DUAL-ANCHOR ACCELERATION IS NEAR-OPTIMAL FOR STOCHASTIC MONOTONE ROOT-FINDING

Anonymous authors
Paper under double-blind review

ABSTRACT

Among distinct optimal acceleration mechanisms for deterministic monotone root-finding problems and fixed-point problems, dual-anchoring has recently been shown to admit a more robust direct stochastic extension than standard anchor acceleration. However, without additional strong monotonicity, the existing stochastic dual-anchoring guarantee by Yoon & Loizou (2026) has two limitations: first, it requires cocoercivity in expectation, and second, it attains only $\mathcal{O}(\epsilon^{-3})$ oracle complexity, leaving a gap to the near-optimal $\widetilde{\mathcal{O}}(\epsilon^{-2})$ complexity achieved by other methods. In this work, we address both of these limitations by combining dual-anchoring with stochastic resolvent approximation and optimized variance control. For unbiased stochastic oracles with variance bounded by σ^2, where sample operators are monotone and uniformly L-Lipschitz, our algorithm finds a point x_ϵ satisfying $\mathbb{E}\left[\|\mathbb{F}(x_\epsilon)\|\right] \leq \epsilon$ with a near-optimal oracle complexity of $\mathcal{O}((LD/\epsilon)\ell + (\sigma^2/\epsilon^2)\ell^2)$, where $\ell = \log(1 + LD/\epsilon)$ and D is the initial distance to a solution. This result improves the best known oracle complexity in the noise-dominated regime under these samplewise assumptions, reducing the poly-logarithmic factor from cubic to quadratic.

1 INTRODUCTION

We consider stochastic root-finding problem for an operator $\mathbb{F}\colon \mathbb{R}^d \to \mathbb{R}^d$:

$$\text{find } x_\star \in \mathbb{R}^d \quad \text{such that} \quad \mathbb{F}(x_\star) = 0 \tag{1}$$

where $\mathbb{F}(x) = \mathbb{E}\left[\mathbb{F}(x;\xi)\right]$ is monotone. This formulation includes convex minimization (Bauschke & Combettes, 2017; Ryu & Yin, 2022), convex-concave minimax optimization (Nemirovski, 2004; Nesterov, 2007; Juditsky et al., 2011; Lin et al., 2020), fixed-point problems (Krasnosel'skii, 1955; Mann, 1953; Halpern, 1967; Wittmann, 1992) or multiplayer game theory (Facchinei & Pang, 2003; Scutari et al., 2014). Our goal is to efficiently find a point with small operator norm $\|\mathbb{F}(\cdot)\|$, using only an unbiased stochastic oracle that returns a sample operator $\mathbb{F}(x;\xi)$. Here, the efficiency is measured by the number of oracle accesses required to achieve $\mathbb{E}\left[\|\mathbb{F}(x)\|\right] \leq \epsilon$.

In deterministic root-finding, Halpern-type or anchoring algorithms reduce the operator norm at the accelerated, optimal rate $\|\mathbb{F}(x_N)\| = \mathcal{O}(1/N)$ (Kim, 2021; Diakonikolas, 2020; Yoon & Ryu, 2021). However, the acceleration achieved by these algorithms does not directly transfer to stochastic settings due to accumulation of noise. Existing work that attempted to extend anchor acceleration to stochastic settings used diminishing variance assumption or variance reduction techniques (Lee & Kim, 2021; Cai et al., 2022; 2024). On the other hand, another type of deterministic acceleration mechanism called dual-anchoring (Yoon et al., 2024) has been recently shown to be more robust against oracle perturbation and stochastic noise, while enjoying the same optimal deterministic convergence guarantee as Halpern-type methods (Yoon & Grimmer, 2026; Yoon & Loizou, 2026). In particular, Yoon & Loizou (2026) showed that a simple stochastic dual-anchoring algorithm called S-DUAL-OHM with constant mini-batch size B achieves

$$\mathbb{E}\left[\|\mathbb{F}(x_{N-1})\|\right]^2 \leq \mathbb{E}\left[\|\mathbb{F}(x_{N-1})\|^2\right] \leq \frac{4L^2\|x_0 - x_\star\|^2}{N^2} + \frac{6\sigma^2}{B} \tag{2}$$

for problems with cocoercive $\mathbb{F}(x;\xi)$, where σ^2 denotes a bound on the oracle variance. Hence, taking $N = \mathcal{O}(\epsilon^{-1})$ and $B = \mathcal{O}(\epsilon^{-2})$ gives $\mathcal{O}(\epsilon^{-3})$ oracle complexity for achieving $\mathbb{E}\left[\|\mathbb{F}(x)\|\right] \leq \epsilon$, which

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.

already offers a clean improvement over stochastic anchoring-based approaches, which required additional algorithmic techniques to achieve the same complexity.

On the other hand, the result (2) was limited to cocoercive problems and did not cover the more general monotone stochastic root-finding problems. Furthermore, the $\mathcal{O}(\epsilon^{-3})$ complexity derived by (2) still has a significant gap to the $\Omega(\epsilon^{-2})$ lower bound from Chen & Luo (2024). *In this work, we address both of these issues and develop a conceptually simple, near-optimal algorithm for monotone stochastic problems based on dual-anchoring.* More precisely, we provide the following technical contributions.

1.1 Main contributions

First, we show that S-Dual-OHM with optimized scheduling of batch sizes B_k in fact attains an improved, near-optimal complexity of $\widetilde{\mathcal{O}}(\epsilon^{-2})$ for stochastic cocoercive root-finding (Theorem 2.2 & Corollary 2.3). Interestingly, this is achieved without fundamentally changing the update rule of S-Dual-OHM or its convergence analysis, and the algorithm remains single-loop.

Second, we extend the near-optimality of S-Dual-OHM with optimized variance control to stochastic monotone root-finding by using subroutines for approximately computing the resolvents. Assuming that each $\mathbb{F}(\cdot;\xi)$ is monotone and L-Lipschitz, the proposed algorithm has the complexity

$$\mathcal{O}\left(\ell + \frac{L\left\|x_0 - x_\star\right\|}{\epsilon}\ell + \frac{\sigma^2}{\epsilon^2}\ell^2\right) \text{ where } \ell = \log\left(1 + \frac{L\left\|x_0 - x_\star\right\|}{\epsilon}\right) \tag{3}$$

(Corollary 3.6). This complexity is near-optimal, and improves upon the previous state-of-the-art algorithm RAIN (Chen & Luo, 2024) that has cubic poly-logarithmic dependence on $\frac{L\|x_0-x_\star\|}{\epsilon}$ in the dominant term, although RAIN only requires monotonicity and Lipschitzness of $\mathbb{F}$.

1.2 Related work

Halpern-type acceleration and its stochastic extension. Halpern iteration (Halpern, 1967) originates in the theory of fixed point problems, which find a point $x_\star$ such that $x_\star = \mathbb{T}(x_\star)$ for nonexpansive $\mathbb{T}\colon \mathbb{R}^d \to \mathbb{R}^d$; it can accelerate the residual convergence of fixed-point iterations (Sabach & Shtern, 2017; Contreras & Cominetti, 2023) and in fact, attain exact optimal complexity matching the lower bound without even a constant factor gap (Lieder, 2021; Park & Ryu, 2022):

$$x_{k+1} = \frac{1}{k+2}x_0 + \frac{k+1}{k+2}\mathbb{T}x_k. \tag{OHM}$$

This idea, also called anchoring in the literature, was adapted to develop accelerated algorithms for monotone root-finding and minimax optimization (Diakonikolas, 2020; Yoon & Ryu, 2021; Lee & Kim, 2021; Tran-Dinh & Luo, 2021; Cai & Zheng, 2023). In stochastic problems, Cai et al. (2022; 2024) combined Halpern iteration with variance reduction to respectively achieve $\mathcal{O}(\epsilon^{-3})$ or improved complexity in certain finite-sum regime.

Algorithms and complexity for stochastic root-finding problems. Using stochastic variants of classical methods such as extragradient (Korpelevich, 1976) or optimistic gradient (Popov, 1980; Rakhlin & Sridharan, 2013a;b) typically yields a suboptimal $\mathcal{O}(\epsilon^{-4})$ oracle complexity (Diakonikolas et al., 2021; Gorbunov et al., 2022; Choudhury et al., 2023). Extending the plain anchoring mechanism of OHM results in accumulation of errors (Lee & Kim, 2021), which is why variance reduction was used in prior work (Cai et al., 2022; 2024) that explored stochastic Halpern-type algorithms. Chen & Luo (2024) proposed the Recursive Anchored Iteration (RAIN) algorithm together with a complexity lower bound, which uses amortized scheduling of regularization strength to achieve a near-optimal complexity $\widetilde{\mathcal{O}}\left(\frac{LD}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\right)$ for stochastic monotone root-finding problems. Yoon & Loizou (2026) considered an alternative approach of extending a distinct form of acceleration from deterministic fixed-point problem setting developed by Yoon et al. (2024):

$$x_{k+1} = x_k + \frac{N-k-1}{N-k}\left(\mathbb{T}(x_k) - \mathbb{T}(x_{k-1})\right) \tag{Dual-OHM}$$

where the total iteration number N is predetermined. They showed that Dual-OHM is more robust to stochastic noise, and $\mathcal{O}(\epsilon^{-3})$ complexity can be more easily attained without variance reduction or

recursive regularization as in the case of stochastic anchoring methods, albeit for a narrower class of cocoercive stochastic root-finding problems.

Proximal view of root-finding problems. Proximal-point methods provide a classical framework for solving monotone inclusions and convex-concave minimax problems using resolvent operations (Martinet, 1970; Rockafellar, 1976). Resolvents are often considered as *implicit* operations that are generally not computable. Approximating resolvents using forward/explicit evaluations of $\mathbb{F}$ is a well-established algorithmic principle in the literature. Extragradient and optimistic gradient methods admit interpretations as approximations to proximal-point iterations (Monteiro & Svaiter, 2010; 2012; Mokhtari et al., 2020a;b; Jiang & Mokhtari, 2025), This viewpoint was also combined with anchor acceleration to design more efficient algorithms (Kim, 2021; Diakonikolas, 2020) or to establish a formal connection between implicit and explicit algorithms (Suh et al., 2023; Yoon & Ryu, 2025). Our algorithmic design combines the idea of approximate resolvent computation with stochastic proximal point methods (Bianchi, 2016; Sadiev et al., 2024).

1.3 Preliminaries and assumptions

We assume throughout the paper that (1) has a solution $x_\star$. We say $\mathbb{F}$ is *monotone* if $\langle \mathbb{F}(x) - \mathbb{F}(y), x - y\rangle \geq 0$ for all $x, y \in \mathbb{R}^d$, and *L-Lipschitz* if $\|\mathbb{F}(x) - \mathbb{F}(y)\| \leq L\|x - y\|$ for all $x, y \in \mathbb{R}^d$. It is $1/L$*-cocoercive* if $\langle \mathbb{F}(x) - \mathbb{F}(y), x - y\rangle \geq \frac{1}{L}\|\mathbb{F}(x) - \mathbb{F}(y)\|^2$ for all $x, y \in \mathbb{R}^d$. While cocoercivity implies monotonicity and Lipschitz continuity, the converse need not hold. We say an operator $\mathbb{T}\colon \mathbb{R}^d \to \mathbb{R}^d$ is *nonexpansive* if $\|\mathbb{T}(x) - \mathbb{T}(y)\| \leq \|x - y\|$ for all $x, y \in \mathbb{R}^d$.

Denote the identity operator by $\mathbb{I}$. For a monotone continuous $\mathbb{F}$, its *resolvent* $\mathbb{J}_{\alpha\mathbb{F}} = (\mathbb{I} + \alpha\mathbb{F})^{-1}$ is well defined operator on $\mathbb{R}^d$ and is (firmly) nonexpansive for any $\alpha > 0$, and fixed points of $\mathbb{J}_{\alpha\mathbb{F}}$ are precisely the zeros of $\mathbb{F}$. When $\mathbb{F}$ is $1/L$-cocoercive, $\mathbb{I} - \alpha\mathbb{F}$ is nonexpansive for $0 < \alpha \leq \frac{2}{L}$ (Bauschke & Combettes, 2017).

Assumption 1.1 (Stochastic oracle)**.** The stochastic oracle operator $\mathbb{F}(x;\xi)$ is jointly measurable, unbiased and has bounded variance. That is, there exists $\sigma > 0$ such that for every $x \in \mathbb{R}^d$,

$$\mathbb{E}\left[\mathbb{F}(x;\xi)\right] = \mathbb{F}(x), \qquad \mathbb{E}\left[\left\|\mathbb{F}(x;\xi) - \mathbb{F}(x)\right\|^2\right] \leq \sigma^2.$$

Furthermore, for each sampled ξ, the stochastic operator $\mathbb{F}(\cdot;\xi)$ can be evaluated at multiple points.

