STATE-AWARE ZEROTH-ORDER SEARCH FOR DR-SUBMODULAR MAXIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

We study zeroth-order continuous DR-submodular maximization over polyhedral constraints and introduce *state-aware direct search* (SDS), a new zeroth-order search paradigm in which local polyhedral geometry determines how search directions are obtained. On nondegenerate iterations, SDS directly polls feasible directions generated from the local geometry, while gradient estimation is used only in degenerate states. We develop a unified analysis framework that extends SDS to non-oblivious objectives whose stationarity yields approximation guarantees. For monotone and nonmonotone maximization over general polytopes, and nonmonotone maximization over down-closed polytopes, SDS achieves approximation factors $1 - e^{-1}$, $(1 - h)/4$, and e^{-1}, respectively, with $O(\epsilon^{-2})$ expected iterations in all three settings. Under exact and stochastic value oracles, the corresponding worst-case expected value-oracle complexities are $O(\epsilon^{-3})$ and $O(\epsilon^{-4})$, respectively. This matches the best known iteration dependence for general polytopes and improves the zeroth-order iteration complexity for the e^{-1} guarantee in nonmonotone down-closed maximization from $O(\epsilon^{-3})$ to $O(\epsilon^{-2})$. Numerical experiments demonstrate the practical effectiveness of SDS, with degenerate states rarely encountered in practice.

1 INTRODUCTION

Continuous DR-submodular maximization arises in applications including machine learning, social-network optimization, and resource allocation (Bian et al., 2017b; Staib & Jegelka, 2020; Bian et al., 2019). Many such problems admit polyhedral feasible regions, ranging from down-closed to more general polytopes (Hassani et al., 2017). While the constraint geometry is explicit, the objective may only be accessible through simulation, sampling, or black-box evaluation, making exact gradients expensive or unavailable (Chen et al., 2020b; 2021). This motivates zeroth-order methods that rely only on function values. We study such methods for continuous DR-submodular maximization over general and down-closed polytopes, under both exact and stochastic value-oracle models.

Recent years have seen substantial progress on zeroth-order DR-submodular maximization. For monotone objectives over general convex sets, Chen et al. (2020a) obtained a $1 - e^{-1}$ approximation from function values, while Lian et al. (2024) further developed variance reduced zeroth-order methods with $O(\epsilon^{-2})$ iteration complexity in the finite-sum setting. For nonmonotone objectives, value-only guarantees have been established over both general convex sets (Pedramfar & Aggarwal, 2024) and down-closed convex sets (Pedramfar et al., 2023), with the latter approximation factor recently improved to $0.385 - O(\epsilon)$ by Chen et al. (2026). Despite these advances, the recent zeroth-order methods discussed above share a common design principle: scalar function values are first converted into a vector-valued first-order surrogate, which then determines the update. From an information perspective, this may be more than is needed to make a single successful move.This motivates our question: *can local constraint geometry provide the search directions, without estimating a gradient at every iteration, while retaining global approximation guarantees?*

We address this question with *state-aware direct search* (SDS). Direct search is appealing because it can choose an update from directional function differences, without first estimating a gradient. Recent unpublished work has also studied direct search for monotone continuous DR-submodular maximization under a directional requirement that is stronger than the κ-descent condition used in concave direct search (Royer, 2025). We take a different route. Direct search applied directly

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.

Table 1: Comparison of zeroth-order methods for continuous DR-submodular maximization.

Setting	Reference	Oracle	Approx. factor	Iter.	VO	Stoch. model
Monotone general convex	Lian et al. (2024)	Exact	$1 - e^{-1} - \epsilon^2$	$O(\epsilon^{-2})$	$O(\epsilon^{-2})$	–
	Chen et al. (2020a)	Stoch.	$1 - e^{-1}$	$O(\epsilon^{-3})$	$O(\epsilon^{-5})$	BV
	Ours	Exact	$1 - e^{-1}$	$O(\epsilon^{-2})$	$O(\epsilon^{-3})$	–
	Ours	Stoch.	$1 - e^{-1}$	$O(\epsilon^{-2})$	$O(\epsilon^{-4})$	CRN
Nonmonotone general convex	Pedramfar & Aggarwal (2024)	Exact	$\frac{1-h}{4}$	$O(\epsilon^{-2})$	$O(\epsilon^{-2})$	–
	Pedramfar & Aggarwal (2024)	Stoch.	$\frac{1-h}{4}$	$O(\epsilon^{-4})$	$O(\epsilon^{-4})$	AB
	Ours	Exact	$\frac{1-h}{4}$	$O(\epsilon^{-2})$	$O(\epsilon^{-3})$	–
	Ours	Stoch.	$\frac{1-h}{4}$	$O(\epsilon^{-2})$	$O(\epsilon^{-4})$	CRN
Nonmonotone down-closed convex	Pedramfar et al. (2023)	Exact	e^{-1}	$O(\epsilon^{-3})$	$O(\epsilon^{-3})$	–
	Pedramfar et al. (2023)	Stoch.	e^{-1}	$O(\epsilon^{-3})$	$O(\epsilon^{-5})$	BV
	Chen et al. (2026)	Exact	$0.385 - O(\epsilon)$	$O(\epsilon^{-3})$	$O(\epsilon^{-3})$	–
	Ours	Exact	e^{-1}	$O(\epsilon^{-2})$	$O(\epsilon^{-3})$	–
	Ours	Stoch.	e^{-1}	$O(\epsilon^{-2})$	$O(\epsilon^{-4})$	CRN

Here $h := \min_{\boldsymbol{x} \in \mathcal{P}} \|\boldsymbol{x}\|_\infty = \|\underline{\boldsymbol{x}}\|_\infty$. The complexity columns report only the dependence on ϵ.Prior results are stated under their original feasible-set assumptions, whereas SDS is analyzed over polytopes. BV denotes an unbiased stochastic value oracle with uniformly bounded variance; AB denotes an unbiased stochastic value oracle with almost-surely bounded outputs; CRN denotes our unbiased oracle satisfying equation 3. For our method, "Iter." denotes the expected stationarity hitting time.

to f only targets stationarity of f, which does not generally give the approximation guarantees we seek. We instead search a suitable non-oblivious objective F, whose stationarity provides an approximation guarantee for f (Zhang et al., 2022; 2024; Lu et al., 2026). This allows SDS to retain the standard tangent-cone directional condition and to cover the nonmonotone settings considered here. SDS does not query F directly. Its directional increments are estimated using only values of the original objective f. The remaining challenge is to realize this directional condition with a small feasible polling set. For general polytopes, a standard approach is to positively generate the approximate tangent cone. When the nearly active constraint normals are linearly dependent, this construction may require extreme-ray enumeration and a large number of polling directions (Lewis et al., 2007; Gratton et al., 2019).

This is precisely where state awareness enters. SDS uses a rank condition on the nearly active constraint normals to distinguish two local states. When this condition holds, the state is nondegenerate and an algebraic construction provides feasible polling directions from the approximate tangent cone. SDS tests these directions sequentially and accepts the first sufficient ascent in F, without forming a gradient estimate. When the rank condition fails, the state is degenerate. SDS instead estimates the reduced gradient of F and projects it onto the approximate tangent cone, avoiding extreme-ray enumeration. Thus, gradient estimation is used only in the degenerate branch. For certain structured polytopes, this branch may disappear entirely, reducing SDS to a pure direct-search method.

Although the two states use different zeroth-order information, the same analysis applies across all three problem regimes. We establish approximation and complexity guarantees under both exact and stochastic value-oracle models. Table 1 summarizes these results and compares them with representative zeroth-order methods under their respective oracle assumptions.

Contribution. In this work, we study zeroth-order continuous DR-submodular maximization over polyhedral constraints. Our main contributions are as follows:

- We introduce *state-aware direct search* (SDS), a zeroth-order paradigm that uses local polyhedral geometry to decide how search directions are obtained. SDS performs direct polling in nondegenerate states and invokes gradient estimation only under degeneracy, avoiding costly extreme-ray enumeration. Numerical experiments show improved value-oracle efficiency, while degenerate states are rarely encountered in practice.
- We develop a modular analysis that connects non-oblivious objectives with SDS. The framework extends to other objectives whenever their stationarity provides an approximation certificate and the required zeroth-order information can be estimated reliably.
- We establish guarantees for monotone general, nonmonotone general, and nonmonotone down-closed polytopes. SDS attains approximation factors $1 - e^{-1}$, $(1 - h)/4$, and e^{-1}, respectively, with $O(\epsilon^{-2})$ expected iterations. Its expected value-oracle complexity is $O(r\epsilon^{-3})$ for exact values and $O(r^2\epsilon^{-4})$ for stochastic values. For nonmonotone down-closed maximization, SDS improves the best known zeroth-order iteration complexity for the e^{-1} approximation from $O(\epsilon^{-3})$ to $O(\epsilon^{-2})$.

Organization. Section 2 presents the problem and approximation certificates. Section 3 develops SDS and its iteration analysis. Section 4 develops the constructions and establishes value-oracle complexity bounds. Section 5 presents the experiments, and Section 6 concludes.

2 PRELIMINARY

In this section, we introduce the problem setting, value-oracle models, and the non-oblivious stationarity certificates used in our approximation analysis.

Problem setting. We study continuous DR-submodular maximization over a compact polytope:

$$\max_{\boldsymbol{x}\in\mathcal{P}} f(\boldsymbol{x}), \qquad \mathcal{P} = \{\boldsymbol{x} \in \mathbb{R}^n : A\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{l} \le A_{\mathcal{I}}\boldsymbol{x} \le \boldsymbol{u}\} \subseteq \mathcal{X} = [0,1]^n. \tag{1}$$

Here $A \in \mathbb{R}^{m\times n}$, $A_{\mathcal{I}} \in \mathbb{R}^{m_{\mathcal{I}}\times n}$, $\boldsymbol{b} \in \mathbb{R}^m$, and $\boldsymbol{l}, \boldsymbol{u} \in \mathbb{R}^{m_{\mathcal{I}}}$, with $\boldsymbol{l} < \boldsymbol{u}$ componentwise. We assume $\mathcal{P}$ is nonempty and A has full row rank when $m > 0$. The deterministic objective f is continuously differentiable, nonnegative, and DR-submodular on $\mathcal{X}$: $\boldsymbol{x} \le \boldsymbol{y}$ implies $\nabla f(\boldsymbol{x}) \ge \nabla f(\boldsymbol{y})$ componentwise (Bian et al., 2017a). It is L-Lipschitz and L_1-smooth on a query domain $\mathcal{U}$ containing a closed neighborhood of $\mathcal{X}$, with $L, L_1 > 0$. In the monotone case, f is coordinatewise nondecreasing and $f(\mathbf{0}) = 0$. We consider both monotone and nonmonotone objectives over general compact polytopes. We also consider the nonmonotone case when $\mathcal{P}$ is *down-closed*, meaning that $\boldsymbol{x} \in \mathcal{P}$ and $\mathbf{0} \le \boldsymbol{y} \le \boldsymbol{x}$ imply $\boldsymbol{y} \in \mathcal{P}$, this additional structure yields a stronger approximation certificate. Vector norms are Euclidean and matrix norms are induced Euclidean norms, unless indicated otherwise.

To enforce the equalities exactly, set $r = n - m \ge 1$ and let $W = [\boldsymbol{w}_1, \ldots, \boldsymbol{w}_r] \in \mathbb{R}^{n\times r}$ have orthonormal columns spanning $\mathrm{Null}(A) := \{\boldsymbol{d} \in \mathbb{R}^n : A\boldsymbol{d} = \mathbf{0}\}$. Thus $W^\top W = I_r$, $\|\boldsymbol{w}_\ell\| = 1$, and $\|W\boldsymbol{d}\| = \|\boldsymbol{d}\|$, set $W = I_n$ if $m = 0$. Fix $\bar{\boldsymbol{x}} \in \mathbb{R}^n$ with $A\bar{\boldsymbol{x}} = \boldsymbol{b}$. Any point satisfying the equality constraints can be written uniquely as $\boldsymbol{x} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{x}}$, and the remaining constraints become

$$\widetilde{\mathcal{P}} = \{\tilde{\boldsymbol{x}} \in \mathbb{R}^r : \widetilde{A}\tilde{\boldsymbol{x}} \le \widetilde{\boldsymbol{b}}\}, \qquad \widetilde{A} = \begin{pmatrix} A_{\mathcal{I}}W \\ -A_{\mathcal{I}}W \end{pmatrix}, \quad \widetilde{\boldsymbol{b}} = \begin{pmatrix} \boldsymbol{u} - A_{\mathcal{I}}\bar{\boldsymbol{x}} \\ A_{\mathcal{I}}\bar{\boldsymbol{x}} - \boldsymbol{l} \end{pmatrix}. \tag{2}$$

Write $\tilde{\boldsymbol{a}}_i^\top$ for row i of $\widetilde{A}$, $i = 1, \ldots, 2m_{\mathcal{I}}$.

Value oracle models. We now state the information available to the algorithm. The constraints are explicit, while the objective is accessed only through scalar evaluations of the original function f.We consider exact and stochastic observations. An exact query returns $f(\boldsymbol{x})$. In the stochastic regime, the oracle returns a stochastic observation $\widehat{f}(\boldsymbol{x}, \xi)$ satisfying $\mathbb{E}_\xi[\widehat{f}(\boldsymbol{x}, \xi)] = f(\boldsymbol{x})$, and different query points can be sampled with the same realization ξ. With the sampling error $\bar{f}(\boldsymbol{x}, \xi) := \widehat{f}(\boldsymbol{x}, \xi) - f(\boldsymbol{x})$,we assume that the sampling errors at nearby points are correlated in the sense that

$$\mathbb{E}_\xi|\bar{f}(\boldsymbol{x}, \xi) - \bar{f}(\boldsymbol{y}, \xi)|^2 \le D_2\|\boldsymbol{x} - \boldsymbol{y}\|^2, \qquad \boldsymbol{x}, \boldsymbol{y} \in \mathcal{U}, \tag{3}$$

for some $D_2 \ge 0$, following Rinaldi et al. (2024). Each scalar evaluation counts as one value-oracle call. The exact regime is the zero-noise special case, but it permits deterministic quadrature and therefore a different oracle complexity.

Non-oblivious objectives. Direct search can use directional differences without estimating the full reduced gradient from function values at every iteration. Applied directly to f, however, it only targets stationary points with weak approximation guarantees under general constraints (Hassani et al., 2017; Chen et al., 2023). We therefore search non-oblivious objectives whose stationarity provides stronger guarantees for f. For each problem regime, we select a non-oblivious objective with the corresponding stationarity guarantee (Zhang et al., 2022; 2024; Lu et al., 2026):

$$F(\boldsymbol{x}) = \begin{cases} \displaystyle\int_0^1 \frac{e^{z-1}}{z} f(z\boldsymbol{x})\, dz, & \text{monotone, general,} \\ \displaystyle\int_0^1 \frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big) - f(\underline{\boldsymbol{x}})}{4z(1-z/2)^3}\, dz, & \text{nonmonotone, general,} \\ \displaystyle\int_0^1 \frac{e^{z-1}}{(1-e^{-1})z} \left[f\big(\mathbf{1} - e^{-z\boldsymbol{x}}\big) - f(\mathbf{0})\right]\, dz, & \text{nonmonotone, down-closed.} \end{cases} \tag{4}$$

Here $\underline{\boldsymbol{x}} \in \arg\min_{\boldsymbol{x}\in\mathcal{P}} \|\boldsymbol{x}\|_\infty$ in the second case, and the exponential is componentwise. The integrands are defined at $z = 0$ by continuity. Let L_f and L_g denote Lipschitz and smoothness constants of the selected F, respectively. Their validity for all three constructions is shown in Appendix A.1. For each problem regime, the corresponding auxiliary objective is fixed throughout the run. We use F to denote this selected objective below.

