,y) = p^\top M_{\mathrm{RPS}} q$. We use $x_\star = y_\star = 0$, $x_0 = (2,-1,-1)$, and $y_0 = (-1,2,-1)$. The stochastic oracle replaces M_{RPS} by

$$M_\xi = M_{\mathrm{RPS}} + 0.3\rho\ \operatorname{diag}(1,-1,0),$$

where ρ is Rademacher, and otherwise uses the same operator. We use $L = 0.5$, $L_\Delta = 0.075$, and $\sigma^2 = 0.045$; the logit-space operator is nonmonotone.

E.5 Rock-paper-scissors gap diagnostic

Softmax saturation can make the logit-space residual small even when the mixed strategies remain exploitable. We therefore also report

$$\begin{aligned}\mathrm{Gap}(p,q) &= \max_{q' \in \Delta_3} p^\top M_{\mathrm{RPS}} q' - \min_{p' \in \Delta_3} p'^\top M_{\mathrm{RPS}} q\\ &= -\min_i (M_{\mathrm{RPS}} p)_i - \min_i (M_{\mathrm{RPS}} q)_i,\end{aligned}$$

which vanishes at the uniform Nash equilibrium.

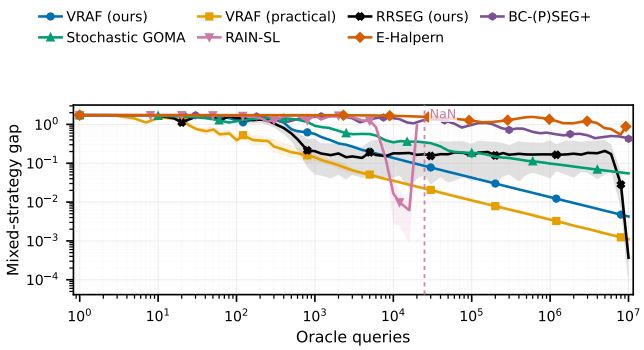

Figure 4: Mixed-strategy saddle-point gap for rock-paper-scissors versus the number of oracle queries. Solid curves and shaded regions use the same 50 runs and quantile convention as Figure 1.

F FURTHER DISCUSSION

F.1 ANYTIME WRAPPER FOR RRSEG

A standard doubling wrapper removes the need to specify a final horizon in advance. Fix $N_0 \geq 19$ and, for $s = 0, 1, \ldots$, run Corollary 6.2 from z_0 with horizon

$$N_s = 2^s N_0.$$

The runs are executed sequentially, and while the run with horizon N_{s+1} is in progress, the method returns the terminal output of the completed run with horizon N_s.

Suppose that a total of T RRSEG iterations have been performed and that the run with horizon N_s is the most recently completed one. Then

$$\sum_{j=0}^{s} N_j \leq T < \sum_{j=0}^{s+1} N_j < 4N_s.$$

Hence $N_s > T/4$. Applying Corollary 6.2 to the returned output gives

$$\mathbb{E}[\mathcal{R}(\widehat{z}_{N_s})^2] = O\left(\frac{L^2 D^2}{T^2} + \frac{\sigma^2 \log^3 T}{T}\right).$$

Each RRSEG iteration still uses two stochastic oracle queries, so the wrapper remains fixed-query.

F.2 VRAF NEED NOT CONVERGE POINTWISE

The residual guarantee of VRAF does not imply convergence of its iterates. The failure already occurs for a state-independent Gaussian oracle when every sampled operator is monotone.

Proposition F.1 (Pointwise failure of VRAF)**.** *Let $\mathcal{H} = \mathbb{R}$, $A = 0$, $F \equiv 0$, and $z_0 = 0$. Use the valid Lipschitz upper bound $L = 1$ and the state-independent Gaussian oracle*

$$F_\xi(z) = \xi, \qquad \xi \sim \mathcal{N}(0, \sigma^2), \qquad \sigma > 0.$$

Then Assumptions 2.1, 2.3, and 2.4 hold with $L_\Delta = 0$, and every sampled operator F_ξ is monotone. VRAF with the schedule of Theorem 4.1 satisfies

$$\mathcal{R}(z_k) = 0 \qquad \textit{for every } k,$$

but (z_k) does not converge in probability. In particular, it does not converge strongly almost surely.