At iteration k, we can draw a mini-batch $\mathcal{B}_k$ containing B_k independent samples, which defines $\mathbb{F}_{\mathcal{B}_k}(x) = \frac{1}{B_k}\sum_{\xi\in\mathcal{B}_k}\mathbb{F}(x;\xi)$. Whenever we use them, mini-batches are independent across iterations, newly drawn samples being independent of all previously drawn ones.

2 Stochastic cocoercive root-finding and fixed-point problems

We first consider stochastic root-finding problems with cocoercive $\mathbb{F} = \mathbb{E}\left[\mathbb{F}(\cdot;\xi)\right]$. This can be recast into stochastic fixed-point problem by taking $\mathbb{T}(\cdot;\xi) = \mathbb{I} - \alpha\mathbb{F}(\cdot;\xi)$ for appropriate α, so that $\mathbb{T} = \mathbb{E}\left[\mathbb{T}(\cdot;\xi)\right]$ is nonexpansive. For the analysis, we will take this fixed-point view. Throughout the section, $x_\star$ will denote a fixed point of $\mathbb{T}$, or equivalently, a zero of $\mathbb{F}$.

The main result of this section is to improve the $\mathcal{O}(\epsilon^{-3})$ complexity bound for S-Dual-OHM provided in Yoon & Loizou (2026) by optimizing the variance term contributed by each iteration number k. For the convergence analysis, it is convenient to consider the following abstraction: at iteration k, consider using a *random map* $\mathbb{T}_k$ at each iteration k such that $\mathbb{E}\left[\mathbb{T}_k(u)\right] = \mathbb{T}(u)$ and $\mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{T}(u)\right\|^2\right] \leq \sigma_k^2$ for all deterministic $u \in \mathbb{R}^d$. In particular, for any $u, v \in \mathbb{R}^d$ it satisfies

$$\mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{T}(v)\right\|^2\right] = \left\|\mathbb{T}(u) - \mathbb{T}(v)\right\|^2 + \mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{T}(u)\right\|^2\right] \leq \|u - v\|^2 + \sigma_k^2. \quad (4)$$

We will additionally require each $\mathbb{T}_k$ to be square-nonexpansive in expectation. Together, we state:

Condition 2.1. The jointly measurable random maps $\mathbb{T}_k$ are independent across iterations, and for every deterministic $u, v \in \mathbb{R}^d$, satisfy

$$\mathbb{E}\left[\|\mathbb{T}_k(u) - \mathbb{T}(v)\|^2\right] \le \|u - v\|^2 + \sigma_k^2, \quad \mathbb{E}\left[\|\mathbb{T}_k(u) - \mathbb{T}_k(v)\|^2\right] \le \|u - v\|^2, \tag{5}$$

where $\sigma_k^2 \ge 0$ is deterministic and the expectation is taken over randomness in $\mathbb{T}_k$.

Now fix $N \ge 2$ and consider the following formulation of S-DUAL-OHM:

$$x_{k+1} = x_k + \frac{N-k-1}{N-k}\left(\mathbb{T}_k(x_k) - \mathbb{T}_{k-1}(x_{k-1})\right) \tag{S-DUAL-OHM}$$

for $k = 0, \dots, N-2$, where $\mathbb{T}_{-1}(x_{-1}) := x_0$. Its final output, x_{N-1}, has small fixed-point residual.

Theorem 2.2. S-DUAL-OHM run with random maps $\mathbb{T}_0, \dots, \mathbb{T}_{N-2}$ satisfying Condition 2.1 and $x_0 \in \mathbb{R}^d$ such that $\|x_0 - x_\star\| \le D$ satisfies: with $a_{N,k} = \frac{1}{(N-k)(N-k-1)}$,

$$\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] \le \frac{4D^2}{N^2} + 6\sum_{k=0}^{N-2} a_{N,k}\sigma_k^2.$$

Observe that $\sum_k a_{N,k} = 1 - \frac{1}{N}$, so if each $\mathbb{T}_k$ is a mini-batch operator with a constant batch-size B, then we can take $\sigma_k^2 \equiv \frac{\alpha^2\sigma^2}{B}$ and the summation of variance terms is bounded by $\frac{6\alpha^2\sigma^2}{B}$, independently of N. This case precisely recovers (Yoon & Loizou, 2026, Theorem 4.1). We note that while Theorem 2.2 is more general than (Yoon & Loizou, 2026, Theorem 4.1), their proof essentially provides all ideas necessary to prove this result.

Our new insight here is that *the noise terms from earlier iterations are much less costly than the later ones* because they are weighted by $a_{N,k} = \frac{1}{(N-k)(N-k-1)}$. Therefore, we may allow larger σ_k for early k and enforce smaller σ_k later. For the special case with mini-batching, this means that we can allocate smaller batches to the earlier iterations without affecting the total noise term much. Optimizing this batch schedule yields a significant oracle complexity improvement, as detailed below.

2.1 MINI-BATCHING AND OPTIMIZED ALLOCATION OF THE ORACLE BUDGET

Consider the case where $\mathbb{T}(\cdot;\xi)$ is an unbiased stochastic oracle such that $\mathbb{E}[\mathbb{T}(x;\xi)] = \mathbb{T}(x)$, $\mathbb{E}\left[\|\mathbb{T}(x;\xi) - \mathbb{T}(x)\|^2\right] \le \tau^2$ and $\mathbb{E}\left[\|\mathbb{T}(x;\xi) - \mathbb{T}(y;\xi)\|^2\right] \le \|x - y\|^2$ for any deterministic $x, y \in \mathbb{R}^d$, and we use the mini-batch operator

$$\mathbb{T}_k(x) = \frac{1}{B_k}\sum_{\xi \in \mathcal{B}_k} \mathbb{T}(x;\xi) \tag{6}$$

where each $\mathcal{B}_k$ is a uniformly and independently sampled mini-batch of size B_k. In this case, Condition 2.1 is satisfied with $\sigma_k^2 = \tau^2/B_k$, and one evaluation of $\mathbb{T}_k$ costs B_k oracle evaluations. The total noise term given by Theorem 2.2 is therefore proportional to $A_N := \sum_{k=0}^{N-2} \frac{a_{N,k}}{B_k}$. Now treating batch-sizes as real numbers, and assuming the total oracle access (complexity) constraint $\sum_{k=0}^{N-2} B_k \le Q$, we can optimize B_k to minimize the noise term, using Cauchy–Schwarz:

$$\left(\sum_{k=0}^{N-2}\sqrt{a_{N,k}}\right)^2 \le \left(\sum_{k=0}^{N-2}\frac{a_{N,k}}{B_k}\right)\left(\sum_{k=0}^{N-2} B_k\right). \tag{7}$$

Equality holds when B_k is proportional to $\sqrt{a_{N,k}}$. Hence $B_k = \frac{Q\sqrt{a_{N,k}}}{S_N}$, where $S_N = \sum_{k=0}^{N-2}\sqrt{a_{N,k}}$, is the optimal allocation satisfying $A_N = \frac{S_N^2}{Q}$. Now rounding up each B_k yields

$$B_k = \left\lceil \frac{Q\sqrt{a_{N,k}}}{S_N} \right\rceil, \qquad \sum_{k=0}^{N-2} B_k \le Q + N - 1, \qquad \sum_{k=0}^{N-2}\frac{a_{N,k}}{B_k} \le \frac{S_N^2}{Q}. \tag{8}$$

Note that $S_N \le \sum_{k=0}^{N-2}\frac{1}{N-k-1} \le 1 + \log N$.

Corollary 2.3. Let $\epsilon > 0$. S-DUAL-OHM run with the mini-batch operators (6), with

$$N = \max\{2, \lceil 4D/\epsilon \rceil\}, \quad Q = \max\{1, 8\tau^2 S_N^2/\epsilon^2\}$$

and the batch schedule (8) attains $\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] \le \epsilon^2$ using

$$\mathcal{O}\left(\frac{D}{\epsilon} + \frac{\tau^2}{\epsilon^2}\left(1 + \log\left(1 + \frac{D}{\epsilon}\right)\right)^2\right) \tag{9}$$

oracle evaluations.

Proof. By Theorem 2.2 and (8), we have

$$\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] \le \frac{4D^2}{N^2} + 6\tau^2 \sum_{k=0}^{N-2} \frac{a_{N,k}}{B_k} \le \frac{4D^2}{N^2} + \frac{6\tau^2 S_N^2}{Q} \le \frac{\epsilon^2}{4} + \frac{3\epsilon^2}{4} = \epsilon^2.$$

Because $N \le 2 + \left\lceil \frac{4D}{\epsilon} \right\rceil$ and $Q \le 1 + \frac{8\tau^2 S_N^2}{\epsilon^2} \le 1 + \frac{8\tau^2(1+\log N)^2}{\epsilon^2}$, the total complexity is at most

$$\sum_{k=0}^{N-2} B_k \le Q + N - 1 \le 2 + \left\lceil \frac{4D}{\epsilon} \right\rceil + \frac{8\tau^2}{\epsilon^2}\left(1 + \log\left(2 + \left\lceil \frac{4D}{\epsilon} \right\rceil\right)\right)^2.$$

□

2.2 SPECIALIZATION TO STOCHASTIC COCOERCIVE ROOT-FINDING

Consider the operator oracle of Assumption 1.1. Suppose additionally that $\mathbb{F}(\cdot;\xi)$ is cocoercive in expectation, i.e., for every $u, v \in \mathbb{R}^d$,

$$\mathbb{E}\left[\|\mathbb{F}(u;\xi) - \mathbb{F}(v;\xi)\|^2\right] \le L \langle \mathbb{F}(u) - \mathbb{F}(v), u - v \rangle . \tag{10}$$

For $0 < \alpha \le 2/L$, taking $\mathbb{T} = \mathbb{I} - \alpha\mathbb{F}$, $\mathbb{T}(\cdot;\xi) = \mathbb{I} - \alpha\mathbb{F}(\cdot;\xi)$ and $\mathbb{T}_k = \mathbb{I} - \alpha\mathbb{F}_{\mathcal{B}_k}$ satisfies Condition 2.1 with $\tau^2 = \alpha^2\sigma^2$ and $\sigma_k^2 = \frac{\alpha^2\sigma^2}{B_k}$, and Theorem 2.2 gives

$$\mathbb{E}\left[\|\mathbb{F}(x_{N-1})\|^2\right] \le \frac{4D^2}{\alpha^2 N^2} + 6\sigma^2 \sum_{k=0}^{N-2} \frac{a_{N,k}}{B_k}.$$

With batch-sizes (8) and $\alpha = \frac{1}{L}$, Corollary 2.3 with $\alpha\epsilon$ in place of ϵ yields $\mathbb{E}\left[\|\mathbb{F}(x_{N-1})\|^2\right] \le \epsilon^2$ in

$$\mathcal{O}\left(\frac{LD}{\epsilon} + \frac{\sigma^2}{\epsilon^2}\log^2\left(\frac{LD}{\epsilon}\right)\right)$$

operator evaluations.

3 STOCHASTIC MONOTONE ROOT-FINDING

Now we extend the result of the previous section to stochastic root-finding for monotone and Lipschitz operators. Recall that the resolvent $\mathbb{T} = \mathbb{J}_{\alpha\mathbb{F}}$ of a maximally monotone operator $\mathbb{F}$ is nonexpansive; we leverage this connection to again take the fixed-point view as in Section 2. Because resolvents are not *exactly* computable, we run inner loops to compute them *approximately*. These will yield random maps $\mathbb{T}_k$, which may no longer be unbiased estimators of $\mathbb{T}$ satisfying Condition 2.1. We instead require the following weaker condition. Later, we show how to select concrete $\mathbb{T}_k$ satisfying it.

Condition 3.1. The jointly measurable random maps $\mathbb{T}_k$ are independent across iterations, and for every deterministic $u, v \in \mathbb{R}^d$, satisfy

$$\mathbb{E}\left[\|\mathbb{T}_k(u) - \mathbb{T}(v)\|^2\right] \le \|u - v\|^2 + \rho_k\left(\|u - x_\star\|^2 + \|v - x_\star\|^2\right) + \sigma_k^2 \tag{11}$$

$$\mathbb{E}\left[\|\mathbb{T}_k(u) - \mathbb{T}_k(v)\|^2\right] \le \|u - v\|^2 + \rho_k\left(\|u - x_\star\|^2 + \|v - x_\star\|^2 + \sigma_k^2\right) \tag{12}$$

where $\sigma_k^2, \rho_k \ge 0$ are deterministic and the expectation is taken over randomness in $\mathbb{T}_k$.

3.1 Outer bound under generalized condition on $\mathbb{T}_k$

Now we present an analogue of Theorem 2.2 under the relaxed Condition 3.1. This result exactly recovers Theorem 2.2 as a special case with $\rho_k \equiv 0$, so the proof of Theorem 3.2, which we present in Appendix A, accounts for the results from the previous section.