Stationarity certificates and approximation. Depending on the objective and constraint structure, the auxiliary objectives above provide different approximation certificates for the original problem. We use a common constrained stationarity measure for all cases. Let $D = \operatorname{diam}(\mathcal{P})$, $D^* = \max\{D, 1\}$, and $\mathrm{OPT} = \max_{\boldsymbol{x}\in\mathcal{P}} f(\boldsymbol{x})$, and define $\pi(\boldsymbol{x}) := \max_{\substack{\boldsymbol{x}+\boldsymbol{d}\in\mathcal{P} \\ \|\boldsymbol{d}\|\le 1}} \langle \nabla F(\boldsymbol{x}), \boldsymbol{d}\rangle$. Since $(\boldsymbol{y} - \boldsymbol{x})/D^*$ is admissible, for every $\boldsymbol{y} \in \mathcal{P}$, $\max_{\boldsymbol{y}\in\mathcal{P}} \langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x}\rangle \le D^*\pi(\boldsymbol{x})$.

For the auxiliary objective selected for each problem regime, the existing non-oblivious inequalities imply that the common stationarity tolerance $\pi(\boldsymbol{x}) \le \varepsilon := \epsilon/D^*$ yields

$$\begin{aligned} f(\boldsymbol{x}) &\ge (1 - e^{-1})\mathrm{OPT} - \epsilon, && \text{monotone, general,} \\ f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right) &\ge \frac{1 - \|\underline{\boldsymbol{x}}\|_\infty}{4}\mathrm{OPT} - \epsilon, && \text{nonmonotone, general,} \\ f\big(\mathbf{1} - e^{-\boldsymbol{x}}\big) &\ge e^{-1}\mathrm{OPT} - \epsilon, && \text{nonmonotone, down-closed.} \end{aligned} \tag{5}$$

The first two guarantees follow from Zhang et al. (2022; 2024), while the last follows from the construction of Lu et al. (2026). Appendix A.2 gives the corresponding stationarity to approximation derivations.

Our remaining task is therefore algorithmic: reach the same stationarity tolerance using only original-function values, without estimating the full reduced auxiliary gradient at every iteration.

3 State-Aware Direct Search and Iteration Complexity

This section develops the state-aware search mechanism and its iteration complexity. We first show how the local polyhedral geometry determines which zeroth-order information is acquired, and then establish a unified $O(\varepsilon^{-2})$ expected iteration bound under conditional information reliability.

3.1 Why State Awareness?

Direct search is attractive in our value-only setting because an update can be driven by directional differences without first estimating the full reduced auxiliary gradient. This advantage, however, relies on a small feasible polling set that sufficiently captures local ascent. Constructing such a set is not always easy, and its size need not remain small. We first characterize polling directions that preserve feasibility at the current step size. Following Gratton et al. (2019), define:

Definition 3.1 (δ-active set and approximate cones)**.** For $\boldsymbol{x} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{x}} \in \mathcal{P}$ and step size $\delta \ge 0$, define the δ-active index set by

$$I(\boldsymbol{x}, \delta) := \left\{ i : \widetilde{b}_i - \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{x}} \le \delta \|\tilde{\boldsymbol{a}}_i\| \right\}. \tag{6}$$

The corresponding δ-approximate normal cone and δ-approximate tangent cone are

$$N(\boldsymbol{x},\delta) := \mathrm{cone}\{\tilde{\boldsymbol{a}}_i : i \in I(\boldsymbol{x},\delta)\}, \qquad T(\boldsymbol{x},\delta) := \left\{\boldsymbol{v} \in \mathbb{R}^r : \boldsymbol{v}^\top \boldsymbol{u} \le 0 \text{ for all } \boldsymbol{u} \in N(\boldsymbol{x},\delta)\right\}. \tag{7}$$

Here cone denotes nonnegative linear combinations.

If $\boldsymbol{d} \in T(\boldsymbol{x},\delta)$, $\|\boldsymbol{d}\| \le \delta$, then $\boldsymbol{x} + W\boldsymbol{d} \in \mathcal{P}$. Thus, choosing $\mathcal{D}(\boldsymbol{x},\delta) \subset T(\boldsymbol{x},\delta)$ guarantees feasible polling. Feasibility alone is not enough: the polling directions should also capture feasible first-order ascent. Adapting the feasible-descent criterion of Gratton et al. (2019) to maximization, we use the cosine measure relative to the approximate tangent cone.

Definition 3.2 (Cosine measure and κ-feasible ascent set)**.** Let $\boldsymbol{x} \in \mathcal{P}$, $\delta > 0$, and let $\mathcal{D} \subset T(\boldsymbol{x},\delta)$ be a nonempty finite set of nonzero directions. For $\boldsymbol{g} \in \mathbb{R}^r$, we call

$$\mathrm{cm}_{T(\boldsymbol{x},\delta)}(\mathcal{D},\boldsymbol{g}) := \max_{\boldsymbol{d}\in\mathcal{D}} \frac{\boldsymbol{d}^\top \boldsymbol{g}}{\|\boldsymbol{d}\|\,\|P_{T(\boldsymbol{x},\delta)}[\boldsymbol{g}]\|} \tag{8}$$

the cosine measure of $\mathcal{D}$ relative to $T(\boldsymbol{x},\delta)$ and $\boldsymbol{g}$, where $P_{T(\boldsymbol{x},\delta)}$ denotes Euclidean projection onto $T(\boldsymbol{x},\delta)$. When $P_{T(\boldsymbol{x},\delta)}[\boldsymbol{g}] = \boldsymbol{0}$, the measure is defined to be 1. For $\kappa \in (0,1)$, $\mathcal{D}$ is called a κ-feasible ascent set for $\boldsymbol{g}$ if $\mathrm{cm}_{T(\boldsymbol{x},\delta)}(\mathcal{D},\boldsymbol{g}) \ge \kappa$.

In our algorithm, this criterion is applied to the reduced auxiliary gradient $\boldsymbol{g}(\boldsymbol{x}) := W^\top \nabla F(\boldsymbol{x})$. A finite set $\mathcal{V} \subset C$ of nonzero vectors is a *positive generating set* of a cone $C \subseteq \mathbb{R}^r$ if $C = \mathrm{cone}(\mathcal{V})$. We construct such generators using the algebraic procedure of Lewis et al. (2007). For a fixed pair $(\boldsymbol{x},\delta)$, write

$$N(\boldsymbol{x},\delta) = \mathrm{cone}\big([\, E(\boldsymbol{x},\delta) \;\; -E(\boldsymbol{x},\delta)\; M(\boldsymbol{x},\delta)\,]\big), \tag{9}$$

where $E(\boldsymbol{x},\delta)$ collects active normal directions whose opposites also occur, and $M(\boldsymbol{x},\delta)$ contains the remaining directions. Hence every $\boldsymbol{d} \in T(\boldsymbol{x},\delta)$ satisfies $E(\boldsymbol{x},\delta)^\top \boldsymbol{d} = \boldsymbol{0}$ and $M(\boldsymbol{x},\delta)^\top \boldsymbol{d} \le \boldsymbol{0}$.

Let $B(\boldsymbol{x},\delta)$ have orthonormal columns forming a basis of $\mathrm{Null}(E(\boldsymbol{x},\delta)^\top)$, and define

$$Q(\boldsymbol{x},\delta)^\top := M(\boldsymbol{x},\delta)^\top B(\boldsymbol{x},\delta). \tag{10}$$

When $Q(\boldsymbol{x},\delta)^\top$ has full row rank, positive generators of $T(\boldsymbol{x},\delta)$ can be constructed directly from this representation.

Proposition 3.3 (Positive generators in the nondegenerate state)**.** *Let $\boldsymbol{x} \in \mathcal{P}$ and $\delta > 0$, and suppose that $Q(\boldsymbol{x},\delta)^\top$ has full row rank. Let R_Q be a right inverse of $Q(\boldsymbol{x},\delta)^\top$, and let the columns of Z_Q form a basis of $\mathrm{Null}(Q(\boldsymbol{x},\delta)^\top)$. Then the columns of*

$$\big[-B(\boldsymbol{x},\delta)R_Q \quad B(\boldsymbol{x},\delta)Z_Q \quad -B(\boldsymbol{x},\delta)Z_Q\big]$$

positively generate $T(\boldsymbol{x},\delta)$. The number of generated directions is at most $2r$.

We call $(\boldsymbol{x},\delta)$ a *nondegenerate state* if $Q(\boldsymbol{x},\delta)^\top$ has full row rank, and a *degenerate state* otherwise.

For each set I that can occur as $I(\boldsymbol{x},\delta)$ in a nondegenerate state with $T(\boldsymbol{x},\delta) \ne \{\boldsymbol{0}\}$, we make a fixed choice of a normalized positive generating set given by Proposition 3.3. Since $I(\boldsymbol{x},\delta) \subseteq \{1,\dots,2m_{\mathcal{I}}\}$, only finitely many such sets can occur, and hence only finitely many fixed generator sets are used. Each has a positive cosine measure lower bound, so their minimum defines a common constant $\kappa \in (0,1)$. Therefore, in every nondegenerate state with $T(\boldsymbol{x},\delta) \ne \{\boldsymbol{0}\}$, the resulting generator set is a κ-feasible ascent set. If $T(\boldsymbol{x},\delta) = \{\boldsymbol{0}\}$, no polling direction is required. In either case, no full reduced-gradient estimate is needed.

When $Q(\boldsymbol{x},\delta)^\top$ is rank deficient, this construction no longer applies. Positive generators can still be obtained by extreme-ray enumeration, for example by the double-description method (Fukuda & Prodon, 1996), but their number may be combinatorial in the number of active constraints. We therefore do not insist on constructing positive generators in every state. Instead, in a degenerate state, we estimate the reduced auxiliary gradient and poll along its projection onto $T(\boldsymbol{x},\delta)$.

Thus, the local geometry determines the information acquired: directional function differences are sufficient in nondegenerate states, whereas the richer reduced-gradient information is acquired only in degenerate states. In particular, if $Q(\boldsymbol{x},\delta)^\top$ has full row rank for every $\boldsymbol{x} \in \mathcal{P}$ and $0 < \delta \le \delta_{\max}$, the degenerate branch is never invoked and SDS reduces to a pure direct-search method. For a small $\delta_{\max}$, this includes simple reduced polytopes (Ziegler, 1995) with an irredundant facet representation, such as boxes and simplices.

3.2 State-Aware Algorithm and Iteration Complexity

Building on the local geometry above, we now turn the state distinction into a state-aware search rule and establish its expected iteration complexity under probabilistically reliable information.

Algorithm 1 starts from $\boldsymbol{x}_0 \in \mathcal{P}$ with polling radius $0 < \delta_0 \leq \delta_{\max}$. The factors $\theta \in (0,1)$ and $\beta > 1$ contract and expand the radius after unsuccessful and successful polls, respectively. Because F is not queried directly, its directional increments are estimated from values of f. $\varepsilon_f > 0$ prescribes their accuracy scale, and $\gamma > 1$ leaves a margin between estimation error and accepted ascent. For a target approximation error $\epsilon > 0$, set $\varepsilon = \epsilon/D^*$.

Algorithm 1 State-aware direct search

Input: $\boldsymbol{x}_0 \in \mathcal{P}$, $0 < \delta_0 \leq \delta_{\max}$, $\theta \in (0,1)$, $\beta, \gamma > 1$, $\varepsilon_f > 0$, and $\epsilon > 0$; set $\varepsilon = \epsilon/D^*$.
for $k = 0, 1, 2, \ldots$ **do**
  Set $T_k = T(\boldsymbol{x}_k, \delta_k)$ and $Q_k^\top = Q(\boldsymbol{x}_k, \delta_k)^\top$.
  if $Q_k^\top$ has full row rank **then**
    Construct $\mathcal{D}_k \subset T_k$ from the normalized positive generators in Section 3.1.
  else
    Construct a reduced-gradient estimate $\boldsymbol{g}_k$ of $\boldsymbol{g}(\boldsymbol{x}_k)$ and set $\widehat{\boldsymbol{q}}_k = P_{T_k}[\boldsymbol{g}_k]$.
    Set $\mathcal{D}_k = \{\widehat{\boldsymbol{q}}_k/\|\widehat{\boldsymbol{q}}_k\|\}$ if $\|\widehat{\boldsymbol{q}}_k\| > \varepsilon/(4D^*)$, and $\mathcal{D}_k = \emptyset$ otherwise.
  end if
  Poll $\boldsymbol{d} \in \mathcal{D}_k$ sequentially and form $\widetilde{a}_{k,\boldsymbol{d}}$ until either $\widetilde{a}_{k,\boldsymbol{d}} \geq \gamma\varepsilon_f\delta_k^2$ or all directions have been tested.
  if a successful direction $\boldsymbol{d}$ is found **then**
    Set $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \delta_k W\boldsymbol{d}$ and $\delta_{k+1} = \min\{\beta\delta_k, \delta_{\max}\}$.
  else
    Set $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k$ and $\delta_{k+1} = \theta\delta_k$.
  end if
end for

The two states are designed to acquire only as much information as the local geometry requires. In a nondegenerate state, we avoid estimating the full reduced gradient and need not poll the entire κ-feasible ascent set, since a possibly randomized subset is sufficient as long as it captures an ascent direction with prescribed probability, and polling can stop as soon as one direction passes the sufficient-ascent test. In a degenerate state, rather than constructing a potentially large positive generating set, we estimate the reduced gradient and use its projection onto T_k. If norm of this projection exceeds $\varepsilon/(4D^*)$, it provides the polling direction; otherwise, we shrink δ_k. The latter is essential because a small projection onto the approximate tangent cone alone does not certify stationarity at a coarse radius. As δ_k decreases, $T(\boldsymbol{x}_k, \delta_k)$ becomes a tighter local model, and a sufficiently small projected gradient can be converted into the target condition $\pi(\boldsymbol{x}_k) \leq \varepsilon$.