Proof. Write $X_k := z_k$ and $W_k := v_k$. Since the same realization is used at z_k and z_{k+1}, the VRAF recursion becomes

$$X_{k+1} = \frac{k}{k+3} X_k - \alpha_k W_k, \qquad W_{k+1} = (1 - \gamma_{k+1}) W_k + \gamma_{k+1} \xi_{k+1}, \tag{64}$$

with $W_0 = \xi_0$. Iterating the first recursion gives, for $n \geq 1$,

$$X_n = -\frac{1}{n(n+1)(n+2)} \sum_{j=0}^{n-1} (j+1)(j+2)(j+3) \alpha_j W_j. \tag{65}$$

For $m \geq 1$, the schedule gives

$$\gamma_m = \frac{4m+9}{4(m+2)(m+3)}, \qquad 1 - \gamma_m = \frac{(2m+3)(2m+5)}{4(m+2)(m+3)}.$$

Hence

$$\gamma_m \geq \frac{1}{2(m+3)}, \qquad 1 - \gamma_m \geq \frac{m}{m+1}. \tag{66}$$

For $1 \le i \le j$, let $r_{j,i}$ be the coefficient of ξ_i in W_j. Equation equation 64 and equation 66 give

$$r_{j,i} = \gamma_i \prod_{m=i+1}^{j} (1-\gamma_m) \ge \frac{i+1}{2(i+3)(j+1)}. \tag{67}$$

Let $q_{n,i}$ denote the coefficient of ξ_i in X_n. By equation 65,

$$|q_{n,i}| = \frac{1}{n(n+1)(n+2)} \sum_{j=i}^{n-1} (j+1)(j+2)(j+3)\alpha_j r_{j,i}. \tag{68}$$

The Wallis estimate equation 27 and $\beta_j = 3/(j+3)$ imply

$$\alpha_j^2 \ge \frac{225\pi}{768(j+3)}.$$

Set $c_\alpha := \sqrt{75\pi/768}$. For $k \ge 6$, take

$$k \le i \le \left\lfloor \frac{5k}{4} \right\rfloor, \qquad \left\lceil \frac{3k}{2} \right\rceil \le j \le 2k-1.$$

Then $j \ge i$, $\alpha_j \ge c_\alpha/\sqrt{k}$, and equation 67 gives $r_{j,i} \ge 1/(8k)$. Moreover,

$$\frac{(j+1)(j+2)(j+3)}{2k(2k+1)(2k+2)} \ge \frac{1}{8}.$$

There are at least $k/3$ admissible indices j. Therefore equation 68 yields

$$|q_{2k,i}| \ge \frac{c_\alpha}{192\sqrt{k}} \qquad \left(k \le i \le \left\lfloor \frac{5k}{4} \right\rfloor\right). \tag{69}$$

Let $\mathcal{F}_k := \sigma(\xi_0, \dots, \xi_{k-1})$. The random variable X_k is $\mathcal{F}_k$-measurable, whereas the noises ξ_i with $i \ge k$ are independent of $\mathcal{F}_k$. Since all noises are Gaussian, $X_{2k} - X_k$ conditioned on $\mathcal{F}_k$ is Gaussian. Equation equation 69 gives

$$\begin{aligned} \mathrm{Var}(X_{2k} - X_k \mid \mathcal{F}_k) &\ge \sigma^2 \sum_{i=k}^{\lfloor 5k/4 \rfloor} q_{2k,i}^2 \\ &\ge \frac{c_\alpha^2}{4 \cdot 192^2} \sigma^2 =: \nu^2 > 0. \end{aligned}$$

For a Gaussian variable with variance at least ν^2, the probability of lying in an interval of length ν is maximized when its mean is the center of the interval. Thus

$$\mathbb{P}\left(|X_{2k} - X_k| \le \frac{\nu}{2} \middle| \mathcal{F}_k \right) \le 2\Phi(1/2) - 1 < 1,$$

where Φ is the standard normal distribution function. Hence $X_{2k} - X_k$ does not converge to zero in probability. The sequence (X_k) is not Cauchy in probability and therefore cannot converge in probability.

Finally, $F \equiv 0$ and $A = 0$, so every point is a solution and $\mathcal{R}(X_k) = 0$ for every k. □

The obstruction is tangential stochastic motion along the solution set. In the counterexample, the residual is identically zero, so no residual estimate can detect this motion. Any pointwise-convergence result for VRAF therefore requires an additional mechanism or stronger assumptions.