Theorem 3.2. S-Dual-OHM run with random maps $\mathbb{T}_0, \dots, \mathbb{T}_{N-2}$ satisfying Condition 3.1 and $x_0 \in \mathbb{R}^d$ such that $\|x_0 - x_\star\| \le D$ satisfies: with $a_{N,k} = \frac{1}{(N-k)(N-k-1)}$ and $R = \sum_{k=0}^{N-2} \rho_k$,

$$\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] \le \frac{4D^2}{N^2} + 6\sum_{k=0}^{N-2} a_{N,k}\sigma_k^2 + 15Re^R\left(D^2 + \sum_{k=0}^{N-2}\sigma_k^2\right).$$

In particular, when $\mathbb{T} = \mathbb{J}_{\alpha\mathbb{F}}$, we obtain the following bound on the operator norm.

Corollary 3.3. Let $\mathbb{F}\colon \mathbb{R}^d \to \mathbb{R}^d$ be monotone and L-Lipschitz, and $\mathbb{T} = \mathbb{J}_{\alpha\mathbb{F}}$ for some $\alpha \in \left(0, \frac{1}{L}\right]$. Fix $N \ge 2$, $x_0 \in \mathbb{R}^d$ and let $D = \|x_0 - x_\star\|$. Let S-Dual-OHM be run with random maps $\mathbb{T}_0, \dots, \mathbb{T}_{N-1}$ satisfying Condition 3.1, and set $y_N = \mathbb{T}_{N-1}(x_{N-1})$. Then, with $R = \sum_{k=0}^{N-1} \rho_k$,

$$\mathbb{E}\left[\|\mathbb{F}(y_N)\|^2\right] \le \frac{8D^2}{\alpha^2 N^2} + \frac{24}{\alpha^2}\sum_{k=0}^{N-1}\frac{\sigma_k^2}{(N-k)^2} + \frac{30}{\alpha^2}Re^R\left(D^2 + \sum_{k=0}^{N-1}\sigma_k^2\right). \tag{13}$$

3.2 Inner loop for stable approximate resolvents $\mathbb{T}_k$

Our choice of inner loop algorithm, i.e., the construction of $\mathbb{T}_k$, will require the following samplewise monotonicity and Lipschitzness in addition to Assumption 1.1. Note that it implies the monotonicity and L-Lipschitzness of the mean operator $\mathbb{F}$, but not vice versa.

Assumption 3.4. $\mathbb{F}(\cdot;\xi)$ is almost surely monotone and L-Lipschitz on $\mathbb{R}^d$, with a common $L > 0$.

In Section 2, we achieved reduced variance σ_k^2 by allocating more sample budgets to iteration k, via mini-batching. Here we aim to choose $\mathbb{T}_k$ similarly, ensuring Condition 3.1 with smaller σ_k^2 given a larger number of subroutine (inner loop) iterations. First, suppose ideally that we can compute stochastic proximal operations exactly. Given sample budget $B_k \ge 1$, first consider $\widehat{\mathbb{T}}_k$ defined by:

$$\begin{cases} z_0 = u \\ z_{j+1} = \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}\big((1-\gamma_j)z_j + \gamma_j u\big) \text{ for } j = 0, \dots, B_k - 1 \\ \widehat{\mathbb{T}}_k(u) = z_{B_k} \end{cases} \tag{14}$$

where we use $\alpha = \frac{1}{2L}$, $\gamma_j = \frac{2}{j+2}$ and each ξ_j is a fresh sample drawn independently of all preceding ones. Following the technical treatment of stochastic proximal point methods similar to Sadiev et al. (2024), we can show that $\widehat{\mathbb{T}}_k$ is almost surely nonexpansive, and additionally satisfies

$$\mathbb{E}\left[\left\|\widehat{\mathbb{T}}_k(u) - \mathbb{T}(v)\right\|^2\right] \le \|u - v\|^2 + \frac{4\alpha^2\sigma^2}{B_k + 1} \tag{15}$$

for any $u, v \in \mathbb{R}^d$, as we show in Appendix B. Hence $\widehat{\mathbb{T}}_k$ satisfies Condition 2.1 with $\sigma_k^2 = \frac{4\alpha^2\sigma^2}{B_k+1}$.

However, because $\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)$ is generally not directly computable, we can only compute it up to errors, via forward evaluations $\mathbb{F}(\cdot;\xi_j)$. Nevertheless, this is not a significant overhead because it is equivalent to finding the zero of a 1-strongly monotone operator $z \mapsto \mathbb{I} + \alpha\gamma_j\mathbb{F}(z;\xi_j) - p_j$, which converges rapidly with simple operator descent using step-size 1. Algorithm 1 presents an approximate version of (14) replacing the resolvent step with this descent loop.

Lemma 3.5. Under Assumptions 1.1 and 3.4, the random map $\mathbb{T}_k$ defined via Algorithm 1 with $\alpha = \frac{1}{2L}$ and $\gamma_j = \frac{2}{j+2}$ satisfies Condition 3.1 with $\sigma_k^2 = \frac{6\alpha^2\sigma^2}{B_k+1}$ and uses $\mathcal{O}\left(B_k\left[1 + \frac{\log(16/\rho_k)}{\log(B_k+1)}\right]\right)$ oracle evaluations.

Algorithm 1 Approximate resolvent $\mathbb{T}_k$

Require: Input u, Sample budget $B_k \geq 1$, Tolerance $\rho_k \in (0, 1]$, Lipschitz constant L, Hyperparameters α, γ_j

$w_0 \leftarrow u$,
for $j = 0, \ldots, B_k - 1$ **do**
  Draw ξ_j
  $p_j \leftarrow (1 - \gamma_j)w_j + \gamma_j u$
  $q_0 \leftarrow p_j, m_{k,j} \leftarrow \left\lceil \frac{\log(16/\rho_k)}{\log(j+2)} \right\rceil$
  for $r = 0, \ldots, m_{k,j}$ **do**
    $q_{r+1} \leftarrow p_j - \alpha\gamma_j \mathbb{F}(q_r; \xi_j)$
  end for
  $w_{j+1} \leftarrow q_{m_{k,j}+1}$
end for
$\mathbb{T}_k(u) \leftarrow w_{B_k}$
return $\mathbb{T}_k(u)$

The parameter ρ_k appears in Condition 3.1 and controls the tolerance to violation of nonexpansiveness in expectation. B_k is the sample budget, controlling the variance bound σ_k^2.

3.3 Optimized allocation of B_k and complexity analysis

Fix $N \geq 2$ and choose $\rho_k \equiv \frac{1}{2N^3}$, so that $Re^R = \mathcal{O}\left(\frac{1}{N^2}\right)$ in Corollary 3.3, so together with Lemma 3.5, the total variance term in (13) is bounded by a constant multiple of $A_N := \sum_{k=0}^{N-1} \frac{1}{(B_k+1)(N-k)^2} \leq \sum_{k=0}^{N-1} \frac{1}{B_k(N-k)^2}$. Lemma 3.5 also shows that the total complexity is

$$\mathcal{O}\left(\sum_{k=0}^{N-1} B_k + \log(eN) \sum_{k=0}^{N-1} \frac{B_k}{\log(B_k+1)}\right) = \mathcal{O}\left(\log(eN)\left[N + \sum_{k=0}^{N-1} \frac{B_k}{\log(eN/(N-k))}\right]\right)$$

where we show in Appendix B.3 that the latter bound holds for any $B_k \geq 1$. Now to optimize the allocation of B_k, similarly as in Section 2.1, we temporarily view B_k as real numbers and use Cauchy-Schwarz to minimize A_N while keeping $Q := \sum_{k=0}^{N-1} \frac{B_k}{\log(eN/(N-k))}$ fixed:

$$\left(\sum_{k=0}^{N-1} \frac{1}{(N-k)\sqrt{\log(eN/(N-k))}}\right)^2 \leq \left(\sum_{k=0}^{N-1} \frac{1}{(N-k)^2 B_k}\right)\left(\sum_{k=0}^{N-1} \frac{B_k}{\log(eN/(N-k))}\right).$$

Equality holds when B_k is proportional to $\frac{1}{N-k}\sqrt{\log \frac{eN}{N-k}}$. More precisely, we choose:

$$B_k = \max\left\{1, \left\lceil \frac{1024\sigma^2}{\epsilon^2(N-k)} \sqrt{\log(eN) \log\left(\frac{eN}{N-k}\right)} \right\rceil\right\} \tag{16}$$

for $k = 0, \ldots, N-1$. The resulting complexity bound is as follows, and we present its detailed proof in Appendix B.3.

Corollary 3.6 (Oracle complexity for monotone root-finding)**.** Under Assumptions 1.1 and 3.4, let $\epsilon > 0$, $\|x_0 - x_\star\| \leq D$. Choose $N = \max\left\{2, \left\lceil \frac{32LD}{\epsilon} \right\rceil\right\}$ and $\rho_k \equiv \frac{1}{2N^3}$. Run S-DUAL-OHM with $\mathbb{T}_k$ computed by Algorithm 1, using batch schedule (16). Then its final output $y_N = \mathbb{T}_{N-1}(x_{N-1})$ attains $\mathbb{E}\left[\|\mathbb{F}(y_N)\|^2\right] \leq \epsilon^2$ using

$$\mathcal{O}\left(\left(1 + \frac{LD}{\epsilon}\right) \log\left(1 + \frac{LD}{\epsilon}\right) + \frac{\sigma^2}{\epsilon^2} \log^2\left(1 + \frac{LD}{\epsilon}\right)\right) \tag{17}$$

operator evaluations.

4 NUMERICAL EXPERIMENTS

We evaluate the empirical effectiveness of S-DUAL-OHM for cocoercive and monotone stochastic root-finding problems. We consider two problems, both with stochastic oracle of the form $\mathbb{F}(x;\xi) = \mathbb{F}(x) + \xi$ where $\xi \sim \mathcal{N}\left(0, \frac{\sigma^2 I_d}{d}\right)$, and use the residual norm $\|\mathbb{F}(x_k)\|$ as a measure of convergence.

Baselines. Denote by $\mathcal{B}_k$ and $\mathcal{B}'_k$ independent mini-batches of fixed size B, and let $\mathbb{F}_{\mathcal{B}_k}$ and $\mathbb{F}_{\mathcal{B}'_k}$ be the corresponding mini-batch operators. We consider the following basic algorithms: Stochastic Gradient Descent-Ascent (SGDA), Stochastic OHM (S-OHM) and Stochastic Extragradient (SEG), where the first two algorithms are used in Experiment 1 where $\mathbb{F}$ is cocoercive, and SEG is used for both experiments. Their update rules are given by

$$\begin{aligned}
\text{SGDA}: &\quad x_{k+1} = x_k - \eta \mathbb{F}_{\mathcal{B}_k}(x_k) \\
\text{SEG}: &\quad x_{k+1/2} = x_k - \eta \mathbb{F}_{\mathcal{B}_k}(x_k), \qquad x_{k+1} = x_k - \eta \mathbb{F}_{\mathcal{B}'_k}(x_{k+1/2}) \\
\text{S-OHM}: &\quad x_{k+1} = \frac{1}{k+2} x_0 + \frac{k+1}{k+2}(x_k - \eta \mathbb{F}_{\mathcal{B}_k}(x_k))
\end{aligned}$$

In Experiment 2, we also include SEG with averaging, which uses separately tuned parameters. Both experiments also use the practical single-loop RAIN implementation from Chen & Luo (2024):

$$x_{k+1/2} = x_k - \eta\left(\mathbb{F}_{\mathcal{B}_k}(x_k) + r_k(x_k)\right), \qquad x_{k+1} = x_k - \eta\left(\mathbb{F}_{\mathcal{B}'_k}(x_{k+1/2}) + r_k(x_{k+1/2})\right)$$

where

$$r_k(v) = \lambda\gamma \sum_{t=0}^{k-1} (1+\gamma)^t (v - x_t), \qquad \gamma = e^{g/T} - 1, \qquad \lambda = \frac{a}{e^g - 1}, \qquad \eta = \frac{c}{L + a}$$

and g, a, c are tunable hyperparameters. For all baseline algorithms, we use a fixed oracle budget Q and take $T = \lfloor Q/B \rfloor$ for SGDA and S-OHM, and $T = \lfloor Q/(2B) \rfloor$ for SEG and RAIN.