We next state the information requirements used in the analysis. All algorithmic randomness is defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Uppercase symbols denote random quantities and lowercase symbols their realizations. In particular, $\boldsymbol{X}_k$, Δ_k, and $\boldsymbol{G}_k$ have realizations $\boldsymbol{x}_k$, δ_k, and $\boldsymbol{g}_k$, respectively. For the polling set, we use $\mathfrak{D}_k$ for the random set and $\mathcal{D}_k$ for its realization. For a realized direction $\boldsymbol{d} \in \mathcal{D}_k$, define the true auxiliary increment $A_{k,\boldsymbol{d}} := F(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}) - F(\boldsymbol{X}_k)$, with realization $a_{k,\boldsymbol{d}}$. Its estimate and realization are denoted by $\widetilde{A}_{k,\boldsymbol{d}}$ and $\widetilde{a}_{k,\boldsymbol{d}}$, respectively.Let $\mathcal{F}_{k-1}$ contain the history before iteration k, so that $\boldsymbol{X}_k$, Δ_k, and the local geometry are already determined. After the state-specific information has been generated, let $\mathcal{F}_{k-1/2} := \sigma(\mathcal{F}_{k-1}, \boldsymbol{G}_k, \mathfrak{D}_k)$, where $\boldsymbol{G}_k$ is included only in a degenerate state and $\mathfrak{D}_k$ is the possibly randomized candidate polling set. Directional-increment estimates are sampled afterwards, and $\mathcal{F}_k$ contains these samples and the resulting update. The stochastic analysis tracks three random ingredients: the accuracy of the estimated auxiliary increments, the directional quality of the realized polling set, and the accuracy of the fallback reduced-gradient estimate. We characterize them by the events J_k, Z_k, and H_k,

respectively. For the fixed directional-quality constant $\kappa \in (0,1)$ introduced above, define

$$J_k := \left\{ \max_{\boldsymbol{d}\in\mathfrak{D}_k} |\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \le \varepsilon_f \Delta_k^2 \right\},$$
$$Z_k := \left\{ \mathrm{cm}_{T(\boldsymbol{X}_k,\Delta_k)}(\mathfrak{D}_k, W^\top \nabla F(\boldsymbol{X}_k)) \ge \kappa \right\}, \qquad H_k := \left\{ \|\boldsymbol{G}_k - W^\top \nabla F(\boldsymbol{X}_k)\| \le \varepsilon_g \right\}. \tag{11}$$

Here $\varepsilon_g := \frac{1-\kappa}{1+\kappa}\frac{\varepsilon}{4D^*}$ is the required accuracy of the reduced-gradient estimate in a degenerate state, and $J_k := \Omega$ when $\mathfrak{D}_k = \emptyset$. For a nondegenerate state with $T(\boldsymbol{X}_k, \Delta_k) = \{\mathbf{0}\}$, we set $Z_k = \Omega$. No polling direction is required in this case. We require these three sources of random information to be sufficiently reliable in the following conditional sense.

Assumption 3.4 (Conditional information reliability)**.** *For fixed $p_d, p_f \in (0,1]$, almost surely,*

$$\mathbb{P}(Z_k \mid \mathcal{F}_{k-1}) \ge p_d, \qquad \textit{on nondegenerate iterations}, \tag{12}$$
$$\mathbb{P}(H_k \mid \mathcal{F}_{k-1}) \ge p_d, \qquad \textit{on degenerate iterations}, \tag{13}$$
$$\mathbb{P}(J_k \mid \mathcal{F}_{k-1/2}) \ge p_f, \tag{14}$$
$$\mathbb{E}\left[\mathbf{1}_{J_k^c} \max_{\boldsymbol{d}\in\mathfrak{D}_k} |\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \,\middle|\, \mathcal{F}_{k-1/2} \right] \le (1-p_f)\varepsilon_f \Delta_k^2. \tag{15}$$

Under Assumption 3.4, Algorithm 1 admits the following unified iteration-complexity guarantee for all three problem regimes.

Theorem 3.5 (Iteration complexity and approximation)**.** *Let $T_\varepsilon := \inf\{k \ge 0 : \pi(\boldsymbol{X}_k) \le \varepsilon\}$. For any target accuracy $\epsilon > 0$, set $\varepsilon = \epsilon/D^*$. Suppose Assumption 3.4 holds and the parameters satisfy*

$$(1-p_f)(\beta^2 - \theta^2) < p_f(\gamma - 1)(1-\theta^2), \qquad p_d p_f \log\beta + (1 - p_d p_f)\log\theta > 0.$$

Then

$$\mathbb{E}[T_\varepsilon] = O(\varepsilon^{-2}) = O(\epsilon^{-2}).$$

Moreover, the corresponding output associated with $\boldsymbol{X}_{T_\varepsilon}$ satisfies equation 5.

These parameter conditions are jointly feasible: for any fixed $\theta \in (0,1)$ and $\beta, \gamma > 1$, they hold by choosing p_f and p_d close to one. The proof relies on a common small-radius property of the two states: before T_ε, a reliable iteration with $\Delta_k \le \delta_\varepsilon/\theta$ must succeed, where $\delta_\varepsilon = \Theta(\varepsilon)$. A nonnegative potential then controls $\sum_{k<T_\varepsilon} \Delta_k^2$, while the stepsize dynamics are compared with a positive-drift random walk whose renewals occur at radius at least δ_ε. A renewal–reward argument therefore gives $\mathbb{E}[T_\varepsilon] = O(\delta_\varepsilon^{-2}) = O(\varepsilon^{-2})$. The approximation guarantees follow from the stationarity certificates in Section 2. The full proof is given in Appendix B.

4 Constructions and Value-Oracle Complexity

This section realizes the reliability requirements of Assumption 3.4 using only values of the original objective and derives the resulting value-oracle complexity. The construction combines randomized polling with exact quadrature or stochastic Monte Carlo estimation, depending on the oracle model.

Randomized polling. In a nondegenerate state, we apply the subspace-polling construction of Gratton et al. (2019) to the positive generators from Proposition 3.3. Rather than polling the complete generator set, we randomly retain directions from its lineality and pointed components. For any prescribed $p_d \in (0,1)$, the constant $\kappa > 0$ introduced above can be decreased, so that $\mathbb{P}(Z_k \mid \mathcal{F}_{k-1}) \ge p_d$, while $|\mathfrak{D}_k| = O(r)$. We use this possibly smaller value of κ throughout the analysis. For $p_d = 1$, we use the full generator set. Construction details are given in Appendix C.1.

Exact value oracle. With exact access to f, we approximate the integrals defining F by midpoint quadrature. Let $z_j = (j - \frac{1}{2})/N$, $j = 1, \dots, N$. For a polling direction $\boldsymbol{d}$, in the monotone case we use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{1}{N}\sum_{j=1}^N \frac{e^{z_j - 1}}{z_j}[f(z_j(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d})) - f(z_j \boldsymbol{X}_k)]$. For the general nonmonotone objective, we use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{1}{N}\sum_{j=1}^N \frac{f(\underline{\boldsymbol{x}} + \frac{z_j}{2}(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}})) - f(\underline{\boldsymbol{x}} + \frac{z_j}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}))}{4 z_j (1 - z_j/2)^3}$. For the down-closed nonmonotone objective, we use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{1}{N}\sum_{j=1}^N \frac{e^{z_j - 1}}{(1-e^{-1})z_j}[f(\mathbf{1} - e^{-z_j(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d})}) - f(\mathbf{1} - e^{-z_j \boldsymbol{X}_k})]$.

In all three cases, $N = O(\Delta_k^{-1})$ is sufficient for $|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \le \varepsilon_f \Delta_k^2$, so J_k holds deterministically. A complete poll therefore uses $O(|\mathfrak{D}_k|\Delta_k^{-1})$ original-value calls. At a degenerate state, $\boldsymbol{G}_k$ is constructed by centered differences at the same midpoint nodes. For $\ell = 1, \dots, r$, the monotone construction is $(\boldsymbol{G}_k)_\ell = \frac{1}{N}\sum_{j=1}^{N} \frac{e^{z_j}-1}{2\mu z_j}[f(z_j(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)) - f(z_j(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell))]$. For the general nonmonotone objective, use $(\boldsymbol{G}_k)_\ell = \frac{1}{N}\sum_{j=1}^{N} \frac{f(\underline{\boldsymbol{x}}+\frac{z_j}{2}(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell-\underline{\boldsymbol{x}}))-f(\underline{\boldsymbol{x}}+\frac{z_j}{2}(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell-\underline{\boldsymbol{x}}))}{8\mu z_j(1-z_j/2)^3}$. For the down-closed objective, use $(\boldsymbol{G}_k)_\ell = \frac{1}{N}\sum_{j=1}^{N} \frac{e^{z_j}-1}{2(1-e^{-1})\mu z_j}[f(\mathbf{1} - e^{-z_j(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell)}) - f(\mathbf{1} - e^{-z_j(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell)})]$. With $N = O(\epsilon^{-1})$ midpoint nodes and the finite-difference radius chosen as specified in Appendix C, we have $\|\boldsymbol{G}_k - W^\top \nabla F(\boldsymbol{X}_k)\| \le \varepsilon_g$. Hence H_k holds deterministically, and each fallback requires $O(r\varepsilon^{-1})$ exact value-oracle calls.

Stochastic value oracle. For stochastic observations, each Monte Carlo replication uses the same realization ξ_i at the two queried points. For a polling direction $\boldsymbol{d}$, in the monotone case we draw Z_i from density $e^{z-1}/(1-e^{-1})$ on $(0,1)$ and use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{1-e^{-1}}{m}\sum_{i=1}^{m}[\widehat{f}(Z_i(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}), \xi_i) - \widehat{f}(Z_i\boldsymbol{X}_k, \xi_i)]/Z_i$. For the general nonmonotone objective, drawing Z_i from density $1/[3(1-z/2)^3]$, we use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{3}{4m}\sum_{i=1}^{m}[\widehat{f}(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}}), \xi_i) - \widehat{f}(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}), \xi_i)]/Z_i$. For the down-closed nonmonotone objective, with Z_i drawn from density $e^{z-1}/(1-e^{-1})$, we use $\widetilde{A}_{k,\boldsymbol{d}} = \frac{1}{m}\sum_{i=1}^{m}[\widehat{f}(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k+\Delta_k W\boldsymbol{d})}, \xi_i) - \widehat{f}(\mathbf{1} - e^{-Z_i\boldsymbol{X}_k}, \xi_i)]/Z_i$. For all three cases, taking $m = O(|\mathfrak{D}_k|\Delta_k^{-2})$ independent replications per polling direction gives $\mathbb{P}(J_k \mid \mathcal{F}_{k-1/2}) \ge p_f$ and the conditional tail bound in Assumption 3.4. Hence a complete stochastic poll uses $O(|\mathfrak{D}_k|^2\Delta_k^{-2})$ original-value calls.At a degenerate state, each coordinate of $\boldsymbol{G}_k$ is estimated by paired centered differences. For $\ell = 1, \dots, r$, the monotone construction is $(\boldsymbol{G}_k)_\ell = \frac{1-e^{-1}}{m}\sum_{i=1}^{m}[\widehat{f}(Z_i(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell), \xi_{i,\ell}) - \widehat{f}(Z_i(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell), \xi_{i,\ell})]/(2\mu Z_i)$, where Z_i has density $e^{z-1}/(1-e^{-1})$. For the general nonmonotone objective, with Z_i drawn from density $1/[3(1-z/2)^3]$, we use $(\boldsymbol{G}_k)_\ell = \frac{3}{4m}\sum_{i=1}^{m}[\widehat{f}(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell}) - \widehat{f}(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell})]/(2\mu Z_i)$. For the down-closed nonmonotone objective, with Z_i drawn from density $e^{z-1}/(1-e^{-1})$, we use $(\boldsymbol{G}_k)_\ell = \frac{1}{m}\sum_{i=1}^{m}[\widehat{f}(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell)}, \xi_{i,\ell}) - \widehat{f}(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell)}, \xi_{i,\ell})]/(2\mu Z_i)$. Using $m = O(r\varepsilon^{-2})$ independent replications per coordinate, with the finite-difference radius chosen as specified in Appendix C, gives $\mathbb{P}(H_k \mid \mathcal{F}_{k-1}) \ge p_d$. Thus each stochastic fallback requires $O(r^2\varepsilon^{-2})$ stochastic value-oracle calls.

Value-oracle complexity. The constructions above verify Assumption 3.4 and specify the value-oracle cost of each iteration. Combining these costs with Theorem 3.5 yields the following total expected value-oracle complexity.

Theorem 4.1 (Worst-case value-oracle complexity)**.** *Suppose the conditions of Theorem 3.5 hold. For the exact value oracle, take $p_f = 1$ and suppose $p_d\beta^{-1} + (1-p_d)\theta^{-1} < 1$. Then the worst-case expected number of value-oracle calls before T_ε is $O(r\varepsilon^{-3})$. For the stochastic value oracle, suppose $p_d p_f \beta^{-2} + (1 - p_d p_f)\theta^{-2} < 1$. Then the worst-case expected number of value-oracle calls before T_ε is $O(r^2\varepsilon^{-4})$.*

The proof combines the per-iteration costs with inverse-stepsize summability and is deferred to Appendix D.These are worst-case bounds: the reduced-gradient fallback is charged only on degenerate iterations, whereas nondegenerate iterations use only directional increment estimates. Thus the actual oracle cost can be substantially smaller when degenerate states are rare.

5 Numerical Experiments

We evaluate SDS on three DR-submodular benchmarks: a quadratic objective, a finite sum of quadratic objectives, and a coverage objective. For each benchmark, we consider the three problem regimes studied above under both exact and stochastic value oracles. For the finite-sum problem, we further vary the dimension over $n \in \{50, 100, 300\}$. We compare objective value against iterations and cumulative value-oracle calls, and record how often SDS enters the degenerate state. Due to space constraints, the main text presents the $n = 100$ finite-sum results, complete experiments and implementation details are deferred to Appendix E.

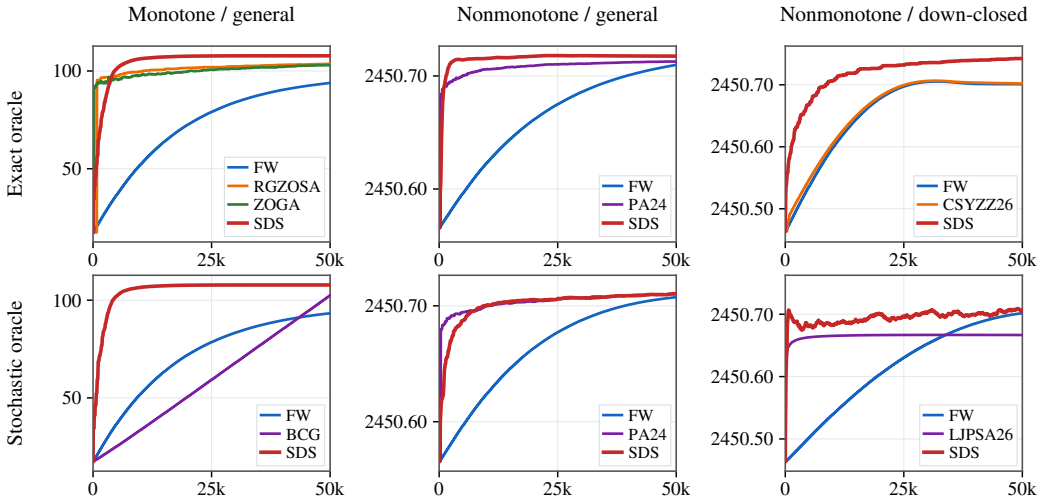


Figure 1: Average objective value versus value-oracle calls for the $n = 100$ finite-sum quadratic problem.

We compare SDS with zeroth-order baselines matched to each problem regime and oracle model. FW denotes the corresponding zeroth-order Frank–Wolfe variant of Pedramfar et al. (2023) and is included in all six settings. For monotone problems, we further compare with RG-ZOSA and ZO-GA (Lian et al., 2024) under the exact oracle, and with BCG (Chen et al., 2020a) under the stochastic oracle. For methods without a standard algorithm name, we use PA24, CSYZZ26, and LJPSA26 to denote Pedramfar & Aggarwal (2024), Chen et al. (2026), and Lu et al. (2026), respectively. PA24 is used for nonmonotone maximization over general polytopes under both oracle models, while CSYZZ26 and LJPSA26 are used for nonmonotone maximization over down-closed polytopes under the exact and stochastic oracles, respectively. All baselines access the objective only through values of the original function f. In the stochastic experiments, all methods are evaluated under the same stochastic value-oracle model. For methods using paired differences, the two queries in each pair share the same oracle realization.

Figure 1 compares the $n = 100$ finite-sum problems across all six settings. SDS reaches the best objective value in all six cases and typically makes substantially faster progress in terms of value-oracle calls, under both exact and stochastic oracles. The efficiency comes from the state-aware design of SDS: on nondegenerate iterations, it directly polls feasible directions from the local geometry, without estimating a full zeroth-order gradient. The results in Appendix E show similar overall trends across the other objective families and dimensions $n = 50, 100$, and 300, together with the corresponding iteration curves. Degenerate states are also rare in practice, occurring in fewer than 2% of iterations in most experimental settings.

6 Conclusion

We introduced state-aware direct search (SDS) for zeroth-order DR-submodular maximization over polytopes. SDS directly polls feasible directions in nondegenerate states and uses reduced-gradient estimation only under degeneracy. Combined with suitable non-oblivious objectives, SDS provides guarantees for monotone general, nonmonotone general, and nonmonotone down-closed polytopes, with $O(\epsilon^{-2})$ expected iterations under both exact and stochastic value oracles. Experiments show favorable value-oracle efficiency and that degenerate states are rarely encountered, supporting the use of local geometry in place of full gradient reconstruction at every iteration.

AI USE STATEMENT

We used generative AI tools, primarily OpenAI ChatGPT, selectively as assistive tools during the research and writing process. They were used for literature retrieval, brainstorming and discussion of research ideas, checking intermediate mathematical derivations, assistance with experimental implementation, drafting or restructuring selected passages, and language polishing. The research methodology, theoretical results, proofs, experimental design, and scientific conclusions were developed and validated by the authors. All mathematical claims, references, experimental procedures, and reported results were independently checked by the authors. The authors take full responsibility for the content of the paper.

REFERENCES

An Bian, Kfir Y. Levy, Andreas Krause, and Joachim M. Buhmann. Continuous DR-submodular maximization: Structure and algorithms. In *Advances in Neural Information Processing Systems*, volume 30, pp. 486–496, 2017a.

Andrew An Bian, Baharan Mirzasoleiman, Joachim M. Buhmann, and Andreas Krause. Guaranteed non-convex optimization: Submodular maximization over continuous domains. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pp. 111–120. PMLR, 2017b.

Yatao Bian, Joachim M. Buhmann, and Andreas Krause. Optimal continuous DR-submodular maximization and applications to provable mean field inference. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 644–653. PMLR, 2019.

Jose H. Blanchet, Coralia Cartis, Matt Menickelly, and Katya Scheinberg. Convergence rate analysis of a stochastic trust-region method via supermartingales. *INFORMS Journal on Optimization*, 1 (2):92–119, 2019. doi: 10.1287/ijoo.2019.0016.

Lin Chen, Mingrui Zhang, Hamed Hassani, and Amin Karbasi. Black box submodular maximization: Discrete and continuous settings. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1058–1070. PMLR, 2020a.

Shengminjie Chen, Donglei Du, Wenguo Yang, Dachuan Xu, and Suixiang Gao. Continuous non-monotone DR-submodular maximization with down-closed convex constraint. *arXiv preprint arXiv:2307.09616*, 2023. URL https://arxiv.org/abs/2307.09616.

Shengminjie Chen, Xiaoming Sun, Wenguo Yang, Jialin Zhang, and Zihan Zhao. DR-submodular maximization with stochastic biased gradients: Classical and quantum gradient algorithms. In *International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=0CZAimzcVr.

Wei Chen, Weizhong Zhang, and Haoyu Zhao. Gradient method for continuous influence maximization with budget-saving considerations. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 43–50, 2020b. doi: 10.1609/aaai.v34i01.5332.

Wei Chen, Xiaoming Sun, Jialin Zhang, and Zhijie Zhang. Network inference and influence maximization from samples. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 1707–1716. PMLR, 2021.

Komei Fukuda and Alain Prodon. Double description method revisited. In Michel Deza, Reinhardt Euler, and Ioannis Manoussakis (eds.), *Combinatorics and Computer Science*, volume 1120 of *Lecture Notes in Computer Science*, pp. 91–111. Springer, 1996. doi: 10.1007/3-540-61576-8_77.

Serge Gratton, Clément W. Royer, Luís Nunes Vicente, and Zaikun Zhang. Direct search based on probabilistic feasible descent for bound and linearly constrained problems. *Computational Optimization and Applications*, 72(3):525–559, 2019. doi: 10.1007/s10589-019-00062-4.

Hamed Hassani, Mahdi Soltanolkotabi, and Amin Karbasi. Gradient methods for submodular maximization. In *Advances in Neural Information Processing Systems*, volume 30, pp. 5841–5851, 2017.

Alan J. Hoffman. On approximate solutions of systems of linear inequalities. *Journal of Research of the National Bureau of Standards*, 49(4):263–265, 1952. doi: 10.6028/jres.049.027.

Tamara G. Kolda, Robert Michael Lewis, and Virginia Torczon. Stationarity results for generating set search for linearly constrained optimization. *SIAM Journal on Optimization*, 17(4):943–968, 2006. doi: 10.1137/S1052623403433638.

Robert Michael Lewis, Anne Shepherd, and Virginia Torczon. Implementing generating set search methods for linearly constrained minimization. *SIAM Journal on Scientific Computing*, 29(6): 2507–2530, 2007. doi: 10.1137/050635432.

Yuefang Lian, Xiao Wang, Dachuan Xu, and Zhongrui Zhao. Zeroth-order stochastic approximation algorithms for DR-submodular optimization. *Journal of Machine Learning Research*, 25(391): 1–55, 2024.

Yiyang Lu, Hareshkumar Jadav, Mohammad Pedramfar, Ranveer Singh, and Vaneet Aggarwal. Upper-linearizability of online non-monotone DR-submodular maximization over down-closed convex sets. In *Proceedings of the 43rd International Conference on Machine Learning*, volume 306 of *Proceedings of Machine Learning Research*. PMLR, 2026. URL https://arxiv.org/abs/2602.20578.

Mohammad Pedramfar and Vaneet Aggarwal. From linear to linearizable optimization: A novel framework with applications to stationary and non-stationary DR-submodular optimization. In *Advances in Neural Information Processing Systems*, volume 37, 2024. doi: 10.52202/079017-1188.

Mohammad Pedramfar, Christopher John Quinn, and Vaneet Aggarwal. A unified approach for maximizing continuous DR-submodular functions. In *Advances in Neural Information Processing Systems*, volume 36, 2023. doi: 10.52202/075280-2671.

Francesco Rinaldi, Luís Nunes Vicente, and Damiano Zeffiro. Stochastic trust-region and direct-search methods: A weak tail bound condition and reduced sample sizing. *SIAM Journal on Optimization*, 34(2):2067–2092, 2024. doi: 10.1137/22M1543446.

Clément W. Royer. A derivative-free algorithm for continuous submodular optimization. COCANA Seminar, 2025. URL https://www.lamsade.dauphine.fr/~croyer/docs/slidesUBC25.pdf. Presentation slides, July 29, 2025. Joint work with Marc Kaspar.

Matthew Staib and Stefanie Jegelka. Robust budget allocation via continuous submodular functions. *Applied Mathematics & Optimization*, 82(3):1049–1079, 2020. doi: 10.1007/s00245-019-09567-0.

Qixin Zhang, Zengde Deng, Zaiyi Chen, Haoyuan Hu, and Yu Yang. Stochastic continuous submodular maximization: Boosting via non-oblivious function. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 26116–26134. PMLR, 2022.

Qixin Zhang, Zongqi Wan, Zengde Deng, Zaiyi Chen, Xiaoming Sun, Jialin Zhang, and Yu Yang. Boosting gradient ascent for continuous DR-submodular maximization. *arXiv preprint arXiv:2401.08330*, 2024. URL https://arxiv.org/abs/2401.08330.

Günter M. Ziegler. *Lectures on Polytopes*, volume 152 of *Graduate Texts in Mathematics*. Springer, New York, 1995. doi: 10.1007/978-1-4613-8431-1.

A NON-OBLIVIOUS OBJECTIVES

We prove the regularity and gap inequalities used in Section 2. Both statements below include all three objectives in equation 4.

A.1 DIFFERENTIATION AND REGULARITY

Choose $\bar{\mu} > 0$ such that $\{\boldsymbol{u} \in \mathbb{R}^n : \operatorname{dist}(\boldsymbol{u}, \mathcal{X}) \le \bar{\mu}\} \subseteq \mathcal{U}$.All transformed finite-difference query points considered below lie in $\mathcal{U}$.

Proposition A.1 (Properties of the non-oblivious objectives)**.** *For the objectives in equation 4, differentiation under the integral sign gives*

$$\nabla F(\boldsymbol{x}) = \begin{cases} \displaystyle\int_0^1 e^{z-1} \nabla f(z\boldsymbol{x})\, dz, & \textit{monotone, general}, \\ \displaystyle\int_0^1 \frac{\nabla f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big)}{8(1 - z/2)^3}\, dz, & \textit{nonmonotone, general}, \\ \displaystyle\int_0^1 \frac{e^{z-1}}{1 - e^{-1}} \operatorname{Diag}(e^{-z\boldsymbol{x}}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{x}})\, dz, & \textit{nonmonotone, down-closed}. \end{cases} \tag{16}$$

Here $\operatorname{Diag}(\boldsymbol{v})$ *denotes the diagonal matrix with diagonal* $\boldsymbol{v}$*. On the finite-difference neighborhood considered below, valid Lipschitz and smoothness constants for the three non-oblivious objectives are*

$$(L_f, L_g) = \begin{cases} ((1 - e^{-1})L, e^{-1}L_1), & \textit{monotone, general}, \\ (3L/8, L_1/8), & \textit{nonmonotone, general}, \\ (eL, (e^2 L_1 + eL)/(e - 1)), & \textit{nonmonotone, down-closed}. \end{cases} \tag{17}$$

Proof. Given equation 16, it remains to verify the stated Lipschitz and smoothness constants on the finite-difference neighborhood.

For the monotone objective, $\|\nabla f\| \le L$ gives $\|\nabla F(\boldsymbol{u})\| \le L \int_0^1 e^{z-1} dz = (1 - e^{-1})L$. Moreover, by L_1-smoothness of f, $\|\nabla f(z\boldsymbol{u}) - \nabla f(z\boldsymbol{v})\| \le L_1 z \|\boldsymbol{u} - \boldsymbol{v}\|$. Hence $\|\nabla F(\boldsymbol{u}) - \nabla F(\boldsymbol{v})\| \le e^{-1} L_1 \|\boldsymbol{u} - \boldsymbol{v}\|$, since $\int_0^1 z e^{z-1} dz = e^{-1}$.

For the general nonmonotone objective, $\int_0^1 [8(1 - z/2)^3]^{-1} dz = 3/8$, and therefore $\|\nabla F(\boldsymbol{u})\| \le 3L/8$. The transformed arguments corresponding to $\boldsymbol{u}$ and $\boldsymbol{v}$ differ by $\frac{z}{2}(\boldsymbol{u} - \boldsymbol{v})$. Thus $\|\nabla f(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{u} - \underline{\boldsymbol{x}})) - \nabla f(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{v} - \underline{\boldsymbol{x}}))\| \le \frac{L_1 z}{2} \|\boldsymbol{u} - \boldsymbol{v}\|$. Using $\int_0^1 z[16(1 - z/2)^3]^{-1} dz = 1/8$ gives $\|\nabla F(\boldsymbol{u}) - \nabla F(\boldsymbol{v})\| \le (L_1/8)\|\boldsymbol{u} - \boldsymbol{v}\|$.

For the down-closed objective, every point in the finite-difference neighborhood has coordinates in $[-1, 2]$, so $\|\operatorname{Diag}(e^{-z\boldsymbol{u}})\| \le e$. Since $e^{z-1}/(1 - e^{-1})$ integrates to one, $\|\nabla F(\boldsymbol{u})\| \le eL$. For smoothness, let $\boldsymbol{u}, \boldsymbol{v}$ lie in the same finite-difference neighborhood. The mean value theorem gives $\|(\mathbf{1} - e^{-z\boldsymbol{u}}) - (\mathbf{1} - e^{-z\boldsymbol{v}})\| \le ez\|\boldsymbol{u} - \boldsymbol{v}\|$ and $\|\operatorname{Diag}(e^{-z\boldsymbol{u}}) - \operatorname{Diag}(e^{-z\boldsymbol{v}})\| \le ez\|\boldsymbol{u} - \boldsymbol{v}\|$. Adding and subtracting $\operatorname{Diag}(e^{-z\boldsymbol{u}})\nabla f(\mathbf{1} - e^{-z\boldsymbol{v}})$, and using $\|\operatorname{Diag}(e^{-z\boldsymbol{u}})\| \le e$, the L_1-smoothness of f, and $\|\nabla f\| \le L$, we obtain

$$\begin{aligned} &\|\operatorname{Diag}(e^{-z\boldsymbol{u}})\nabla f(\mathbf{1} - e^{-z\boldsymbol{u}}) - \operatorname{Diag}(e^{-z\boldsymbol{v}})\nabla f(\mathbf{1} - e^{-z\boldsymbol{v}})\| \\ &\quad \le eL_1 \|(\mathbf{1} - e^{-z\boldsymbol{u}}) - (\mathbf{1} - e^{-z\boldsymbol{v}})\| + L\|\operatorname{Diag}(e^{-z\boldsymbol{u}}) - \operatorname{Diag}(e^{-z\boldsymbol{v}})\| \\ &\quad \le z(e^2 L_1 + eL)\|\boldsymbol{u} - \boldsymbol{v}\|. \end{aligned}$$

Integrating with weight $e^{z-1}/(1 - e^{-1})$ and using $\int_0^1 z e^{z-1}/(1 - e^{-1})\, dz = 1/(e - 1)$ gives

$$\|\nabla F(\boldsymbol{u}) - \nabla F(\boldsymbol{v})\| \le \frac{e^2 L_1 + eL}{e - 1} \|\boldsymbol{u} - \boldsymbol{v}\|.$$

Since the finite-difference neighborhood is convex, the gradient-norm bounds give the stated Lipschitz constants, while the bounds above give the stated smoothness constants. □

In particular, $F_{\max} := \max_{\boldsymbol{x} \in \mathcal{P}} F(\boldsymbol{x})$ is finite and attained. The iteration proof requires this boundedness, but not nonnegativity or DR-submodularity of F.