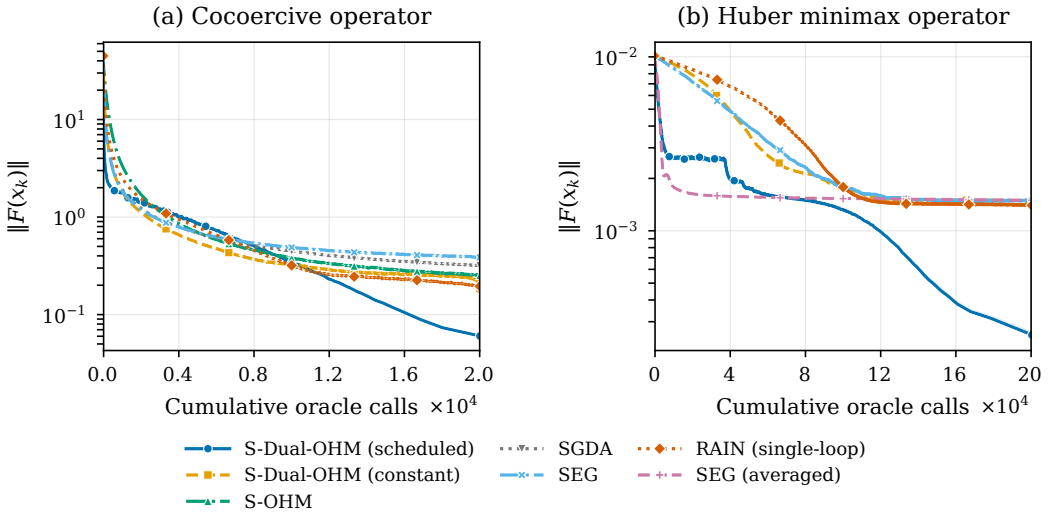

Figure 1: Mean residual norms over ten independent runs. The horizontal axis counts every sampled-operator evaluation, including all samples in mini-batches and inner solves.

Experiment 1: Cocoercive problem. We use the worst-case affine operator construction from Park & Ryu (2022), which has also been considered in Yoon & Loizou (2026). Let $d = 1001$, and define

$$\mathbb{H}(x) = \left(x_d - \frac{2}{\sqrt{d}}, -x_1, \ldots, -x_{d-1}\right), \qquad \mathbb{F}(x) = \mathbb{H}(x - s) + (x - s), \qquad \mathbb{T} = \mathbb{I} - \eta\mathbb{F}$$

where $s \sim \mathcal{N}(0, I_d)$ is drawn and fixed across all runs. The operator $\mathbb{F}$ is $1/2$-cocoercive and has a unique zero $x_\star = s + \frac{1}{\sqrt{d}}\mathbf{1}$. We use $x_0 = 0$, $\sigma = 2$, and a budget of 2×10^4 oracle

calls. Our proposed algorithm is the *scheduled* version of S-DUAL-OHM, which uses the allocation profile in (8), normalized to the oracle budget. We provide the full details in Appendix C. By contrast, S-DUAL-OHM *(constant)* denotes the version using a constant batch-size B. For that algorithm, SGDA and S-OHM, we tune B over the grid $\{1, 2, 4, 8, 16, 32, 64, 128, 256\}$, and η over $\{0.005, 0.01, 0.025, 0.05, 0.1, 0.25, 0.5, 0.75, 1\}$. SEG uses the same grid for B and $\eta \in \{0.005, 0.01, 0.025, 0.05, 0.1, 0.2, 0.3, 0.45\}$. RAIN uses $B \in \{1, 4, 16, 64\}$, $g \in \{1, 3, 6\}$, $a \in \{0.001, 0.01, 0.1, 1\}$, and $c \in \{0.02, 0.05, 0.1, 0.2, 0.5, 0.8\}$. We observe that RAIN and S-DUAL-OHM (constant) are the most competitive among the baselines, and notably, S-DUAL-OHM (scheduled) outperforms its constant-batch counterpart by a large margin (Fig. 1(a)), even though they are sharing the same designing principle and only the batch schedule is chosen differently.

Experiment 2: Beyond cocoercivity. We use the Huber minimax problem used in Chen & Luo (2024): writing $c_\nu(t) = \min\{\nu, \max\{-\nu, t\}\}$ componentwise, its saddle operator is

$$\mathbb{F}(x^{(1)}, x^{(2)}) = \Big((1-\delta)c_\nu\big(x^{(1)}\big) + \delta x^{(2)}, (1-\delta)c_\nu\big(x^{(2)}\big) - \delta x^{(1)}\Big)$$

where $x = (x^{(1)}, x^{(2)}) \in \mathbb{R}^{1024}$ and we use $\delta = 0.01, \nu = 5 \times 10^{-5}$. This operator is monotone and 1-Lipschitz but not cocoercive. We choose a fixed random unit vector x_0, $\sigma = 0.032$, and a budget of 2×10^5 calls. Our scheduled S-DUAL-OHM uses $\alpha = \frac{1}{2}$ and the allocation profile in (16), normalized to the oracle budget for each candidate N. We provide the full details in Appendix C. For Fig. 1(b), we use a practical variant of Algorithm 1 where we use constant-length loop steps ($m_{k,j} \equiv m$), while we also report below the residual value from Algorithm 1 run with $m_{k,j}$ set precisely as the theory prescribes, which are more conservative at the tested oracle budget.

After a joint search over m and N, we fix $m = 1$ (two descent iterations per inner sample) and refine N separately for the two variants. We select $N = 24000$ for S-DUAL-OHM (scheduled) and $N = 3200$ for S-DUAL-OHM (constant). SEG selects $B \in \{1, 4, 16, 64, 256\}$ and $\eta \in \{0.005, 0.01, 0.025, 0.05, 0.1, 0.2, 0.4, 0.7, 0.9, 1, 2, 5, 10\}$. For RAIN, we first sweep the grid $(B, g, a, c) \in \{1, 8, 64\} \times \{1, 3, 6\} \times \{0.0001, 0.001, 0.01, 0.1, 1\} \times \{0.02, 0.05, 0.1, 0.2, 0.5, 0.8, 1, 2, 5, 10\}$, and later refine the search using the union of $\{32, 64, 128\} \times \{2, 3, 4, 6\} \times \{0.003, 0.01, 0.03\} \times \{0.35, 0.5, 0.7\}$ and $\{1\} \times \{6\} \times \{0.03, 0.1, 0.3\} \times \{0.005, 0.01, 0.015, 0.02, 0.03, 0.05\}$.

Fig. 1(b) shows that the practical variant of scheduled S-DUAL-OHM achieves the smallest residual of 2.55×10^{-4}. The final residual values were 1.47×10^{-3} for S-DUAL-OHM (constant) and 1.41×10^{-3} for RAIN. We also test the prescribed schedule in Algorithm 1 and Lemma 3.5, and observe that it reaches final residual norm 1.46×10^{-3} at its tuned horizon $N = 2400$.

5 CONCLUSION

In this work, we show that the robustness of dual-anchor acceleration observed in recent prior work can be translated into near-optimal oracle efficiency for stochastic root-finding. This is achieved by carefully controlling how errors at different iterations contribute to the residual bound. In the cocoercive setting, without materially changing the underlying dual-anchor acceleration mechanism but only changing the batch-size schedules, we close the polynomial complexity gap to the lower bound. We show that the same principle extends to samplewise monotone and Lipschitz problems via stochastic resolvent approximations. These results strengthen the motivation for studying distinct deterministic acceleration mechanisms; their different responses to noise can provide alternative, clean routes to designing efficient stochastic algorithms.

Our analysis highlights the separation between choosing the acceleration mechanism governing the outer loop and constructing random maps that meet suitable requirements. While we have used inner loops to handle these requirements, the broader framework we propose is not necessarily confined to algorithms with nested loop structures. Natural directions for future work include designing a single-loop algorithm for monotone problems from the same principles, removing the logarithmic factors from oracle complexity, developing adaptive budget schedules that require less knowledge of the problem, or extending our framework to horizon-free algorithms.

AI USE STATEMENT

We used generative AI tools (ChatGPT and Codex) to assist with mathematical derivations, manuscript drafting and editing, and implementation of the numerical experiments. The authors checked and rederived the mathematical arguments, verified the references, revised the manuscript, and take full responsibility for the final content.

ETHICS STATEMENT

This work develops theoretical algorithms for stochastic root-finding and evaluates them on synthetic problems. It involves no human participants or personal data. We do not identify ethical concerns specific to the research presented here.

REPRODUCIBILITY STATEMENT

The assumptions and algorithms are specified in Sections 1.3–3, with complete proofs provided in the appendix. The experimental setup and hyperparameter search grids are described in Section 4.

REFERENCES

Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer International Publishing, 2nd edition, 2017.

Pascal Bianchi. Ergodic convergence of a stochastic proximal point algorithm. *SIAM Journal on Optimization*, 26(4):2235–2260, 2016.

Xufeng Cai, Chaobing Song, Cristóbal A Guzmán, and Jelena Diakonikolas. Stochastic halpern iteration with variance reduction for stochastic monotone inclusions. *Neural Information Processing Systems*, 2022.

Xufeng Cai, Ahmet Alacaoglu, and Jelena Diakonikolas. Variance reduced halpern iteration for finite-sum monotone inclusions. *International Conference on Learning Representations*, 2024.

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

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.

Sayantan Choudhury, Eduard Gorbunov, and Nicolas Loizou. Single-call stochastic extragradient methods for structured non-monotone variational inequalities: Improved analysis under weaker conditions. *Neural Information Processing Systems*, 2023.

Juan Pablo Contreras and Roberto Cominetti. Optimal error bounds for non-expansive fixed-point iterations in normed spaces. *Mathematical Programming*, 199(1):343–374, 2023.

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

Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient Methods for Structured Nonconvex-Nonconcave Min-Max Optimization. *International Conference on Artificial Intelligence and Statistics*, 2021.

Francisco Facchinei and Jong-Shi Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems*. Springer-Verlag, 2003.

Eduard Gorbunov, Hugo Berard, Gauthier Gidel, and Nicolas Loizou. Stochastic extragradient: General analysis and improved rates. *International Conference on Artificial Intelligence and Statistics*, 2022.

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

Ruichen Jiang and Aryan Mokhtari. Generalized optimistic methods for convex-concave saddle point problems. *SIAM Journal on Optimization*, 35(3):2066–2097, 2025.

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.

G. M. Korpelevich. The extragradient method for finding saddle points and other problems. *Ekonomika i Matematicheskie Metody*, 12(4):747–756, 1976.

M. A. Krasnosel'skii. Two remarks on the method of successive approximations. *Uspekhi Matematicheskikh Nauk*, 10(1):123–127, 1955.

Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex–nonconcave minimax problems. *Neural Information Processing Systems*, 2021.

Felix Lieder. On the convergence rate of the Halpern-iteration. *Optimization Letters*, 15(2):405–418, 2021.

Tianyi Lin, Chi Jin, and Michael I. Jordan. Near-optimal algorithms for minimax optimization. *Conference on Learning Theory*, 125, 2020.

William Robert Mann. Mean value methods in iteration. *Proceedings of the American Mathematical Society*, 4(3):506–510, 1953.

B. Martinet. Régularisation d'inéquations variationnelles par approximations successives. *Revue Française d'Informatique et de Recherche Opérationnelle, Série Rouge*, 4(3):154–158, 1970.

Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. *International Conference on Artificial Intelligence and Statistics*, 2020a.

Aryan Mokhtari, Asuman E. Ozdaglar, and Sarath Pattathil. Convergence rate of O(1/k) for optimistic gradient and extragradient methods in smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 30(4):3230–3251, 2020b.

Renato D. C. Monteiro and B. F. Svaiter. On the complexity of the hybrid proximal extragradient method for the iterates and the ergodic mean. *SIAM Journal on Optimization*, 20(6):2755–2787, 2010.

Renato D. C. Monteiro and Benar F. Svaiter. Iteration-complexity of a newton proximal extragradient method for monotone variational inequalities and inclusion problems. *SIAM Journal on Optimization*, 22(3):914–935, 2012.

Arkadi Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004.

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

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

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.

Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. *Conference on Learning Theory*, 2013a.

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

R. Tyrrell Rockafellar. Monotone operators and the proximal point algorithm. *SIAM Journal on Control and Optimization*, 14(5):877–898, 1976.

Ernest K Ryu and Wotao Yin. *Large-Scale Convex Optimization via Monotone Operators*. Cambridge University Press, 2022.

Shoham Sabach and Shimrit Shtern. A first order method for solving convex bilevel optimization problems. *SIAM Journal on Optimization*, 27(2):640–660, 2017.

Abdurakhmon Sadiev, Laurent Condat, and Peter Richtárik. Stochastic Proximal Point Methods for Monotone Inclusions under Expected Similarity. *arXiv:2405.14255*, 2024.

Gesualdo Scutari, Francisco Facchinei, Jong-Shi Pang, and Daniel P. Palomar. Real and complex monotone communication games. *IEEE Transactions on Information Theory*, 60(7):4197–4231, 2014.

Jaewook J. Suh, Jisun Park, and Ernest K. Ryu. Continuous-time analysis of anchor acceleration. *Neural Information Processing Systems*, 2023.

Quoc Tran-Dinh and Yang Luo. Halpern-type accelerated and splitting algorithms for monotone inclusions. *arXiv:2110.08150*, 2021.

Rainer Wittmann. Approximation of fixed points of nonexpansive mappings. *Archiv der Mathematik*, 58(5):486–491, 1992.