A.2 STATIONARITY TO APPROXIMATION GUARANTEES

For vectors in $\mathcal{X}$, $\boldsymbol{x} \wedge \boldsymbol{y}$ and $\boldsymbol{x} \vee \boldsymbol{y}$ denote the componentwise minimum and maximum. We use the following consequence of nonnegativity and DR-submodularity (Bian et al., 2017a).

Lemma A.2. *For every $\boldsymbol{p}, \boldsymbol{q} \in \mathcal{X}$,*

$$f(\boldsymbol{p} \vee \boldsymbol{q}) \geq (1 - \|\boldsymbol{p}\|_\infty) f(\boldsymbol{q}).$$

Proof. If $\|\boldsymbol{p}\|_\infty = 0$, the claim is immediate. Otherwise, $(\boldsymbol{p} \vee \boldsymbol{q} - \boldsymbol{q})/\|\boldsymbol{p}\|_\infty \geq \mathbf{0}$. Moreover, $\boldsymbol{q} + (\boldsymbol{p} \vee \boldsymbol{q} - \boldsymbol{q})/\|\boldsymbol{p}\|_\infty \in \mathcal{X}$: whenever $p_i > q_i$, $p_i - q_i \leq \|\boldsymbol{p}\|_\infty - q_i \leq \|\boldsymbol{p}\|_\infty(1 - q_i)$, while the other coordinates are unchanged. DR-submodularity implies concavity along nonnegative directions. Since $\boldsymbol{p} \vee \boldsymbol{q}$ lies on the segment between $\boldsymbol{q}$ and $\boldsymbol{q} + (\boldsymbol{p} \vee \boldsymbol{q} - \boldsymbol{q})/\|\boldsymbol{p}\|_\infty$, with weight $\|\boldsymbol{p}\|_\infty$, we obtain $f(\boldsymbol{p} \vee \boldsymbol{q}) \geq (1 - \|\boldsymbol{p}\|_\infty) f(\boldsymbol{q}) + \|\boldsymbol{p}\|_\infty f(\boldsymbol{q} + (\boldsymbol{p} \vee \boldsymbol{q} - \boldsymbol{q})/\|\boldsymbol{p}\|_\infty)$. The result follows from the nonnegativity of f. □

Theorem A.3 (Non-oblivious gap estimates)**.** *Let $\boldsymbol{x}, \boldsymbol{y} \in \mathcal{P}$, and select F as in equation 4. The corresponding inequalities are*

$$\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \geq (1 - e^{-1}) f(\boldsymbol{y}) - f(\boldsymbol{x}), \qquad \textit{monotone, general}, \tag{18}$$

$$\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \geq \frac{1 - \|\underline{\boldsymbol{x}}\|_\infty}{4} f(\boldsymbol{y}) - f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right), \quad \textit{nonmonotone, general}, \tag{19}$$

$$(1 - e^{-1}) \langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \geq e^{-1} f(\boldsymbol{y}) - f(\mathbf{1} - e^{-\boldsymbol{x}}), \qquad \textit{nonmonotone, down-closed}. \tag{20}$$

Consequently, $\pi(\boldsymbol{x}) \leq \varepsilon = \epsilon / D^*$ gives all three feasible-output guarantees in equation 5.

Proof. We first consider the monotone case. By equation 16,

$$\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle = \int_0^1 e^{z-1} \langle \nabla f(z\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \, dz.$$

Since f is monotone and DR-submodular,

$$\langle \nabla f(z\boldsymbol{x}), \boldsymbol{y} \rangle \geq \langle \nabla f(z\boldsymbol{x}), \boldsymbol{y} \vee (z\boldsymbol{x}) - z\boldsymbol{x} \rangle \geq f(\boldsymbol{y} \vee (z\boldsymbol{x})) - f(z\boldsymbol{x}) \geq f(\boldsymbol{y}) - f(z\boldsymbol{x}).$$

Moreover,

$$\begin{aligned} \int_0^1 e^{z-1} \langle \nabla f(z\boldsymbol{x}), \boldsymbol{x} \rangle \, dz &= \int_0^1 e^{z-1} \frac{d}{dz} f(z\boldsymbol{x}) \, dz \\ &= f(\boldsymbol{x}) - \int_0^1 e^{z-1} f(z\boldsymbol{x}) \, dz, \end{aligned}$$

where the boundary term at $z = 0$ is zero because $f(\mathbf{0}) = 0$. Therefore,

$$\begin{aligned} \langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle &\geq \int_0^1 e^{z-1} \big(f(\boldsymbol{y}) - f(z\boldsymbol{x}) \big) \, dz - \int_0^1 e^{z-1} \langle \nabla f(z\boldsymbol{x}), \boldsymbol{x} \rangle \, dz \\ &= (1 - e^{-1}) f(\boldsymbol{y}) - f(\boldsymbol{x}). \end{aligned}$$

This proves equation 18.

We next consider the general nonmonotone case. By equation 16,

$$\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle = \int_0^1 \frac{\left\langle \nabla f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big), \boldsymbol{y} - \boldsymbol{x} \right\rangle}{8(1 - z/2)^3} \, dz.$$

We split $\boldsymbol{y} - \boldsymbol{x} = (\boldsymbol{y} - \underline{\boldsymbol{x}} - \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})) - (\boldsymbol{x} - \underline{\boldsymbol{x}} - \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}}))$. Since f is DR-submodular and nonnegative on $\mathcal{X}$,

$$\begin{aligned} &\left\langle \nabla f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big), \boldsymbol{y} - \underline{\boldsymbol{x}} - \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}}) \right\rangle \\ &\geq f\Big(\boldsymbol{y} \vee \Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big)\Big) - 2f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big). \end{aligned}$$

Indeed, decompose the displacement into its nonpositive part $\boldsymbol{y} \wedge (\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})) - (\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}}))$ and nonnegative part $\boldsymbol{y} \vee (\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})) - (\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}}))$, and apply concavity along each direction.

Moreover,

$$\begin{aligned}
&\int_0^1 \frac{\left\langle \nabla f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big), \boldsymbol{x} - \underline{\boldsymbol{x}} - \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}}) \right\rangle}{8(1 - z/2)^3}\, dz \\
&= f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right) - \frac{1}{4} f(\underline{\boldsymbol{x}}) - \int_0^1 \frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big)}{4(1 - z/2)^3}\, dz,
\end{aligned}$$

where the equality follows from

$$\frac{d}{dz} f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big) = \frac{1}{2} \left\langle \nabla f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big), \boldsymbol{x} - \underline{\boldsymbol{x}} \right\rangle$$

and integration by parts.

Combining the two bounds gives

$$\begin{aligned}
\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle &\geq \int_0^1 \frac{f\big(\boldsymbol{y} \vee \big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big)\big)}{8(1 - z/2)^3}\, dz \\
&\quad - f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right) + \frac{1}{4} f(\underline{\boldsymbol{x}}) \\
&\geq \int_0^1 \frac{f\big(\boldsymbol{y} \vee \big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big)\big)}{8(1 - z/2)^3}\, dz - f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right).
\end{aligned}$$

By Lemma A.2,

$$f\Big(\boldsymbol{y} \vee \Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big)\Big) \geq \Big(1 - \Big\|\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big\|_\infty\Big) f(\boldsymbol{y}).$$

Since $\boldsymbol{x}, \underline{\boldsymbol{x}} \in [0,1]^n$,

$$\Big\|\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\Big\|_\infty \leq \Big(1 - \frac{z}{2}\Big) \|\underline{\boldsymbol{x}}\|_\infty + \frac{z}{2}.$$

Therefore,

$$\begin{aligned}
\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle &\geq (1 - \|\underline{\boldsymbol{x}}\|_\infty) f(\boldsymbol{y}) \int_0^1 \frac{dz}{8(1 - z/2)^2} - f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right) \\
&= \frac{1 - \|\underline{\boldsymbol{x}}\|_\infty}{4} f(\boldsymbol{y}) - f\left(\frac{\boldsymbol{x} + \underline{\boldsymbol{x}}}{2}\right).
\end{aligned}$$

This proves equation 19.

Finally, consider the down-closed nonmonotone case. We verify the inequality of Lu et al. (2026) with the normalization in equation 4. For fixed $z \in [0,1]$, if $\|\mathbf{1} - e^{-z\boldsymbol{x}}\|_\infty > 0$, then

$$\boldsymbol{y} + \frac{\mathrm{Diag}(\mathbf{1} - \boldsymbol{y})(\mathbf{1} - e^{-z\boldsymbol{x}})}{\|\mathbf{1} - e^{-z\boldsymbol{x}}\|_\infty} \in \mathcal{X}.$$

Indeed, every coordinate of the added vector is nonnegative and at most $1 - y_i$. Concavity along this nonnegative direction and nonnegativity of f therefore give

$$\begin{aligned}
f\big(\mathbf{1} - e^{-z\boldsymbol{x}} + \mathrm{Diag}(e^{-z\boldsymbol{x}})\boldsymbol{y}\big) &\geq \big(1 - \|\mathbf{1} - e^{-z\boldsymbol{x}}\|_\infty\big) f(\boldsymbol{y}) \\
&= e^{-z\|\boldsymbol{x}\|_\infty} f(\boldsymbol{y}) \geq e^{-z} f(\boldsymbol{y}).
\end{aligned}$$

The same inequality is immediate when $\|\mathbf{1} - e^{-z\boldsymbol{x}}\|_\infty = 0$.

Applying concavity once more, now along the nonnegative direction $\mathrm{Diag}(e^{-z\boldsymbol{x}})\boldsymbol{y}$, yields

$$\big\langle \nabla f(\mathbf{1} - e^{-z\boldsymbol{x}}), \mathrm{Diag}(e^{-z\boldsymbol{x}})\boldsymbol{y} \big\rangle \geq e^{-z} f(\boldsymbol{y}) - f(\mathbf{1} - e^{-z\boldsymbol{x}}).$$

Using equation 16 and integrating gives

$$(1 - e^{-1}) \langle \nabla F(\boldsymbol{x}), \boldsymbol{y} \rangle \geq e^{-1} f(\boldsymbol{y}) - \int_0^1 e^{z-1} f(\mathbf{1} - e^{-z\boldsymbol{x}})\, dz.$$

On the other hand,

$$(1-e^{-1})\langle \nabla F(\boldsymbol{x}), -\boldsymbol{x}\rangle = -\int_0^1 e^{z-1}\frac{d}{dz} f(\mathbf{1}-e^{-z\boldsymbol{x}})\,dz$$
$$= -f(\mathbf{1}-e^{-\boldsymbol{x}}) + e^{-1}f(\mathbf{0}) + \int_0^1 e^{z-1} f(\mathbf{1}-e^{-z\boldsymbol{x}})\,dz,$$

where the last equality follows from integration by parts. Adding the two inequalities and using $f(\mathbf{0}) \geq 0$ proves equation 20. Hence

$$\max_{\boldsymbol{y}\in\mathcal{P}} \langle \nabla F(\boldsymbol{x}), \boldsymbol{y}-\boldsymbol{x}\rangle \leq D^*\pi(\boldsymbol{x}). \tag{21}$$

Choose a maximizer of f over $\mathcal{P}$. Combining equation 21 with the appropriate gap inequality gives the first two guarantees directly. In the third case,

$$f(\mathbf{1}-e^{-\boldsymbol{x}}) \geq e^{-1}\mathrm{OPT} - (1-e^{-1})D^*\pi(\boldsymbol{x}) \geq e^{-1}\mathrm{OPT} - \epsilon.$$

The midpoint output is feasible by convexity. The exponential output is feasible because $\mathbf{0} \leq \mathbf{1} - e^{-\boldsymbol{x}} \leq \boldsymbol{x}$ and $\mathcal{P}$ is down-closed. This proves all three assertions. □

B Proof of the iteration-complexity theorem

We prove Theorem 3.5 in several steps. Section B.1 first develops the polyhedral bounds needed to connect the local polling geometry with feasibility and stationarity. Section B.2 then introduces the random quantities and filtration used throughout the stochastic analysis, together with the basic consequences of accurate function-difference estimates. Section B.3 defines the good event and shows that it occurs with conditional probability at least $p_d p_f$. Section B.4 uses this event to compare the actual stepsize dynamics with an auxiliary random walk: before reaching ε-stationarity, a sufficiently small good iteration cannot contract the radius. Section B.5 constructs a nonnegative potential and proves a uniform expected decrease proportional to Δ_k^2. These two properties are combined through the renewal–reward bound, whose proof is given in Section B.6. Finally, Section B.7 applies this bound to obtain the expected hitting-time estimate and then invokes the stationarity certificates to conclude the approximation guarantees.

B.1 Polyhedral bounds used in the proof

We use the reduced representation in equation 2. As usual, $\mathrm{cone}(\emptyset) = \{\mathbf{0}\}$. We first verify that polling directions chosen from the approximate tangent cone preserve feasibility.

Lemma B.1 (Feasible tangent steps)**.** *Let $\boldsymbol{x} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{x}} \in \mathcal{P}$, $\delta > 0$, and $\boldsymbol{d} \in T(\boldsymbol{x}, \delta)$. If $\|\boldsymbol{d}\| \leq \delta$, then $\boldsymbol{x} + W\boldsymbol{d} \in \mathcal{P}$.*

Proof. Since the columns of W lie in $\mathrm{Null}(A)$, $A(\boldsymbol{x} + W\boldsymbol{d}) = A\boldsymbol{x} = \boldsymbol{b}$. It remains to verify the reduced inequalities. For an index $i \in I(\boldsymbol{x}, \delta)$, the definition of the approximate tangent cone gives $\tilde{\boldsymbol{a}}_i^\top \boldsymbol{d} \leq 0$. Hence $\tilde{\boldsymbol{a}}_i^\top(\tilde{\boldsymbol{x}} + \boldsymbol{d}) \leq \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{x}} \leq \widetilde{b}_i$. If $i \notin I(\boldsymbol{x}, \delta)$, then

$$\tilde{\boldsymbol{a}}_i^\top \boldsymbol{d} \leq \|\tilde{\boldsymbol{a}}_i\|\,\|\boldsymbol{d}\| \leq \delta\|\tilde{\boldsymbol{a}}_i\| < \widetilde{b}_i - \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{x}}.$$

Thus $\tilde{\boldsymbol{a}}_i^\top(\tilde{\boldsymbol{x}} + \boldsymbol{d}) \leq \widetilde{b}_i$ also holds for every such index. All constraints are therefore satisfied. □

In particular, each unit polling direction in $T(\boldsymbol{x}, \delta)$ gives a feasible trial point $\boldsymbol{x} + \delta W\boldsymbol{d}$. Since $\mathcal{P}$ is convex, the segment joining the current and trial points is feasible as well. This permits the smoothness estimate used in the failed-poll argument below.

We next relate the local projected gradient to the stationarity measure $\pi(\boldsymbol{x})$. Since $T(\boldsymbol{x}, \delta)$ is only an approximate tangent cone, a small value of $\|P_{T(\boldsymbol{x},\delta)}[W^\top \nabla F(\boldsymbol{x})]\|$ does not control all feasible displacements unless δ is also small. We use Hoffman's error bound (Hoffman, 1952) to compare each feasible displacement with a nearby direction in $T(\boldsymbol{x}, \delta)$, which gives the required local-to-global stationarity bound.

We use $H_{\mathcal{P}} < \infty$ to denote a common Hoffman constant, over all $J \subseteq \{1, \dots, 2m_{\mathcal{I}}\}$, for the coefficient matrices associated with the systems

$$\widetilde{A}\tilde{\boldsymbol{d}} \le \boldsymbol{s}, \qquad \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{d}} \le 0, \quad i \in J, \qquad \boldsymbol{s} \in \mathbb{R}^{2m_{\mathcal{I}}}.$$

For each nonempty system, the constant is uniform in its right-hand side $\boldsymbol{s}$, since Hoffman's bound depends only on the coefficient matrix and the chosen norms. A common constant exists because there are finitely many subsets of rows. We next use Hoffman's error bound to quantify when a small projected gradient, together with a sufficiently small polling radius, certifies ε-stationarity.

Lemma B.2 (Hoffman-type stationarity bound)**.** *Let $\boldsymbol{x} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{x}} \in \mathcal{P}$, $\delta \ge 0$, and $T = T(\boldsymbol{x}, \delta)$. Set $C_H = \max\{1, L_f H_{\mathcal{P}} \max_{1 \le i \le 2m_{\mathcal{I}}} \|\tilde{\boldsymbol{a}}_i\|\}$. Then, for every $\boldsymbol{y} \in \mathcal{P}$,*

$$\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \le D\|P_T[\boldsymbol{g}(\boldsymbol{x})]\| + C_H\delta. \tag{22}$$

In particular, if $\|P_T[\boldsymbol{g}(\boldsymbol{x})]\| \le \varepsilon/(2D^*)$ and $\delta \le \varepsilon/(2C_H)$, then $\langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle \le \varepsilon$ for every $\boldsymbol{y} \in \mathcal{P}$, and consequently $\pi(\boldsymbol{x}) \le \varepsilon$.

Proof. Fix $\boldsymbol{y} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{y}} \in \mathcal{P}$, and set $\boldsymbol{v} := \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{x}}$. Consider

$$\mathcal{C} := T(\boldsymbol{x}, \delta) \cap \{\boldsymbol{d} : \widetilde{A}\boldsymbol{d} \le \widetilde{\boldsymbol{b}} - \widetilde{A}\tilde{\boldsymbol{x}}\}.$$

This set is closed and convex, and is nonempty because it contains zero. We apply Hoffman's error bound to the linear system defining $\mathcal{C}$, using the Euclidean norm for the distance on the left-hand side and the infinity norm for the residual on the right-hand side.

Since $\boldsymbol{y} \in \mathcal{P}$, $\boldsymbol{v}$ satisfies

$$\widetilde{A}\boldsymbol{v} \le \widetilde{\boldsymbol{b}} - \widetilde{A}\tilde{\boldsymbol{x}}.$$

Moreover, for every $i \in I(\boldsymbol{x}, \delta)$,

$$\tilde{\boldsymbol{a}}_i^\top \boldsymbol{v} = \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{y}} - \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{x}} \le \widetilde{b}_i - \tilde{\boldsymbol{a}}_i^\top \tilde{\boldsymbol{x}} \le \delta\|\tilde{\boldsymbol{a}}_i\|.$$

By Hoffman's error bound, with $\hat{\boldsymbol{v}} := P_{\mathcal{C}}\boldsymbol{v}$,

$$\begin{aligned} \|\boldsymbol{v} - \hat{\boldsymbol{v}}\| &\le H_{\mathcal{P}} \left(\left\| \left[\widetilde{A}\boldsymbol{v} - (\widetilde{\boldsymbol{b}} - \widetilde{A}\tilde{\boldsymbol{x}})\right]_+ \right\|_\infty + \max_{i \in I(\boldsymbol{x}, \delta)} [\tilde{\boldsymbol{a}}_i^\top \boldsymbol{v}]_+ \right) \\ &\le H_{\mathcal{P}}\delta \max_{1 \le i \le 2m_{\mathcal{I}}} \|\tilde{\boldsymbol{a}}_i\|. \end{aligned} \tag{23}$$

Here $[\cdot]_+$ denotes the componentwise positive part; the maximum over an empty active set is zero. The first residual vanishes because $\boldsymbol{v}$ satisfies the original reduced inequalities.

Using $\boldsymbol{g}(\boldsymbol{x}) = W^\top \nabla F(\boldsymbol{x})$,

$$\langle \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{v} \rangle = \langle \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{v} - \hat{\boldsymbol{v}} \rangle + \langle \boldsymbol{g}(\boldsymbol{x}), \hat{\boldsymbol{v}} \rangle.$$

For the first term, using the L_f-Lipschitz continuity of F,

$$\langle \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{v} - \hat{\boldsymbol{v}} \rangle \le L_f H_{\mathcal{P}}\delta \max_{1 \le i \le 2m_{\mathcal{I}}} \|\tilde{\boldsymbol{a}}_i\| \le C_H\delta,$$

where $\|\boldsymbol{g}(\boldsymbol{x})\| \le \|W^\top\| \, \|\nabla F(\boldsymbol{x})\| \le L_f$ and $\|W^\top\| = 1$.

For the second term, using the projection decomposition induced by the polar cones T and $N(\boldsymbol{x}, \delta)$, and since $\hat{\boldsymbol{v}} \in T$,

$$\langle \boldsymbol{g}(\boldsymbol{x}), \hat{\boldsymbol{v}} \rangle \le \langle P_T[\boldsymbol{g}(\boldsymbol{x})], \hat{\boldsymbol{v}} \rangle \le \|P_T[\boldsymbol{g}(\boldsymbol{x})]\| \, \|\hat{\boldsymbol{v}}\|.$$

As $\boldsymbol{0} \in \mathcal{C}$, we have $P_{\mathcal{C}}[\boldsymbol{0}] = \boldsymbol{0}$. The nonexpansiveness of Euclidean projection therefore gives

$$\|\hat{\boldsymbol{v}}\| \le \|\boldsymbol{v}\| = \|W\boldsymbol{v}\| = \|\boldsymbol{y} - \boldsymbol{x}\| \le D.$$

Therefore,

$$\langle \boldsymbol{g}(\boldsymbol{x}), \hat{\boldsymbol{v}} \rangle \le D\|P_T[\boldsymbol{g}(\boldsymbol{x})]\|.$$

Combining the two bounds and using $\boldsymbol{y} - \boldsymbol{x} = W\boldsymbol{v}$, we obtain

$$\begin{aligned} \langle \nabla F(\boldsymbol{x}), \boldsymbol{y} - \boldsymbol{x} \rangle &= \langle \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{v} \rangle \\ &= \langle \boldsymbol{g}(\boldsymbol{x}), \boldsymbol{v} - \hat{\boldsymbol{v}} \rangle + \langle \boldsymbol{g}(\boldsymbol{x}), \hat{\boldsymbol{v}} \rangle \\ &\le C_H\delta + D\|P_T[\boldsymbol{g}(\boldsymbol{x})]\|. \end{aligned}$$

This proves equation 22. Under the two additional bounds in the statement, the right-hand side is at most $\varepsilon/2 + D\varepsilon/(2D^*) \le \varepsilon$. The proof is complete. □

The Hoffman constant is used only in the preceding deterministic transfer from projected-gradient smallness to first-order stationarity. It is not used by the algorithm. In particular, the state-aware method does not need to estimate $H_{\mathcal{P}}$ in order to generate polling directions or make acceptance decisions.

We also need a uniform bound on the normal component of a feasible reduced displacement. For every nonzero approximate normal cone $N(\boldsymbol{x}, \delta)$, the active normals $\{\tilde{\boldsymbol{a}}_i : i \in I(\boldsymbol{x}, \delta)\}$ form a finite positive generating set. Since only finitely many subsets of the rows of $\widetilde{A}$ can occur, the corresponding cosine measures admit a common positive lower bound, denoted by $\eta_{\min} > 0$. The following result is Proposition B.1 of Kolda et al. (2006), applied to the reduced polytope $\widetilde{\mathcal{P}}$.

Lemma B.3 (Normal-component bound)**.** *Let $\boldsymbol{x} = \bar{\boldsymbol{x}} + W\tilde{\boldsymbol{x}} \in \mathcal{P}$, $\delta > 0$, and let $\tilde{\boldsymbol{d}}$ satisfy $\boldsymbol{x} + W\tilde{\boldsymbol{d}} \in \mathcal{P}$. Then*

$$\|P_{N(\boldsymbol{x},\delta)}[\tilde{\boldsymbol{d}}]\| \leq \frac{\delta}{\eta_{\min}}. \tag{24}$$

B.2 Random quantities and filtration

We use the random quantities introduced in the main text. All randomness is defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Uppercase symbols denote random quantities and lowercase symbols their realizations. In particular, $\boldsymbol{X}_k$, Δ_k, and $\boldsymbol{G}_k$ have realizations $\boldsymbol{x}_k$, δ_k, and $\boldsymbol{g}_k$, respectively, while $\mathfrak{D}_k$ denotes the random polling set and $\mathcal{D}_k$ its realization. Recall that $\boldsymbol{g}(\boldsymbol{X}_k) = W^\top \nabla F(\boldsymbol{X}_k)$ is the true reduced gradient, whereas $\boldsymbol{G}_k$ is its estimate. For the analysis, define $I_k^{\mathrm{deg}} := \mathbf{1}\{Q(\boldsymbol{X}_k, \Delta_k)^\top \text{ is rank deficient}\}$ and $I_k^{\mathrm{poll}} := \mathbf{1}\{\mathfrak{D}_k \neq \emptyset\}$. The initial point and radius are deterministic, and we set $\mathcal{F}_{-1} = \sigma(\boldsymbol{X}_0, \Delta_0)$. For $k \geq 0$, $\mathcal{F}_{k-1}$ contains all information generated before iteration k. Hence $\boldsymbol{X}_k$, Δ_k, the local geometry, and I_k^{deg} are $\mathcal{F}_{k-1}$-measurable.

Within iteration k, the state-specific information is generated before the directional-increment estimates. In a degenerate iteration, this includes the reduced-gradient estimate $\boldsymbol{G}_k$; when no gradient estimate is required, we take $\boldsymbol{G}_k = \mathbf{0}$. The polling set $\mathfrak{D}_k$ is then formed, and $\mathcal{F}_{k-1/2} := \sigma(\mathcal{F}_{k-1}, \boldsymbol{G}_k, \mathfrak{D}_k)$. Thus $\boldsymbol{G}_k$, $\mathfrak{D}_k$, I_k^{poll}, H_k, and Z_k are $\mathcal{F}_{k-1/2}$-measurable. The estimates $\widetilde{A}_{k,\boldsymbol{d}}$, $\boldsymbol{d} \in \mathfrak{D}_k$, are generated afterwards, and $\mathcal{F}_k$ contains these estimates together with the resulting update $\boldsymbol{X}_{k+1}, \Delta_{k+1}$. Consequently, J_k is $\mathcal{F}_k$-measurable.

For opportunistic polling, fix the order of the candidate directions before their increment estimates are generated. For the analysis, the estimates associated with all candidates are defined on the underlying probability space, although the algorithm reveals them only until the first successful direction is found. Hence J_k is well defined, while directions not actually tested incur no value-oracle calls.

When $\mathfrak{D}_k = \emptyset$, we use the convention $J_k = \Omega$. For a nondegenerate state with $T(\boldsymbol{X}_k, \Delta_k) = \{\mathbf{0}\}$, we set $Z_k = \Omega$. Such a state produces an empty poll and contracts the radius, the directional-reliability condition is therefore needed only for nondegenerate states with a nonzero approximate tangent cone.

Finally, since F is continuously differentiable and $\mathcal{P}$ is compact, the stationarity measure π is Borel measurable and T_ε is a stopping time with respect to the filtration $\{\mathcal{F}_{k-1}\}_{k\geq 0}$. Since the algorithm only observes estimates of non-oblivious function differences, we first record how accurate estimates translate the sufficient-ascent test into statements about the true non-oblivious objective.

Lemma B.4 (Accurate poll decisions)**.** *Suppose that $I_k^{\mathrm{poll}} = 1$ and J_k occurs. If iteration k is successful, then $F(\boldsymbol{X}_{k+1}) - F(\boldsymbol{X}_k) \geq (\gamma - 1)\varepsilon_f \Delta_k^2$. If iteration k is unsuccessful, then $F(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}) - F(\boldsymbol{X}_k) < (\gamma + 1)\varepsilon_f \Delta_k^2$ for every $\boldsymbol{d} \in \mathfrak{D}_k$.*

Proof. On a successful iteration, let $\widetilde{\boldsymbol{D}}_k$ denote the accepted reduced direction. The acceptance test and J_k imply

$$\begin{aligned} F(\boldsymbol{X}_{k+1}) - F(\boldsymbol{X}_k) &= A_{k,\widetilde{\boldsymbol{D}}_k} \\ &\geq \widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - |\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - A_{k,\widetilde{\boldsymbol{D}}_k}| \\ &\geq \gamma\varepsilon_f \Delta_k^2 - \varepsilon_f \Delta_k^2 = (\gamma - 1)\varepsilon_f \Delta_k^2. \end{aligned}$$

If iteration k is unsuccessful, every candidate has been tested, and $\widetilde{A}_{k,\boldsymbol{d}} < \gamma\varepsilon_f\Delta_k^2$ for every $\boldsymbol{d} \in \mathfrak{D}_k$. On J_k, therefore, $A_{k,\boldsymbol{d}} \le \widetilde{A}_{k,\boldsymbol{d}} + |\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| < (\gamma+1)\varepsilon_f\Delta_k^2$. □

Thus, on the accuracy event J_k, a successful iteration produces a true increase in the non-oblivious objective. This is the reason for imposing $\gamma > 1$ in the sufficient-ascent test.

We next use a renewal–reward argument to bound the expected hitting time T_ε. The potential decrease controls a cumulative reward before the stopping time, while the stepsize updates are compared with a random walk whose upward moves occur with conditional probability at least $p_d p_f$. Positive conditional drift ensures that the stepsize is above δ_ε often enough, and each such renewal contributes a fixed amount to the accumulated reward. We record the bound in the form needed below.