TaeHo Yoon and Benjamin Grimmer. A Theory of Composition and Duality of Extremal Optimal Fixed-Point Algorithms. *arXiv:2605.02231*, 2026.

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

TaeHo Yoon and Ernest K. Ryu. Accelerated algorithms for smooth convex-concave minimax problems with $\mathcal{O}(1/k^2)$ rate on squared gradient norm. *International Conference on Machine Learning*, 2021.

TaeHo Yoon and Ernest K. Ryu. Accelerated minimax algorithms flock together. *SIAM Journal on Optimization*, 35(1):180–209, 2025.

TaeHo Yoon, Jaeyeon Kim, Jaewook J. Suh, and Ernest K. Ryu. Optimal acceleration for minimax and fixed-point problems is not unique. *International Conference on Machine Learning*, 2024.

A CONVERGENCE OF S-DUAL-OHM WITH GENERAL RANDOM MAPS

We prove Theorem 3.2 and 3.3. With $\rho_k = 0$, Theorem 3.2 immediately recovers Theorem 2.2. Throughout the section, $\mathbb{T}$ is nonexpansive and $\mathbb{T}(x_\star) = x_\star$.

A.1 IDENTITY FOR ARBITRARY EVALUATION MAPS

The following identity holds for S-DUAL-OHM run with arbitrary maps $\mathbb{T}_0, \dots, \mathbb{T}_{N-2}$.

Lemma A.1 (Yoon & Loizou (2026), Lemma 4.2)**.** Set $a_{N,k} = \frac{1}{(N-k)(N-k-1)}$ and

$$Q_k = \frac{1}{2}\left(\|x_k - x_{N-1}\|^2 - \|\mathbb{T}_k(x_k) - \mathbb{T}(x_{N-1})\|^2\right).$$

Then we have

$$0 = \frac{N-1}{4}\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2 + \frac{1}{2}\langle x_{N-1} - \mathbb{T}(x_{N-1}), x_{N-1} - x_0\rangle + \frac{N}{2}\sum_{k=0}^{N-2} a_{N,k} Q_k. \tag{18}$$

Proof. Define the potential

$$\begin{aligned} U_k = &-\frac{N-k-1}{N-k}\|\mathbb{T}_{k-1}(x_{k-1}) - x_k + x_{N-1} - \mathbb{T}(x_{N-1})\|^2 \\ &+ \frac{2}{N-k}\langle \mathbb{T}_{k-1}(x_{k-1}) - x_k + x_{N-1} - \mathbb{T}(x_{N-1}), x_k - x_{N-1}\rangle. \end{aligned}$$

One can directly verify

$$U_k - U_{k+1} = 2a_{N,k}Q_k \tag{19}$$

by substituting the update rule S-DUAL-OHM and simplifying. This has been done in an equivalent form using $\mathbb{T} = \mathbb{I} - \alpha\mathbb{F}$ for cocoercive operator $\mathbb{F}$, in Yoon & Loizou (2026, Lemma 4.2). Because we have $U_{N-1} = 0$, using $\mathbb{T}_{-1}(x_{-1}) = x_0$, we have

$$U_0 = -\frac{N-1}{N}\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2 + \frac{2}{N}\langle x_{N-1} - \mathbb{T}(x_{N-1}), x_0 - x_{N-1}\rangle.$$

Thus, summing (19) for $k = 0, \dots, N-2$ and multiplying by $N/4$ proves the identity (18). □

Observe that

$$\begin{aligned} &\langle x_{N-1} - \mathbb{T}(x_{N-1}), x_{N-1} - x_\star\rangle - \frac{1}{2}\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2 \\ &= \frac{1}{2}\left(\|x_{N-1} - x_\star\|^2 - \|\mathbb{T}(x_{N-1}) - x_\star\|^2\right) \geq 0 \end{aligned}$$

Using the above inequality and taking total expectation in (18) gives

$$\begin{aligned} 0 &\geq \mathbb{E}\left[\frac{N}{4}\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2 + \frac{1}{2}\langle x_{N-1} - \mathbb{T}(x_{N-1}), x_\star - x_0\rangle\right] + \frac{N}{2}\sum_{k=0}^{N-2} a_{N,k}\mathbb{E}[Q_k] \\ &\geq \frac{N}{8}\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] - \frac{1}{2N}\|x_0 - x_\star\|^2 + \frac{N}{2}\sum_{k=0}^{N-2} a_{N,k}\mathbb{E}[Q_k] \end{aligned}$$

and rearranging, we get

$$\mathbb{E}\left[\|x_{N-1} - \mathbb{T}(x_{N-1})\|^2\right] \leq \frac{4D^2}{N^2} - 4\sum_{k=0}^{N-2} a_{N,k}\mathbb{E}[Q_k]. \tag{20}$$

A.2 Proof of Theorem 3.2

Unrolling the S-DUAL-OHM recursion gives

$$x_j = \frac{1}{N}x_0 + \sum_{t=0}^{j-2} a_{N,t}\mathbb{T}_t(x_t) + \frac{N-j}{N-j+1}\mathbb{T}_{j-1}(x_{j-1}) \tag{21}$$

for $j = 1, \dots, N-1$; see (Yoon & Loizou, 2026, Lemma 4.4). Note that this is a convex combination. Define the common moment bound in advance: with $R = \sum_{t=0}^{N-2} \rho_t$, set

$$M^2 = e^R \left(D^2 + \sum_{t=0}^{N-2} \sigma_t^2 \right).$$

Taking $v = x_\star$ in (11) gives

$$\mathbb{E}\left[\|\mathbb{T}_k(u) - x_\star\|^2 \right] \le (1+\rho_k)\|u - x_\star\|^2 + \sigma_k^2$$

for $k = 0, 1, \dots, N-2$. Therefore, using induction and the fact that x_j is a convex combination of x_0 and $\mathbb{T}_t(x_t)$ for $t = 0, \dots, j-1$, we obtain

$$\mathbb{E}\left[\|x_j - x_\star\|^2 \right] \le \prod_{t=0}^{j-1}(1+\rho_t) \left(D^2 + \sum_{t=0}^{j-1} \sigma_t^2 \right) \le M^2. \tag{22}$$

Now, write $\mathbb{S}_k = \mathbb{E}[\mathbb{T}_k]$ for the deterministic mean map. By setting $u = v$ in (11) and using the moment bound (22), we see that

$$\nu_k^2 := \mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\|^2 \right] \le \mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{T}(x_k)\|^2 \right] \le \sigma_k^2 + 2\rho_k M^2 \tag{23}$$

where the first inequality follows from $\mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{T}(x_k)\|^2 \right] = \mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\|^2 \right] + \mathbb{E}\left[\|\mathbb{S}_k(x_k) - \mathbb{T}(x_k)\|^2 \right]$.

Now suppose that we replace only $\mathbb{T}_k$ by an independent copy $\mathbb{T}'_k$, and keep all the other maps the same. Denote the resulting iterates by $x_j^{(k)}$. The two trajectories agree through time k, and the copied trajectory has the same distribution as the original one, so

$$x_j^{(k)} = x_j \quad (0 \le j \le k), \qquad \mathbb{E}\left[\left\| x_t^{(k)} - x_\star \right\|^2 \right] \le M^2 \quad (0 \le t \le N-1).$$

Conditional on x_k, the evaluations $\mathbb{T}_k(x_k)$ and $\mathbb{T}'_k(x_k)$ are independent and have the same mean $\mathbb{S}_k(x_k)$. Therefore,

$$\mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{T}'_k(x_k)\|^2 \right] = \mathbb{E}\left[\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\|^2 \right] + \mathbb{E}\left[\|\mathbb{T}'_k(x_k) - \mathbb{S}_k(x_k)\|^2 \right] = 2\nu_k^2.$$

For $t > k$, the pair $(x_t, x_t^{(k)})$ is independent of the $\mathbb{T}_t$, which is shared between the two trajectories. Thus we can use (12) and apply (22) to obtain

$$\begin{aligned}
\mathbb{E}\left[\left\| \mathbb{T}_t(x_t) - \mathbb{T}_t(x_t^{(k)}) \right\|^2 \right] &= \mathbb{E}\left[\mathbb{E}\left[\left\| \mathbb{T}_t(x_t) - \mathbb{T}_t(x_t^{(k)}) \right\|^2 \,\middle|\, x_t, x_t^{(k)} \right] \right] \\
&\le \mathbb{E}\left[\left\| x_t - x_t^{(k)} \right\|^2 \right] + \rho_t \left(\mathbb{E}\left[\|x_t - x_\star\|^2 \right] + \mathbb{E}\left[\left\| x_t^{(k)} - x_\star \right\|^2 \right] + \sigma_t^2 \right) \\
&\le \mathbb{E}\left[\left\| x_t - x_t^{(k)} \right\|^2 \right] + \rho_t(2M^2 + \sigma_t^2).
\end{aligned} \tag{24}$$

Again, (21) is a convex combination, and squared norm $\|\cdot\|^2$ is convex, so using (24) we obtain

$$\begin{aligned}
\mathbb{E}\left[\left\| x_j - x_j^{(k)} \right\|^2 \right] \le 2a_{N,k}\nu_k^2 &+ \sum_{t=k+1}^{j-2} a_{N,t} \left(\mathbb{E}\left[\left\| x_t - x_t^{(k)} \right\|^2 \right] + \rho_t(2M^2 + \sigma_t^2) \right) \\
&+ \frac{N-j}{N-j+1} \left(\mathbb{E}\left[\left\| x_{j-1} - x_{j-1}^{(k)} \right\|^2 \right] + \rho_{j-1}(2M^2 + \sigma_{j-1}^2) \right)
\end{aligned} \tag{25}$$

for $j > k+1$. Using this bound, we will show by induction that, for $j = k+1, \dots, N-1$,

$$\begin{aligned}&\mathbb{E}\left[\left\|x_j - x_j^{(k)}\right\|^2\right] \\ &\leq \nu_k^2 \left(1 + a_{N,k}(N-j)(N-j-1)\right) + \frac{1}{2}\sum_{t=k+1}^{j-1} \rho_t(2M^2+\sigma_t^2)\left(1 + a_{N,t}(N-j)(N-j-1)\right).\end{aligned} \tag{26}$$

For $j = k+1$, because x_k and $\mathbb{T}_{k-1}(x_{k-1})$ are shared, we have

$$\begin{aligned}\mathbb{E}\left[\left\|x_{k+1} - x_{k+1}^{(k)}\right\|^2\right] &\leq \frac{N-k-1}{N-k}\mathbb{E}\left[\left\|\mathbb{T}_k(x_k) - \mathbb{T}_k'(x_k)\right\|^2\right] \\ &= \frac{2(N-k-1)}{N-k}\nu_k^2 \\ &= \nu_k^2\left(1 + a_{N,k}(N-k-1)(N-k-2)\right),\end{aligned}$$

which proves (26) at the base case. Now fix $j \geq k+2$ and suppose the bound holds through $j-1$. Substituting it into (25) and interchanging the finite sums gives

$$\begin{aligned}&\mathbb{E}\left[\left\|x_j - x_j^{(k)}\right\|^2\right] \\ &\quad\leq \nu_k^2\Bigg[2a_{N,k} + \sum_{t=k+1}^{j-2} a_{N,t}\left(1 + a_{N,k}(N-t)(N-t-1)\right) \\ &\qquad\qquad + \frac{N-j}{N-j+1}\left(1 + a_{N,k}(N-j+1)(N-j)\right)\Bigg] \\ &\qquad + \frac{1}{2}\sum_{s=k+1}^{j-2}\rho_s(2M^2+\sigma_s^2)\Bigg[2a_{N,s} + \sum_{t=s+1}^{j-2} a_{N,t}\left(1 + a_{N,s}(N-t)(N-t-1)\right) \\ &\qquad\qquad + \frac{N-j}{N-j+1}\left(1 + a_{N,s}(N-j+1)(N-j)\right)\Bigg] \\ &\qquad + \frac{N-j}{N-j+1}\rho_{j-1}(2M^2+\sigma_{j-1}^2).\end{aligned} \tag{27}$$

Note that for every $k \leq s \leq j-2$, the identity $a_{N,t} = 1/(N-t-1) - 1/(N-t)$ gives

$$\begin{aligned}2a_{N,s} + \sum_{t=s+1}^{j-2} a_{N,t} + \frac{N-j}{N-j+1} &= 2a_{N,s} + \frac{1}{N-j+1} - \frac{1}{N-s-1} + \frac{N-j}{N-j+1} \\ &= 1 - \frac{1}{N-s} + a_{N,s},\end{aligned}$$

and

$$\begin{aligned}&\sum_{t=s+1}^{j-2} a_{N,t}a_{N,s}(N-t)(N-t-1) + \frac{N-j}{N-j+1}a_{N,s}(N-j+1)(N-j) \\ &= a_{N,s}\left(\sum_{t=s+1}^{j-2} 1 + (N-j)^2\right) \\ &= a_{N,s}\left(j-s-2+(N-j)^2\right) \\ &= a_{N,s}\left(N-s-2+(N-j)(N-j-1)\right).\end{aligned}$$

Consequently, each square bracket simplifies to

$$2a_{N,s} + \sum_{t=s+1}^{j-2} a_{N,t}\left(1 + a_{N,s}(N-t)(N-t-1)\right) + \frac{N-j}{N-j+1}\left(1 + a_{N,s}(N-j+1)(N-j)\right)$$
$$= 1 - \frac{1}{N-s} + a_{N,s} + a_{N,s}\left(N-s-2+(N-j)(N-j-1)\right)$$
$$= 1 + a_{N,s}(N-j)(N-j-1)$$

and the coefficient of the last term also satisfies

$$\frac{N-j}{N-j+1} = \frac{1}{2}\left(1 + a_{N,j-1}(N-j)(N-j-1)\right).$$

Substituting these identities into (27) proves (26), and completes the induction.