Theorem B.5 (Renewal–reward bound)**.** *Let $\{\mathcal{F}_k\}_{k\ge -1}$ be a filtration. Suppose that Δ_0 and Φ_0 are deterministic and $\Delta_0 \ge \delta_\varepsilon$. Let $\{\Phi_k\}_{k\ge 0}$ and $\{\Delta_k\}_{k\ge 0}$ be nonnegative processes such that Φ_k and Δ_k are $\mathcal{F}_{k-1}$-measurable for every $k \ge 0$. Let $W_k \in \{\log\beta, \log\theta\}$ be $\mathcal{F}_k$-measurable and satisfy*

$$\mathbb{P}(W_k = \log\beta \mid \mathcal{F}_{k-1}) \ge p, \qquad p := p_d p_f.$$

Assume $p\log\beta + (1-p)\log\theta > 0$. Let T_ε be a stopping time with respect to $\{\mathcal{F}_{k-1}\}_{k\ge 0}$. Suppose that $\Delta_k \le \delta_{\max}$ and, for some $\delta_\varepsilon > 0$,

$$\mathbf{1}_{\{T_\varepsilon > k\}}\Delta_{k+1} \ge \mathbf{1}_{\{T_\varepsilon > k\}}\min\{\Delta_k e^{W_k}, \delta_\varepsilon\}.$$

Finally, suppose that there exists $M > 0$ such that

$$\mathbb{E}\big[(\Phi_{k+1} - \Phi_k)\mathbf{1}_{\{T_\varepsilon > k\}} \,\big|\, \mathcal{F}_{k-1}\big] \le -M\Delta_k^2\mathbf{1}_{\{T_\varepsilon > k\}}.$$

Then

$$\mathbb{E}[T_\varepsilon] \le \frac{\log\beta}{p\log\beta + (1-p)\log\theta}\frac{\Phi_0}{M\delta_\varepsilon^2} + 1.$$

The proof follows the renewal–reward construction of Blanchet et al. (2019), with logarithmic moves $\log\beta$ and $\log\theta$. We give the complete proof in Section B.6.

B.3 The good event and its probability

We now verify the hypotheses of Theorem B.5. The first step is to identify the iterations that play the role of upward moves in the auxiliary stepsize process. These are not defined directly by success of the algorithm, but by the occurrence of the random information needed to justify the deterministic polling argument. For iteration k, define the good event

$$\begin{aligned}\mathcal{G}_k := {} & \{I_k^{\mathrm{deg}} = 0\} \cap Z_k \cap J_k \\ & \cup \{I_k^{\mathrm{deg}} = 1,\ I_k^{\mathrm{poll}} = 1\} \cap H_k \cap J_k \\ & \cup \{I_k^{\mathrm{deg}} = 1,\ I_k^{\mathrm{poll}} = 0\} \cap H_k.\end{aligned} \tag{25}$$

Thus, in a nondegenerate iteration, the event $\mathcal{G}_k$ requires both a feasible-ascent polling set and sufficiently accurate estimates of the non-oblivious function differences. In a degenerate iteration with a polling direction, it requires a sufficiently accurate zeroth-order gradient estimate and sufficiently accurate estimates of the non-oblivious function differences. In a degenerate iteration without a polling direction, it only requires the zeroth-order gradient estimate to be sufficiently accurate.

We can now define the auxiliary random variable used in Theorem B.5. It records whether the random information at iteration k is good enough for the deterministic polling argument:

$$W_k := \begin{cases}\log\beta, & \text{on } \mathcal{G}_k, \\ \log\theta, & \text{on } \mathcal{G}_k^c.\end{cases} \tag{26}$$

The event $\mathcal{G}_k$ is $\mathcal{F}_k$-measurable, and hence W_k is $\mathcal{F}_k$-measurable.

Lemma B.6 (Probability of the good event)**.** *For every $k \ge 0$,*

$$\mathbb{P}(W_k = \log\beta \mid \mathcal{F}_{k-1}) = \mathbb{P}(\mathcal{G}_k \mid \mathcal{F}_{k-1}) \ge p_d p_f.$$

Proof. The three cases in equation 25 are disjoint, and hence

$$\begin{aligned}\mathbf{1}_{\mathcal{G}_k} = \big[(1 - I_k^{\mathrm{deg}})\mathbf{1}_{Z_k} + I_k^{\mathrm{deg}} I_k^{\mathrm{poll}} \mathbf{1}_{H_k}\big]\mathbf{1}_{J_k} \\ + I_k^{\mathrm{deg}}(1 - I_k^{\mathrm{poll}})\mathbf{1}_{H_k}.\end{aligned}$$

The terms multiplying $\mathbf{1}_{J_k}$ are $\mathcal{F}_{k-1/2}$-measurable. Therefore, by the tower property and Assumption 3.4,

$$\begin{aligned}\mathbb{E}[\mathbf{1}_{\mathcal{G}_k} \mid \mathcal{F}_{k-1}] &\geq p_f\, \mathbb{E}\Big[(1 - I_k^{\mathrm{deg}})\mathbf{1}_{Z_k} + I_k^{\mathrm{deg}} I_k^{\mathrm{poll}} \mathbf{1}_{H_k} \,\Big|\, \mathcal{F}_{k-1}\Big] \\ &\quad + \mathbb{E}\Big[I_k^{\mathrm{deg}}(1 - I_k^{\mathrm{poll}})\mathbf{1}_{H_k} \,\Big|\, \mathcal{F}_{k-1}\Big] \\ &\geq p_f\, \mathbb{E}\Big[(1 - I_k^{\mathrm{deg}})\mathbf{1}_{Z_k} + I_k^{\mathrm{deg}} \mathbf{1}_{H_k} \,\Big|\, \mathcal{F}_{k-1}\Big].\end{aligned}$$

The second inequality uses $p_f \leq 1$. Since I_k^{deg} is $\mathcal{F}_{k-1}$-measurable, the reliability bounds for Z_k and H_k, together with the convention $Z_k = \Omega$ on a nondegenerate zero cone, give

$$\begin{aligned}\mathbb{E}[\mathbf{1}_{\mathcal{G}_k} \mid \mathcal{F}_{k-1}] &\geq p_f\big[p_d(1 - I_k^{\mathrm{deg}}) + p_d I_k^{\mathrm{deg}}\big] \\ &= p_d p_f.\end{aligned}$$

This proves the claim. □

B.4 THE STEPSIZE COMPARISON

Thus the conditional probability condition in Theorem B.5 is satisfied with $p = p_d p_f$. We next verify the stepsize comparison condition. The threshold δ_ε will be chosen below with $0 < \delta_\varepsilon \leq \delta_0 \leq \delta_{\max}$. By the definition of W_k, it is enough to compare the actual update of Δ_k with the auxiliary update using e^{W_k}.

On $\{T_\varepsilon > k\}$, if iteration k is successful, then $\Delta_{k+1} = \min\{\beta\Delta_k, \delta_{\max}\}$. Hence the stepsize comparison holds immediately. If iteration k is unsuccessful, then $\Delta_{k+1} = \theta\Delta_k$. When $W_k = \log\theta$, the comparison again holds directly. Therefore, the only case requiring analysis is an unsuccessful iteration on the good event $\mathcal{G}_k$, where $W_k = \log\beta$. If $\Delta_k \geq \delta_\varepsilon/\theta$, then $\Delta_{k+1} = \theta\Delta_k \geq \delta_\varepsilon$, and the desired comparison still holds. It remains to rule out the case in which an iteration is unsuccessful, the good event occurs, and $\Delta_k \leq \frac{\delta_\varepsilon}{\theta}$. We consider this small-stepsize case according to whether a polling set is available, that is, according to the value of I_k^{poll}.

We next consider the case $I_k^{\mathrm{poll}} = 1$, where a polling set is available. This includes both nondegenerate iterations, in which the polling set is generated from a positive generating set of the approximate tangent cone, and degenerate iterations, in which the polling set consists of the projected zeroth-order gradient direction. The next lemma shows that, when the feasible-ascent event holds and the non-oblivious function differences are sufficiently accurate, an unsuccessful poll forces the projected reduced gradient to be at most on the order of the stepsize.

Lemma B.7 (Accurate failed poll)**.** *Suppose that $I_k^{\mathrm{poll}} = 1$, $Z_k \cap J_k$ occurs, and iteration k is unsuccessful. Then*

$$\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\| \leq \frac{1}{\kappa}\left(\frac{L_g}{2} + (\gamma+1)\varepsilon_f\right)\Delta_k. \tag{27}$$

Proof. If $P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)] = \mathbf{0}$, the claim is immediate.

Otherwise, let $\tilde{\boldsymbol{d}}_k \in \mathfrak{D}_k$ satisfy the feasible-ascent condition in Z_k. By the L_g-smoothness of F, feasibility of the trial segment, and $\|W\tilde{\boldsymbol{d}}_k\| = 1$,

$$\begin{aligned}F(\boldsymbol{X}_k + \Delta_k W\tilde{\boldsymbol{d}}_k) - F(\boldsymbol{X}_k) &\geq \Delta_k\left\langle \boldsymbol{g}(\boldsymbol{X}_k), \tilde{\boldsymbol{d}}_k\right\rangle - \frac{L_g}{2}\Delta_k^2 \\ &\geq \kappa\Delta_k\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\| - \frac{L_g}{2}\Delta_k^2,\end{aligned}$$

where the second inequality follows from Z_k. Since the poll is unsuccessful and J_k occurs, Lemma B.4 gives

$$F(\boldsymbol{X}_k + \Delta_k W \tilde{\boldsymbol{d}}_k) - F(\boldsymbol{X}_k) < (\gamma + 1)\varepsilon_f \Delta_k^2.$$

Combining the two inequalities gives equation 27. □

In a nondegenerate iteration with a nonempty poll, the good event already contains $Z_k \cap J_k$, and hence the preceding lemma applies directly. It remains to check that the same feasible-ascent condition also holds in the degenerate polling branch. This is the content of the next lemma.

Lemma B.8 (Projected direction in degenerate polling)**.** *Suppose that $I_k^{\mathrm{deg}} = 1$, $I_k^{\mathrm{poll}} = 1$, and H_k occurs. Then Z_k occurs.*

Proof. By the definition of the degenerate polling branch,

$$\mathfrak{D}_k = \left\{ \frac{P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|} \right\}, \qquad \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| > \frac{\varepsilon}{4D^*}.$$

If $P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)] = \mathbf{0}$, the cosine measure is equal to 1 by convention. We therefore assume that $P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)] \neq \mathbf{0}$.

On H_k, we have

$$\|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\| \leq \varepsilon_g = \frac{1-\kappa}{1+\kappa}\frac{\varepsilon}{4D^*} < \frac{1-\kappa}{1+\kappa}\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|.$$

Using the projection theorem for the cone $T(\boldsymbol{X}_k, \Delta_k)$,

$$\left\langle P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k], \boldsymbol{G}_k \right\rangle = \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|^2.$$

Hence,

$$\begin{aligned}
&\mathrm{cm}_{T(\boldsymbol{X}_k,\Delta_k)}\big(\mathfrak{D}_k, \boldsymbol{g}(\boldsymbol{X}_k)\big) \\
&\quad= \frac{\left\langle P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k], \boldsymbol{g}(\boldsymbol{X}_k)\right\rangle}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\,\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\|} \\
&\quad= \frac{\left\langle P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k], \boldsymbol{G}_k\right\rangle + \left\langle P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k], \boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\right\rangle}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\,\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\|} \\
&\quad\geq \frac{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|^2 - \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\,\|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\|}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\left(\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| + \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k]\|\right)} \\
&\quad\geq \frac{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|^2 - \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\,\|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\|}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|\left(\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| + \|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\|\right)} \\
&\quad= \frac{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| - \|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\|}{\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| + \|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\|} \\
&\quad\geq \kappa.
\end{aligned}$$

Here the first inequality uses the projection theorem, the Cauchy–Schwarz inequality. The second inequality uses the nonexpansiveness of projection. For the last inequality, the preceding error bound gives

$$(1+\kappa)\|\boldsymbol{g}(\boldsymbol{X}_k) - \boldsymbol{G}_k\| \leq (1-\kappa)\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\|.$$

Rearranging this inequality proves that the last ratio is at least κ. Since $I_k^{\mathrm{poll}} = 1$, this is precisely Z_k.

□

We now choose the stepsize threshold so that an unsuccessful good iteration at a sufficiently small radius can be converted into a stationarity bound. Let C_H be the constant in Lemma B.2, and define

$$L_\gamma := \frac{L_g}{2\kappa} + \frac{(\gamma+1)\varepsilon_f}{\kappa} + \frac{L_f}{\eta_{\min}},$$
$$\delta_\varepsilon := \min\left\{\delta_0,\ \frac{\theta\varepsilon}{L_\gamma},\ \frac{\theta\varepsilon}{2C_H}\right\}. \tag{28}$$

Lemma B.9 (Good small steps cannot fail)**.** *Suppose that*

$$\{T_\varepsilon > k\} \cap \mathcal{G}_k \cap \left\{\Delta_k \le \frac{\delta_\varepsilon}{\theta}\right\} \cap \{I_k^{\mathrm{poll}} = 1\}$$

is nonempty. Then, on this event, iteration k is successful.

Proof. We argue by contradiction. Fix a realization in $\{T_\varepsilon > k\} \cap \mathcal{G}_k \cap \left\{\Delta_k \le \frac{\delta_\varepsilon}{\theta}\right\} \cap \{I_k^{\mathrm{poll}} = 1\}$, and suppose that iteration k is unsuccessful. We write the realization as $\boldsymbol{x}_k$, δ_k, $T_k = T(\boldsymbol{x}_k, \delta_k)$, and $N_k = N(\boldsymbol{x}_k, \delta_k)$.