In particular, setting $j = N-1$ in (26) and using $\sigma_t^2 \le M^2$ and $\sum_{t=k+1}^{N-2} \rho_t \le R$ gives

$$\mathbb{E}\left[\left\|x_{N-1} - x_{N-1}^{(k)}\right\|^2\right] \le \nu_k^2 + \frac{1}{2}\sum_{t=k+1}^{N-2} \rho_t(2M^2 + \sigma_t^2) \le \nu_k^2 + \frac{3}{2}RM^2. \tag{28}$$

Now, conditional on the maps before iteration k, the error $\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)$ has mean zero and is independent of $x_{N-1}^{(k)}$. Thus, Cauchy–Schwarz inequality, nonexpansivity of $\mathbb{T}$, (23) and (28) imply

$$\left|\mathbb{E}\left[\langle \mathbb{T}(x_{N-1}), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\rangle\right]\right| = \left|\mathbb{E}\left[\left\langle \mathbb{T}(x_{N-1}) - \mathbb{T}(x_{N-1}^{(k)}), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\rangle\right]\right|$$
$$\le \sqrt{\nu_k^2\left(\nu_k^2 + \frac{3}{2}RM^2\right)} \le \nu_k^2 + \frac{3}{4}RM^2. \tag{29}$$

Now for any $u, v \in \mathbb{R}^d$, the identity $\mathbb{E}\left[\mathbb{T}_k(u) - \mathbb{S}_k(u)\right] = 0$ gives the variance decomposition

$$\mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{T}(v)\right\|^2\right] = \left\|\mathbb{S}_k(u) - \mathbb{T}(v)\right\|^2 + \mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{S}_k(u)\right\|^2\right].$$

Subtracting the last term from (11) therefore yields

$$\left\|\mathbb{S}_k(u) - \mathbb{T}(v)\right\|^2 \le \left\|u - v\right\|^2 + \rho_k\left(\left\|u - x_\star\right\|^2 + \left\|v - x_\star\right\|^2\right) + \sigma_k^2 - \mathbb{E}\left[\left\|\mathbb{T}_k(u) - \mathbb{S}_k(u)\right\|^2\right]. \tag{30}$$

Here the expectation is taken over $\mathbb{T}_k$ with u fixed, so both sides are deterministic functions of (u, v), and in particular (30) holds pathwise with $(u, v) = (x_k, x_{N-1})$. Thus, almost surely,

$$\left\|\mathbb{S}_k(x_k) - \mathbb{T}(x_{N-1})\right\|^2$$
$$\le \left\|x_k - x_{N-1}\right\|^2 + \sigma_k^2 + \rho_k\left(\left\|x_k - x_\star\right\|^2 + \left\|x_{N-1} - x_\star\right\|^2\right) - \mathbb{E}\left[\left\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\|^2 \,\middle|\, x_k\right]. \tag{31}$$

Now we manipulate the quantities Q_k appearing in (20) as follows:

$$Q_k = \frac{1}{2}\left(\left\|x_k - x_{N-1}\right\|^2 - \left\|\mathbb{T}_k(x_k) - \mathbb{T}(x_{N-1})\right\|^2\right)$$
$$= \frac{1}{2}\left(\left\|x_k - x_{N-1}\right\|^2 - \left\|\mathbb{S}_k(x_k) - \mathbb{T}(x_{N-1})\right\|^2\right) - \frac{1}{2}\left\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\|^2$$
$$- \left\langle \mathbb{S}_k(x_k), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\rangle + \left\langle \mathbb{T}(x_{N-1}), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\rangle.$$

By independence of x_k and $\mathbb{T}_k$, we have $\mathbb{E}\left[\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k) \,|\, x_k\right] = 0$, which implies

$$\mathbb{E}\left[\langle \mathbb{S}_k(x_k), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\rangle\right] = \mathbb{E}\left[\langle \mathbb{S}_k(x_k), \mathbb{E}\left[\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k) \,|\, x_k\right]\rangle\right] = 0$$
$$\mathbb{E}\left[\mathbb{E}\left[\left\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\|^2 \,\middle|\, x_k\right]\right] = \mathbb{E}\left[\left\|\mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\right\|^2\right] = \nu_k^2.$$

Taking expectations in the expansion of Q_k and applying (31) now shows explicitly how the variance terms cancel:

$$\mathbb{E}\left[Q_k\right] \ge -\frac{\sigma_k^2}{2} - \frac{\rho_k}{2}\left(\mathbb{E}\left[\left\|x_k - x_\star\right\|^2\right] + \mathbb{E}\left[\left\|x_{N-1} - x_\star\right\|^2\right]\right) + \frac{\nu_k^2}{2} - \frac{\nu_k^2}{2}$$
$$+ \mathbb{E}\left[\langle \mathbb{T}(x_{N-1}), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\rangle\right]$$
$$\ge -\frac{\sigma_k^2}{2} - \rho_k M^2 + \mathbb{E}\left[\langle \mathbb{T}(x_{N-1}), \mathbb{T}_k(x_k) - \mathbb{S}_k(x_k)\rangle\right],$$

where the last inequality uses (22). Finally, (29) and (23) yield

$$\begin{aligned}\mathbb{E}\left[Q_k\right] &\geq -\frac{\sigma_k^2}{2} - \rho_k M^2 - \nu_k^2 - \frac{3}{4}RM^2 \\ &\geq -\frac{\sigma_k^2}{2} - \rho_k M^2 - \left(\sigma_k^2 + 2\rho_k M^2\right) - \frac{3}{4}RM^2 \\ &= -\frac{3}{2}\sigma_k^2 - 3\rho_k M^2 - \frac{3}{4}RM^2.\end{aligned}$$

Using this estimate in (20) gives

$$\begin{aligned}\mathbb{E}\left[\left\|x_{N-1} - \mathbb{T}(x_{N-1})\right\|^2\right] \leq \frac{4D^2}{N^2} + 6\sum_{k=0}^{N-2} a_{N,k}\sigma_k^2 \\ + 12M^2\sum_{k=0}^{N-2} a_{N,k}\rho_k + 3RM^2\sum_{k=0}^{N-2} a_{N,k}.\end{aligned}$$

Since $a_{N,k} \leq 1$ and $\sum_k a_{N,k} = 1 - 1/N < 1$, the last two terms are bounded by $15RM^2$, proving Theorem 3.2.

A.3 Proof of Corollary 3.3

For this result, we evaluate one more random map $\mathbb{T}_{N-1}$ because $y_N = \mathbb{T}_{N-1}(x_{N-1})$ is the output of the algorithm. Therefore, we let $R = \sum_{k=0}^{N-1}\rho_k$ and $M^2 = e^R\left(D^2 + \sum_{k=0}^{N-1}\sigma_k^2\right)$. The same moment bound (22) still applies to x_{N-1}. Now by (11) we have

$$\mathbb{E}\left[\left\|\mathbb{T}_{N-1}(x_{N-1}) - \mathbb{T}(x_{N-1})\right\|^2\right] \leq 2\rho_{N-1}M^2 + \sigma_{N-1}^2.$$

Observe that because $\mathbb{T} = \mathbb{J}_{\alpha\mathbb{F}}$, we have $x_{N-1} - \mathbb{T}(x_{N-1}) = \alpha\mathbb{F}(\mathbb{T}(x_{N-1}))$, so

$$\begin{aligned}\left\|\mathbb{F}(y_N)\right\|^2 &\leq 2\left\|\mathbb{F}(y_N) - \mathbb{F}(\mathbb{T}(x_{N-1}))\right\|^2 + 2\left\|\mathbb{F}(\mathbb{T}(x_{N-1}))\right\|^2 \\ &\leq 2L^2\left\|y_N - \mathbb{T}(x_{N-1})\right\|^2 + \frac{2}{\alpha^2}\left\|x_{N-1} - \mathbb{T}(x_{N-1})\right\|^2 \\ &= 2L^2\left\|\mathbb{T}_{N-1}(x_{N-1}) - \mathbb{T}(x_{N-1})\right\|^2 + \frac{2}{\alpha^2}\left\|x_{N-1} - \mathbb{T}(x_{N-1})\right\|^2.\end{aligned}$$

Taking expectation and applying the bounds from Theorem 3.2, we obtain

$$\begin{aligned}\mathbb{E}\left[\left\|\mathbb{F}(y_N)\right\|^2\right] &\leq \frac{2}{\alpha^2}\mathbb{E}\left[\left\|x_{N-1} - \mathbb{T}(x_{N-1})\right\|^2\right] + 2L^2\mathbb{E}\left[\left\|\mathbb{T}_{N-1}(x_{N-1}) - \mathbb{T}(x_{N-1})\right\|^2\right] \\ &\leq \frac{8D^2}{\alpha^2 N^2} + \frac{12}{\alpha^2}\sum_{k=0}^{N-2} a_{N,k}\sigma_k^2 + 2L^2\sigma_{N-1}^2 + M^2\left(\frac{30}{\alpha^2}\sum_{k=0}^{N-2}\rho_k + 4L^2\rho_{N-1}\right).\end{aligned}$$

Because $\alpha L \leq 1$ and $a_{N,k} \leq \frac{2}{(N-k)^2}$, the last line is upper-bounded by (13), completing the proof.

B Analysis of approximate resolvents for monotone problems

The analysis below verifies Condition 3.1 for the finite maps computed by Algorithm 1. Exact sampled resolvents appear only in the proof of these map properties.

B.1 Basic handy lemmas

Lemma B.1. Let $\mathbb{G}\colon \mathbb{R}^d \to \mathbb{R}^d$ be single-valued and maximally monotone. For any $x \in \mathbb{R}^d$, we have $\|x - \mathbb{J}_{\mathbb{G}}(x)\| \le \|\mathbb{G}(x)\|$.

Proof. Note that $x = \mathbb{J}_{\mathbb{G}}(x + \mathbb{G}(x))$, so

$$\|x - \mathbb{J}_{\mathbb{G}}(x)\| = \|\mathbb{J}_{\mathbb{G}}(x + \mathbb{G}(x)) - \mathbb{J}_{\mathbb{G}}(x)\| \le \|(x + \mathbb{G}(x)) - x\| = \|\mathbb{G}(x)\|.$$

□

Lemma B.2. Under Assumptions 1.1 and 3.4, for $\beta > 0$ and any deterministic $u, v \in \mathbb{R}^d$,

$$\mathbb{E}\left[\left\|\mathbb{J}_{\beta\mathbb{F}(\cdot;\xi)}(u) - \mathbb{J}_{\beta\mathbb{F}}(v)\right\|^2\right] \le \|u - v\|^2 + \beta^2\sigma^2. \tag{32}$$

Proof. Let $y = \mathbb{J}_{\beta\mathbb{F}}(v)$ and $\delta_\xi = \mathbb{F}(y;\xi) - \mathbb{F}(y)$. Then we have $v = y + \beta\mathbb{F}(y) = y + \beta\mathbb{F}(y;\xi) - \beta\delta_\xi$. Therefore,

$$y + \beta\mathbb{F}(y;\xi) = v + \beta\delta_\xi \iff \mathbb{J}_{\beta\mathbb{F}}(v) = y = \mathbb{J}_{\beta\mathbb{F}(\cdot;\xi)}(v + \beta\delta_\xi).$$

This implies, for any $u, v \in \mathbb{R}^d$,

$$\begin{aligned}
\mathbb{E}\left[\left\|\mathbb{J}_{\beta\mathbb{F}(\cdot;\xi)(u)} - \mathbb{J}_{\beta\mathbb{F}}(v)\right\|^2\right] &= \mathbb{E}\left[\left\|\mathbb{J}_{\beta\mathbb{F}(\cdot;\xi)(u)} - \mathbb{J}_{\beta\mathbb{F}(\cdot;\xi)}(v + \beta\delta_\xi)\right\|^2\right] \\
&\le \mathbb{E}\left[\left\|u - v - \beta\delta_\xi\right\|^2\right] \\
&\le \|u - v\|^2 + \beta^2\sigma^2
\end{aligned}$$

where the last line uses $\mathbb{E}[\delta_\xi] = 0$ and $\mathbb{E}\left[\|\delta_\xi\|^2\right] \le \sigma^2$.

□

B.2 Proof of Lemma 3.5

Let us denote the ideal iterates from (14) by $z_j(u)$, so that $z_{B_k}(u) = \widehat{\mathbb{T}}_k(u)$. Because $\mathbb{T} = \mathbb{J}_{\alpha\mathbb{F}}$, for any $v \in \mathbb{R}^d$, we have $v = \mathbb{T}(v) + \alpha\mathbb{F}(\mathbb{T}(v))$. Multiplying γ_j throughout and rearranging, we obtain

$$\mathbb{T}(v) + \alpha\gamma_j\mathbb{F}(\mathbb{T}(v)) = (1-\gamma_j)\mathbb{T}(v) + \gamma_j v \iff \mathbb{T}(v) = \mathbb{J}_{\alpha\gamma_j\mathbb{F}}((1-\gamma_j)\mathbb{T}(v) + \gamma_j v).$$

Therefore, applying (32) with $\beta = \alpha\gamma_j$, we obtain

$$\begin{aligned}
&\mathbb{E}\left[\|z_{j+1}(u) - \mathbb{T}(v)\|^2\right] \\
&= \mathbb{E}\left[\left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}\left((1-\gamma_j)z_j(u) + \gamma_j u\right) - \mathbb{J}_{\alpha\gamma_j\mathbb{F}}\left((1-\gamma_j)\mathbb{T}(v) + \gamma_j v\right)\right\|^2\right] \\
&\le \|(1-\gamma_j)z_j(u) + \gamma_j u - ((1-\gamma_j)\mathbb{T}(v) + \gamma_j v)\|^2 + \alpha^2\gamma_j^2\sigma^2 \\
&= \|(1-\gamma_j)(z_j(u) - \mathbb{T}(v)) + \gamma_j(u - v)\|^2 + \alpha^2\gamma_j^2\sigma^2 \\
&\le (1-\gamma_j)\|z_j(u) - \mathbb{T}(v)\|^2 + \gamma_j\|u - v\|^2 + \alpha^2\gamma_j^2\sigma^2
\end{aligned}$$

where the expectation is conditioned on the randomness revealed before ξ_j, and the last line uses convexity of $\|\cdot\|^2$. Now subtracting $\|u - v\|^2$ from both sides of the above, using $\gamma_j = \frac{2}{j+2}$ and taking total expectation, we obtain

$$\mathbb{E}\left[\|z_{j+1}(u) - \mathbb{T}(v)\|^2\right] - \|u - v\|^2 \le \frac{j}{j+2}\left(\mathbb{E}\left[\|z_j(u) - \mathbb{T}(v)\|^2\right] - \|u - v\|^2\right) + \frac{4\alpha^2\sigma^2}{(j+2)^2}.$$

Multiplying by $(j+1)(j+2)$ and chaining the resulting inequality for $j = 0, \dots, B_k - 1$ we obtain

$$B_k(B_k+1)\left(\mathbb{E}\left[\|z_{B_k}(u) - \mathbb{T}(v)\|^2\right] - \|u-v\|^2\right) \le 4\alpha^2\sigma^2 \sum_{j=0}^{B_k-1} \frac{j+1}{j+2} \le 4\alpha^2\sigma^2 B_k.$$

Dividing both sides by $B_k(B_k+1)$ and using $z_{B_k}(u) = \widehat{\mathbb{T}}(u)$, we obtain (15).

Fix k and $u \in \mathbb{R}^d$. Consider the iterates w_j from Algorithm 1 and the ideal iterates $z_j = z_j(u)$ from (14), obtained by using a same sequence of samples. Below, we drop the u-dependence in z_j and w_j. Recall that $p_j = (1-\gamma_j)w_j + \gamma_j u$ and w_{j+1} is obtained by the following iteration: $q_0 = p_j$,

$$q_{r+1} = p_j - \alpha\gamma_j \mathbb{F}(q_r; \xi_j)$$

for $j = 0, \dots, m_{k,j}$, which is the Picard iteration with respect to the operator $x \mapsto p_j - \alpha\gamma_j\mathbb{F}(x;\xi_j)$, which has the contraction factor $\frac{1}{j+2}$ because

$$\begin{aligned}&\|(p_j - \alpha\gamma_j\mathbb{F}(x;\xi_j)) - (p_j - \alpha\gamma_j\mathbb{F}(y;\xi_j))\| \\ &\quad\le \alpha\gamma_j \|\mathbb{F}(x;\xi_j) - \mathbb{F}(y;\xi_j)\| \le \alpha L\gamma_j \|x-y\| = \frac{1}{j+2}\|x-y\|\end{aligned}$$

for any $x, y \in \mathbb{R}^d$. Therefore,

$$\begin{aligned}\left\|w_{j+1} - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\| &\le \left(\frac{1}{j+2}\right)^{m_{k,j}+1} \left\|p_j - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\| \\ &\le \frac{\rho_k}{16(j+2)} \left\|p_j - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\| \\ &\le \frac{\rho_k\alpha\gamma_j}{16(j+2)} \|\mathbb{F}(p_j;\xi_j)\| \end{aligned} \tag{33}$$

where the second line uses $m_{k,j} = \left\lceil \frac{\log(16/\rho_k)}{\log(j+2)} \right\rceil$ and the last line uses Lemma B.1 with $\mathbb{G} = \alpha\gamma_j\mathbb{F}(\cdot;\xi_j)$. On the other hand, we have

$$\begin{aligned}\left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j) - x_\star\right\| &\le \left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j) - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(x_\star)\right\| + \left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(x_\star) - x_\star\right\| \\ &\le \|p_j - x_\star\| + \alpha\gamma_j \|\mathbb{F}(x_\star;\xi_j)\| \end{aligned} \tag{34}$$

where the last line again uses Lemma B.1 with $\mathbb{G} = \alpha\gamma_j\mathbb{F}(\cdot;\xi_j)$, and

$$\|\mathbb{F}(p_j;\xi_j)\| \le \|\mathbb{F}(p_j;\xi_j) - \mathbb{F}(x_\star;\xi_j)\| + \|\mathbb{F}(x_\star;\xi_j)\| \le L\|p_j - x_\star\| + \|\mathbb{F}(x_\star;\xi_j)\|. \tag{35}$$

Now for any random quantities, consider their L^2-norms $\|X\|_2 := \mathbb{E}\left[\|X\|^2\right]^{1/2}$. Using Minkowski inequality and (33)–(35), we have

$$\begin{aligned}\|w_{j+1} - x_\star\|_2 &\le \left\|w_{j+1} - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\|_2 + \left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j) - x_\star\right\|_2 \\ &\le \frac{\rho_k\alpha\gamma_j}{16(j+2)} \|\mathbb{F}(p_j;\xi_j)\|_2 + \|p_j - x_\star\|_2 + \alpha\gamma_j \|\mathbb{F}(x_\star;\xi_j)\|_2 \\ &\le \frac{\rho_k\alpha\gamma_j}{16(j+2)} \left(L\|p_j - x_\star\|_2 + \|\mathbb{F}(x_\star;\xi_j)\|_2\right) + \|p_j - x_\star\|_2 + \alpha\gamma_j\sigma \\ &\le \left(1 + \frac{\rho_k\gamma_j}{32}\right)\|p_j - x_\star\|_2 + \left(1 + \frac{\rho_k}{16}\right)\alpha\gamma_j\sigma \\ &\le \left(1 + \frac{\rho_k\gamma_j}{32}\right)\left((1-\gamma_j)\|w_j - x_\star\|_2 + \gamma_j\|u - x_\star\|_2\right) + \left(1 + \frac{\rho_k}{16}\right)\alpha\gamma_j\sigma \\ &\le \left(1 - \frac{3\gamma_j}{4}\right)\|w_j - x_\star\|_2 + \frac{5\gamma_j}{4}\|u - x_\star\|_2 + \frac{3\alpha\gamma_j\sigma}{2}\end{aligned}$$

where we use $\frac{\alpha L}{j+2} = \frac{1}{2(j+2)} \le \frac{1}{2}$,

$$\|\mathbb{F}(x_\star;\xi_j)\|_2 = \mathbb{E}\left[\|\mathbb{F}(x_\star;\xi_j) - \mathbb{F}(x_\star)\|^2\right]^{1/2} \le \sigma,$$

and $\rho_k \le 1$. Noting that p_{j+1} is a convex combination of w_{j+1} and u, we can use induction starting from $w_0 = u$ to obtain

$$\max\left\{\|w_j - x_\star\|_2, \|p_j - x_\star\|_2\right\} \le 2\left(\|u - x_\star\| + \alpha\sigma\right). \tag{36}$$

Now

$$\begin{aligned}
&\|w_{j+1} - z_{j+1}\|_2 \\
&\le \left\|w_{j+1} - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\|_2 + \left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j) - z_{j+1}\right\|_2 \\
&\le \left\|w_{j+1} - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}(p_j)\right\|_2 + \left\|\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}\left((1-\gamma_j)w_j + \gamma_j u\right) - \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}\left((1-\gamma_j)z_j + \gamma_j u\right)\right\|_2 \\
&\le \frac{\rho_k\alpha\gamma_j}{16(j+2)}\left\|\mathbb{F}(p_j;\xi_j)\right\|_2 + \left\|(1-\gamma_j)(w_j - z_j)\right\|_2 \\
&\le (1-\gamma_j)\left\|w_j - z_j\right\|_2 + \frac{\rho_k\gamma_j}{16(j+2)}(\|u - x_\star\| + 2\alpha\sigma),
\end{aligned}$$

where the third inequality uses (33), and the last inequality combines (35) with (36). Chaining this result for $j = 0, \ldots, B_k - 1$ and using $w_0 = z_0$ and $\prod_{s=j+1}^{B_k-1}(1-\gamma_s) = \frac{(j+1)(j+2)}{B_k(B_k+1)}$, we obtain

$$\begin{aligned}
\sqrt{\mathbb{E}\left[\left\|\mathbb{T}_k(u) - \widehat{\mathbb{T}}_k(u)\right\|^2\right]} &= \|w_{B_k} - z_{B_k}\|_2 \\
&\le \sum_{j=0}^{B_k-1}\frac{j+1}{j+2}\frac{\rho_k}{8B_k(B_k+1)}(\|u - x_\star\| + 2\alpha\sigma) \\
&\le \frac{\rho_k}{8(B_k+1)}(\|u - x_\star\| + 2\alpha\sigma).
\end{aligned} \tag{37}$$

Now we use Young's inequality with (15) and (37):

$$\begin{aligned}
\mathbb{E}\left[\|\mathbb{T}_k(u) - \mathbb{T}(v)\|^2\right] &\le \left(1 + \frac{\rho_k}{8}\right)\mathbb{E}\left[\left\|\widehat{\mathbb{T}}_k(u) - \mathbb{T}(v)\right\|^2\right] + \left(1 + \frac{8}{\rho_k}\right)\mathbb{E}\left[\left\|\mathbb{T}_k(u) - \widehat{\mathbb{T}}_k(u)\right\|^2\right] \\
&\le \left(1 + \frac{\rho_k}{8}\right)\left(\|u - v\|^2 + \frac{4\alpha^2\sigma^2}{B_k+1}\right) + \frac{\rho_k^2\left(1 + \frac{8}{\rho_k}\right)}{64(B_k+1)^2}(\|u - x_\star\| + 2\alpha\sigma)^2 \\
&\le \left(1 + \frac{\rho_k}{8}\right)\|u - v\|^2 + \frac{\rho_k\left(1 + \frac{\rho_k}{8}\right)}{4(B_k+1)^2}\|u - x_\star\|^2 \\
&\quad + \left(1 + \frac{\rho_k}{8}\right)\left(\frac{4}{B_k+1} + \frac{\rho_k}{(B_k+1)^2}\right)\alpha^2\sigma^2 \\
&\le \|u - v\|^2 + \frac{\rho_k}{4}\left(1 + \frac{\left(1 + \frac{\rho_k}{8}\right)}{(B_k+1)^2}\right)\|u - x_\star\|^2 + \frac{\rho_k}{4}\|v - x_\star\|^2 \\
&\quad + \left(1 + \frac{\rho_k}{8}\right)\frac{4 + \frac{\rho_k}{2}}{B_k+1}\alpha^2\sigma^2 \\
&\le \|u - v\|^2 + \rho_k\left(\|u - x_\star\|^2 + \|v - x_\star\|^2\right) + \frac{6\alpha^2\sigma^2}{B_k+1}
\end{aligned}$$

where the third inequality uses $\frac{\rho_k^2\left(1+\frac{8}{\rho_k}\right)}{64(B_k+1)^2} = \frac{\rho_k\left(1+\frac{\rho_k}{8}\right)}{8(B_k+1)^2}$ and $(\|u - x_\star\| + 2\alpha\sigma)^2 \le 2\|u - x_\star\|^2 + 8\alpha^2\sigma^2$, the fourth inequality uses $\frac{\rho_k}{8}\|u - v\|^2 \le \frac{\rho_k}{4}\|u - x_\star\|^2 + \frac{\rho_k}{4}\|v - x_\star\|^2$, and the last inequality follows from $B_k \ge 1$ and $0 < \rho_k \le 1$, which proves (11) with $\sigma_k^2 = \frac{6\alpha^2\sigma^2}{B_k+1}$.

Now we turn to showing (12). First, since $z_{j+1}(u) = \mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}\left((1-\gamma_j)z_j(u) + \gamma_j u\right)$, we have

$$\|z_{j+1}(u) - z_{j+1}(v)\| \le (1-\gamma_j)\|z_j(u) - z_j(v)\| + \gamma_j\|u - v\|$$

for any $u, v \in \mathbb{R}^d$ by nonexpansivity of $\mathbb{J}_{\alpha\gamma_j\mathbb{F}(\cdot;\xi_j)}$. Induction starting from $z_0(u) = u$ and $z_0(v) = v$ shows that $\widehat{\mathbb{T}}_k$ is pathwise nonexpansive, i.e.,

$$\left\|\widehat{\mathbb{T}}_k(u) - \widehat{\mathbb{T}}_k(v)\right\| = \|z_{B_k}(u) - z_{B_k}(v)\| \le \|u - v\|.$$

We use Young's inequality multiple times with (37):

$$
\begin{aligned}
&\mathbb{E}\left[\|\mathbb{T}_k(u)-\mathbb{T}_k(v)\|^2\right]\\
&\le \left(1+\frac{\rho_k}{8}\right)\mathbb{E}\left[\left\|\widehat{\mathbb{T}}_k(u)-\widehat{\mathbb{T}}_k(v)\right\|^2\right]+\left(1+\frac{8}{\rho_k}\right)\mathbb{E}\left[\left\|\left(\mathbb{T}_k(u)-\widehat{\mathbb{T}}_k(u)\right)-\left(\mathbb{T}_k(v)-\widehat{\mathbb{T}}_k(v)\right)\right\|^2\right]\\
&\le \left(1+\frac{\rho_k}{8}\right)\|u-v\|^2+2\left(1+\frac{8}{\rho_k}\right)\left(\mathbb{E}\left[\left\|\mathbb{T}_k(u)-\widehat{\mathbb{T}}_k(u)\right\|^2\right]+\mathbb{E}\left[\left\|\mathbb{T}_k(v)-\widehat{\mathbb{T}}_k(v)\right\|^2\right]\right)\\
&\le \left(1+\frac{\rho_k}{8}\right)\|u-v\|^2+\frac{2\rho_k^2\left(1+\frac{8}{\rho_k}\right)}{64(B_k+1)^2}\left((\|u-x_\star\|+2\alpha\sigma)^2+(\|v-x_\star\|+2\alpha\sigma)^2\right)\\
&\le \left(1+\frac{\rho_k}{8}\right)\|u-v\|^2+\frac{\rho_k\left(1+\frac{\rho_k}{8}\right)}{2(B_k+1)^2}\left(\|u-x_\star\|^2+\|v-x_\star\|^2+8\alpha^2\sigma^2\right)\\
&\le \|u-v\|^2+\frac{\rho_k}{4}\left(1+\frac{2\left(1+\frac{\rho_k}{8}\right)}{(B_k+1)^2}\right)\left(\|u-x_\star\|^2+\|v-x_\star\|^2\right)+\frac{4\rho_k\left(1+\frac{\rho_k}{8}\right)}{(B_k+1)^2}\alpha^2\sigma^2\\
&\le \|u-v\|^2+\rho_k\left(\|u-x_\star\|^2+\|v-x_\star\|^2\right)+\frac{6\rho_k\alpha^2\sigma^2}{B_k+1}
\end{aligned}
$$

where the last line again uses $B_k \ge 1$ and $0 < \rho_k \le 1$.

Finally, it remains to bound the number of oracle evaluations at iteration k, which is

$$
\sum_{j=0}^{B_k-1}(m_{k,j}+1)=\sum_{j=0}^{B_k-1}\left(1+\left\lceil\frac{\log(16/\rho_k)}{\log(j+2)}\right\rceil\right)\le 2B_k+\log\frac{16}{\rho_k}\sum_{j=0}^{B_k-1}\frac{1}{\log(j+2)}.
$$

We can bound the final summation using $\sum_{j=0}^{B_k-1}\frac{1}{\log(j+2)} \le \frac{3(B_k+1)}{\log(B_k+1)}$, which completes the proof.

B.3 Proof of Corollary 3.6

First fix $N \ge 2$ and $\rho_k = 1/(2N^3)$, and consider arbitrary deterministic integer lengths $B_k \ge 1$. Since $\log\frac{16}{\rho_k} = \log(32N^3) \le 4\log(eN)$, Lemma 3.5 bounds the total oracle cost by

$$
\mathcal{O}\left(\sum_{k=0}^{N-1}B_k+\log(eN)\sum_{k=0}^{N-1}\frac{B_k}{\log(B_k+1)}\right).
$$

For any $b, t \ge 1$,

$$
\frac{b}{\log(b+1)} \le \frac{2b}{t}+\frac{e^{t/2}}{\log 2}.
$$

Indeed, the first term suffices when $b \ge e^{t/2}$, and the second suffices otherwise. Applying this with $b = B_k$ and $t = \log\frac{eN}{N-k}$, and using $\sum_{k=0}^{N-1}(N-k)^{-1/2} \le 2\sqrt{N}$, we get

$$
\sum_{k=0}^{N-1}\frac{B_k}{\log(B_k+1)} \le 2\sum_{k=0}^{N-1}\frac{B_k}{\log(eN/(N-k))}+\frac{2\sqrt{e}}{\log 2}N.
$$

Also, $\log\frac{eN}{N-k} \le \log(eN)$ implies

$$
\sum_{k=0}^{N-1}B_k \le \log(eN)\sum_{k=0}^{N-1}\frac{B_k}{\log(eN/(N-k))}.
$$

Consequently, for any choice of batch-sizes $B_k \ge 1$, the total oracle cost is bounded by

$$
\mathcal{O}\left(\log(eN)\left[N+\sum_{k=0}^{N-1}\frac{B_k}{\log(eN/(N-k))}\right]\right). \tag{38}
$$

We now specialize to B_k in (16) and N in Corollary 3.6. Note that $R = \sum_{k=0}^{N-1} \frac{1}{2N^3} = \frac{1}{2N^2} \le \frac{1}{8}$. We use Corollary 3.3 and Lemma 3.5 with $\alpha = \frac{1}{2L}$, and use $e^{1/8} < \frac{8}{7}$ and

$$R \sum_{k=0}^{N-1} \sigma_k^2 = \frac{1}{2N^2} \sum_{k=0}^{N-1} \sigma_k^2 \le \sum_{k=0}^{N-1} \frac{\sigma_k^2}{2(N-k)^2},$$

to obtain

$$\begin{aligned}\mathbb{E}\left[\|\mathbb{F}(y_N)\|^2\right] &\le \frac{704}{7}\frac{L^2D^2}{N^2} + \frac{1728}{7}\sigma^2 \sum_{k=0}^{N-1} \frac{1}{(N-k)^2(B_k+1)} \\ &\le 256\left(\frac{L^2D^2}{N^2} + \sigma^2 \sum_{k=0}^{N-1} \frac{1}{(N-k)^2(B_k+1)}\right).\end{aligned}$$

Because $x \mapsto 1/(x\sqrt{\log(eN/x)})$ is a decreasing function on $[1, N]$,

$$\begin{aligned}\sum_{k=0}^{N-1} \frac{1}{(N-k)\sqrt{\log(eN/(N-k))}} &\le \frac{1}{\sqrt{\log(eN)}} + \int_1^N \frac{dx}{x\sqrt{\log(eN/x)}} \\ &= \frac{1}{\sqrt{\log(eN)}} + 2\big(\sqrt{\log(eN)} - 1\big) \le 2\sqrt{\log(eN)}.\end{aligned} \tag{39}$$

The batch allocation (16) therefore satisfies

$$\sigma^2 \sum_{k=0}^{N-1} \frac{1}{(N-k)^2(B_k+1)} \le \frac{\epsilon^2}{1024\sqrt{\log(eN)}} \sum_{k=0}^{N-1} \frac{1}{(N-k)\sqrt{\log(eN/(N-k))}} \le \frac{\epsilon^2}{512}. \tag{40}$$

Since $N \ge \frac{32LD}{\epsilon}$, we conclude

$$\mathbb{E}\left[\|\mathbb{F}(y_N)\|^2\right] \le \frac{256L^2D^2}{N^2} + 256\sigma^2 \sum_{k=0}^{N-1} \frac{1}{(N-k)^2(B_k+1)} \le \frac{\epsilon^2}{4} + \frac{\epsilon^2}{2} \le \epsilon^2.$$

To bound the complexity, observe that

$$B_k \le 1 + \frac{1024\sigma^2\sqrt{\log(eN)}}{\epsilon^2}\frac{\sqrt{\log(eN/(N-k))}}{N-k}.$$

Since $\log \frac{eN}{N-k} \ge 1$, we obtain from (39):

$$\begin{aligned}\sum_{k=0}^{N-1} \frac{B_k}{\log(eN/(N-k))} &\le N + \frac{1024\sigma^2\sqrt{\log(eN)}}{\epsilon^2} \sum_{k=0}^{N-1} \frac{1}{(N-k)\sqrt{\log(eN/(N-k))}} \\ &\le N + \frac{2048\sigma^2}{\epsilon^2}\log(eN).\end{aligned}$$

Substituting this into (38) therefore bounds the total number of oracle evaluations by

$$\mathcal{O}\left(N\log(eN) + \frac{\sigma^2}{\epsilon^2}\log^2(eN)\right).$$

Since $N = \mathcal{O}\left(1 + \frac{LD}{\epsilon}\right)$ and $\log(eN) = \Theta\left(\log\left(1 + \frac{LD}{\epsilon}\right)\right)$ in the regime $\frac{LD}{\epsilon} \gg 1$, this proves (17).

C EXPERIMENTAL DETAILS

We further specify the tuning and budget normalization of S-DUAL-OHM used in Section 4.

Experiment 1. For scheduled S-DUAL-OHM, we search the same stepsize grid as its constant-batch counterpart, $\eta \in \{0.005, 0.01, 0.025, 0.05, 0.1, 0.25, 0.5, 0.75, 1\}$, and $N - 1 \in \{32, 47, 69, 101, 148, 217, 317, 465, 682, 1000\}$. For each candidate N, we use

$$B_k = \left\lceil \frac{C}{\sqrt{(N-k)(N-k-1)}} \right\rceil, \quad k = 0, \dots, N-2, \qquad \sum_{k=0}^{N-2} B_k \le Q = 20000,$$

and choose the largest feasible $C > 0$. This preserves the allocation in (8) while adhering the oracle budget constraint. Every candidate pair is evaluated on all three tuning seeds, and the selected values are $\eta = 0.75$ and $N = 1001$. The constant-batch counterpart selects $\eta = 0.025$ and $B = 2$, giving $N = 10001$ and the same total oracle cost.

Experiment 2. For the practical variant of S-DUAL-OHM (scheduled), we set

$$B_k = \max\left\{1, \left\lceil C \frac{\sqrt{\log(eN/(N-k))}}{N-k} \right\rceil \right\}, \quad k = 0, \dots, N-1, \quad (m+1) \sum_{k=0}^{N-1} B_k \le Q = 2{\times}10^5.$$

For each candidate (m, N), we choose the largest feasible C. The constant-batch counterpart uses $B = \lfloor Q/((m+1)N) \rfloor$.

For both versions, we test $m \in \{1, 2, 4, 7\}$ and $N \in \{1200, 2400, 4800, 9600\}$ on a test seed. We then fix $m = 1$ and refine $N \in \{4800, 9600, 16000, 24000, 40000\}$ for the scheduled method and $N \in \{1600, 2400, 2500, 3200, 3333, 4800, 5000\}$ for the constant-batch method, using all three tuning seeds. The selected scheduled run has $N = 24000$, and the selected constant run has $N = 3200$ and $B = 31$.

For the theory-prescribed version of S-DUAL-OHM, we set $\rho_k = 1/(2N^3)$ and use the B_k as above, and choose largest C satisfying $\sum_{k=0}^{N-1} \sum_{j=0}^{B_k-1} (m_{k,j} + 1) \le Q$, with $m_{k,j}$ given by Algorithm 1. The final search uses $N \in \{1400, 1600, \dots, 3200\}$ on the tuning seeds and selects $N = 2400$.