If the iteration is nondegenerate, then the definition of $\mathcal{G}_k$ implies $Z_k \cap J_k$. If the iteration is degenerate, then $\mathcal{G}_k \cap \{I_k^{\mathrm{poll}} = 1\}$ implies $H_k \cap J_k$, and Lemma B.8 gives Z_k. In both cases, Lemma B.7 applies, and hence

$$\|P_{T_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\| \le \left(\frac{L_g}{2\kappa} + \frac{(\gamma+1)\varepsilon_f}{\kappa}\right)\delta_k. \tag{29}$$

We now bound the stationarity measure. Using the reduced representation of feasible directions,

$$\begin{aligned}
\pi(\boldsymbol{x}_k) &= \max_{\substack{\boldsymbol{x}_k+\boldsymbol{d}\in\mathcal{P}\\ \|\boldsymbol{d}\|\le 1}} \langle \boldsymbol{d}, \nabla F(\boldsymbol{x}_k)\rangle \\
&= \max_{\substack{\tilde{\boldsymbol{x}}_k+\tilde{\boldsymbol{d}}\in\widetilde{\mathcal{P}}\\ \|\tilde{\boldsymbol{d}}\|\le 1}} \left\langle \tilde{\boldsymbol{d}}, \boldsymbol{g}(\boldsymbol{x}_k)\right\rangle \\
&= \max_{\substack{\tilde{\boldsymbol{x}}_k+\tilde{\boldsymbol{d}}\in\widetilde{\mathcal{P}}\\ \|\tilde{\boldsymbol{d}}\|\le 1}} \left\langle P_{T_k}[\tilde{\boldsymbol{d}}] + P_{N_k}[\tilde{\boldsymbol{d}}],\ P_{T_k}[\boldsymbol{g}(\boldsymbol{x}_k)] + P_{N_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\right\rangle \\
&\le \max_{\substack{\tilde{\boldsymbol{x}}_k+\tilde{\boldsymbol{d}}\in\widetilde{\mathcal{P}}\\ \|\tilde{\boldsymbol{d}}\|\le 1}} \left(\|P_{T_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\| + \left\|P_{N_k}[\tilde{\boldsymbol{d}}]\right\|\, \|P_{N_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\|\right) \\
&\le \max_{\substack{\tilde{\boldsymbol{x}}_k+\tilde{\boldsymbol{d}}\in\widetilde{\mathcal{P}}\\ \|\tilde{\boldsymbol{d}}\|\le 1}} \left(\|P_{T_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\| + \left\|P_{N_k}[\tilde{\boldsymbol{d}}]\right\|\, \|\nabla F(\boldsymbol{x}_k)\|\right).
\end{aligned}$$

Here the second equality follows from the reduced parametrization $\boldsymbol{d} = W\tilde{\boldsymbol{d}}$. The third equality is the Moreau decomposition for the polar cones T_k and N_k. The first inequality uses the polarity of T_k and N_k, the Cauchy–Schwarz inequality, and $\|P_{T_k}[\tilde{\boldsymbol{d}}]\| \le \|\tilde{\boldsymbol{d}}\| \le 1$. The last inequality uses $\|P_{N_k}[\boldsymbol{g}(\boldsymbol{x}_k)]\| \le \|\boldsymbol{g}(\boldsymbol{x}_k)\| \le \|\nabla F(\boldsymbol{x}_k)\|$.

By Lemma B.7, Lemma B.3, and $\|\nabla F(\boldsymbol{x}_k)\| \le L_f$, it follows that

$$\pi(\boldsymbol{x}_k) \le \left[\frac{L_g}{2\kappa} + \frac{(\gamma+1)\varepsilon_f}{\kappa} + \frac{L_f}{\eta_{\min}}\right]\delta_k = L_\gamma\delta_k.$$

Since $\delta_k \le \delta_\varepsilon/\theta$, the definition of δ_ε gives

$$\pi(\boldsymbol{x}_k) \le L_\gamma\frac{\delta_\varepsilon}{\theta} \le \varepsilon.$$

This contradicts $T_\varepsilon > k$. Therefore the iteration must be successful. □

Thus, on the good event, an unsuccessful iteration with $I_k^{\mathrm{poll}} = 1$ cannot occur when the algorithm has not yet reached ε-stationarity and the stepsize is at most $\delta_\varepsilon/\theta$. It remains to consider the case

$I_k^{\text{poll}} = 0$. By the definition of δ_ε, the small-stepsize condition in Lemma B.2 is satisfied whenever $\Delta_k \le \delta_\varepsilon/\theta$. To apply Lemma B.2, it remains to verify the corresponding projected-gradient bound. This is provided by the next lemma.

Lemma B.10 (No-poll degenerate iterations)**.** *Suppose that $I_k^{\text{deg}} = 1$, $I_k^{\text{poll}} = 0$, and H_k occurs. Then*

$$\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\| \le \frac{\varepsilon}{2D^*}.$$

Proof. If $I_k^{\text{poll}} = 0$, then the degenerate branch of Algorithm 1 gives

$$\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| \le \frac{\varepsilon}{4D^*}.$$

Therefore, by the nonexpansiveness of projection and H_k,

$$\begin{aligned}\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\| &\le \|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{G}_k]\| + \|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \\ &\le \frac{\varepsilon}{4D^*} + \varepsilon_g \le \frac{\varepsilon}{2D^*}.\end{aligned}$$

The last inequality uses $\varepsilon_g = \frac{1-\kappa}{1+\kappa}\frac{\varepsilon}{4D^*} \le \varepsilon/(4D^*)$.

□

Thus, if $I_k^{\text{deg}} = 1$, $I_k^{\text{poll}} = 0$, $\mathcal{G}_k$ occurs, and $\Delta_k \le \delta_\varepsilon/\theta$, then Lemma B.10 gives the required projected gradient bound, while equation 28 gives $\Delta_k \le \varepsilon/(2C_H)$. The two hypotheses of Lemma B.2 are therefore satisfied, and $\pi(\boldsymbol{X}_k) \le \varepsilon$. This contradicts $T_\varepsilon > k$.

If $I_k^{\text{deg}} = 0$ and $I_k^{\text{poll}} = 0$, then $T(\boldsymbol{X}_k, \Delta_k) = \{\mathbf{0}\}$, and hence

$$\|P_{T(\boldsymbol{X}_k,\Delta_k)}[\boldsymbol{g}(\boldsymbol{X}_k)]\| = 0.$$

If also $\Delta_k \le \delta_\varepsilon/\theta$, then Lemma B.2 gives $\pi(\boldsymbol{X}_k) \le \varepsilon$, contradicting $T_\varepsilon > k$.

Together with Lemma B.9, the empty-poll arguments show that, on

$$\{T_\varepsilon > k\} \cap \mathcal{G}_k \cap \{\Delta_k \le \delta_\varepsilon/\theta\},$$

iteration k must be successful.

Lemma B.11 (Stepsize dynamic)**.** *Let δ_ε be defined by equation 28. Then, for every $k \ge 0$,*

$$\mathbf{1}_{\{T_\varepsilon > k\}}\Delta_{k+1} \ge \mathbf{1}_{\{T_\varepsilon > k\}} \min\{\Delta_k e^{W_k}, \delta_\varepsilon\}. \tag{30}$$

Proof. It suffices to consider $\{T_\varepsilon > k\}$. If $W_k = \log\theta$, the update rule and $\delta_\varepsilon \le \delta_{\max}$ give

$$\Delta_{k+1} \ge \min\{\theta\Delta_k, \delta_\varepsilon\} = \min\{\Delta_k e^{W_k}, \delta_\varepsilon\}.$$

Suppose now that $W_k = \log\beta$, equivalently, $\mathcal{G}_k$ occurs. If $\Delta_k \le \delta_\varepsilon/\theta$, the preceding result implies that iteration k is successful, and hence

$$\Delta_{k+1} = \min\{\beta\Delta_k, \delta_{\max}\} \ge \min\{\beta\Delta_k, \delta_\varepsilon\}.$$

If $\Delta_k > \delta_\varepsilon/\theta$, the same inequality holds on a successful iteration, while on an unsuccessful iteration $\Delta_{k+1} = \theta\Delta_k > \delta_\varepsilon$. Thus equation 30 holds in every case. □

B.5 Expected potential decrease

It remains to verify the potential-decrease condition in Theorem B.5. We use a potential that balances two quantities: the remaining gap in the non-oblivious objective and the current stepsize. The first term decreases on accurate successful iterations, while the second term decreases on unsuccessful iterations. A suitable choice of the weight ν ensures that the combined potential decreases in conditional expectation.

For $\nu \in (0,1)$, define

$$\Phi_k := \frac{\nu}{\varepsilon_f}\big(F_{\max} - F(\boldsymbol{X}_k)\big) + (1-\nu)\Delta_k^2, \qquad F_{\max} := \max_{\boldsymbol{x}\in\mathcal{P}} F(\boldsymbol{x}). \tag{31}$$

Continuity of F and compactness of $\mathcal{P}$ imply that $F_{\max}$ is finite. Since all iterates are feasible, $F_{\max} - F(\boldsymbol{X}_k) \geq 0$. Also $\Delta_k \leq \delta_{\max}$, so Φ_k is nonnegative and bounded, and is therefore integrable. It is $\mathcal{F}_{k-1}$-measurable because $\boldsymbol{X}_k$ and Δ_k are known before iteration k.

We use the sufficient parameter condition stated in the main theorem,

$$(1-p_f)(\beta^2-\theta^2) < p_f(\gamma-1)(1-\theta^2). \tag{32}$$

Choose $\nu = (\beta^2-\theta^2)/(\beta^2-\theta^2+\gamma-1)$. Since $\beta > 1$, $0 < \theta < 1$, and $\gamma > 1$, this choice lies in $(0,1)$, and satisfies $\nu(\gamma-1) = (1-\nu)(\beta^2-\theta^2)$. The resulting drift constant is

$$M := p_f(1-\nu)(1-\theta^2) - \nu(1-p_f) > 0. \tag{33}$$

Lemma B.12 (Expected potential decrease)**.** *Suppose Assumption 3.4 and equation 32 hold. With the choices above, for every $k \geq 0$,*

$$\mathbb{E}\big[(\Phi_{k+1}-\Phi_k)\mathbf{1}_{\{T_\varepsilon > k\}} \,\big|\, \mathcal{F}_{k-1}\big] \leq -M\Delta_k^2 \mathbf{1}_{\{T_\varepsilon > k\}}. \tag{34}$$

Proof. Since $\{T_\varepsilon > k\} \in \mathcal{F}_{k-1}$, it suffices to prove

$$\mathbb{E}[\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1}] \leq -M\Delta_k^2.$$

We first condition on $\mathcal{F}_{k-1/2}$.

If $I_k^{\text{poll}} = 0$, then $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k$ and $\Delta_{k+1} = \theta\Delta_k$. Hence

$$\Phi_{k+1} - \Phi_k = -(1-\nu)(1-\theta^2)\Delta_k^2 \leq -M\Delta_k^2,$$

because $M \leq (1-\nu)(1-\theta^2)$.

It remains to consider the case $I_k^{\text{poll}} = 1$. On the event J_k, if iteration k is successful, then Lemma B.4 gives

$$F(\boldsymbol{X}_{k+1}) - F(\boldsymbol{X}_k) \geq (\gamma-1)\varepsilon_f\Delta_k^2.$$

Moreover,

$$\Delta_{k+1}^2 - \Delta_k^2 \leq (\beta^2-1)\Delta_k^2.$$

Therefore,

$$\begin{aligned}\Phi_{k+1} - \Phi_k &= \frac{\nu}{\varepsilon_f}\big[F(\boldsymbol{X}_k) - F(\boldsymbol{X}_{k+1})\big] + (1-\nu)(\Delta_{k+1}^2 - \Delta_k^2) \\ &\leq -\nu(\gamma-1)\Delta_k^2 + (1-\nu)(\beta^2-1)\Delta_k^2 \\ &\leq -(1-\nu)(1-\theta^2)\Delta_k^2.\end{aligned}$$

The last line follows from $\nu(\gamma-1) = (1-\nu)(\beta^2-\theta^2)$. If iteration k is unsuccessful, then $\boldsymbol{X}_{k+1} = \boldsymbol{X}_k$ and $\Delta_{k+1} = \theta\Delta_k$, and hence

$$\Phi_{k+1} - \Phi_k = -(1-\nu)(1-\theta^2)\Delta_k^2.$$

Thus, on J_k,

$$\Phi_{k+1} - \Phi_k \leq -(1-\nu)(1-\theta^2)\Delta_k^2.$$

Using

$$\mathbb{P}(J_k \mid \mathcal{F}_{k-1/2}) \geq p_f,$$

we obtain

$$\mathbb{E}\big[(\Phi_{k+1}-\Phi_k)\mathbf{1}_{J_k} \mid \mathcal{F}_{k-1/2}\big] \leq -p_f(1-\nu)(1-\theta^2)\Delta_k^2. \tag{35}$$

We next bound the contribution from J_k^c. If iteration k is unsuccessful, then

$$\Phi_{k+1} - \Phi_k = -(1-\nu)(1-\theta^2)\Delta_k^2 \leq 0.$$

If iteration k is successful, let $\widetilde{\boldsymbol{D}}_k$ be the accepted direction. Since

$$\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} \geq \gamma\varepsilon_f\Delta_k^2,$$

we have

$$\begin{aligned} -A_{k,\widetilde{\boldsymbol{D}}_k} &= F(\boldsymbol{X}_k) - F(\boldsymbol{X}_{k+1}) \\ &\le |\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - A_{k,\widetilde{\boldsymbol{D}}_k}| - \widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} \\ &\le |\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - A_{k,\widetilde{\boldsymbol{D}}_k}| - \gamma\varepsilon_f\Delta_k^2. \end{aligned}$$

Consequently,

$$\begin{aligned} \Phi_{k+1} - \Phi_k &\le \frac{\nu}{\varepsilon_f}|\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - A_{k,\widetilde{\boldsymbol{D}}_k}| - \nu\gamma\Delta_k^2 + (1-\nu)(\beta^2-1)\Delta_k^2 \\ &\le \frac{\nu}{\varepsilon_f}|\widetilde{A}_{k,\widetilde{\boldsymbol{D}}_k} - A_{k,\widetilde{\boldsymbol{D}}_k}|. \end{aligned}$$

The same error upper bound holds on an unsuccessful iteration because its potential change is non-positive. Therefore,

$$(\Phi_{k+1} - \Phi_k)\mathbf{1}_{J_k^c} \le \frac{\nu}{\varepsilon_f}\mathbf{1}_{J_k^c}\max_{\tilde{\boldsymbol{d}}\in\mathfrak{D}_k}|\widetilde{A}_{k,\tilde{\boldsymbol{d}}} - A_{k,\tilde{\boldsymbol{d}}}|.$$

By the conditional tail bound in Assumption 3.4,

$$\frac{\nu}{\varepsilon_f}\mathbb{E}\left[\mathbf{1}_{J_k^c}\max_{\boldsymbol{d}\in\mathfrak{D}_k}|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \,\middle|\, \mathcal{F}_{k-1/2}\right] \le \nu(1-p_f)\Delta_k^2.$$

Combining the estimates on J_k and J_k^c, we obtain, on $\{I_k^{\mathrm{poll}} = 1\}$,

$$\begin{aligned} \mathbb{E}\big[\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1/2}\big] &\le -p_f(1-\nu)(1-\theta^2)\Delta_k^2 + \nu(1-p_f)\Delta_k^2 \\ &= -M\Delta_k^2. \end{aligned}$$

Thus, both on $\{I_k^{\mathrm{poll}} = 0\}$ and on $\{I_k^{\mathrm{poll}} = 1\}$,

$$\mathbb{E}\big[\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1/2}\big] \le -M\Delta_k^2.$$

Taking conditional expectation with respect to $\mathcal{F}_{k-1}$ and using the tower property gives

$$\begin{aligned} \mathbb{E}[\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1}] &= \mathbb{E}\big[\mathbb{E}[\Phi_{k+1} - \Phi_k \mid \mathcal{F}_{k-1/2}] \,\big|\, \mathcal{F}_{k-1}\big] \\ &\le \mathbb{E}[-M\Delta_k^2 \mid \mathcal{F}_{k-1}] = -M\Delta_k^2, \end{aligned}$$

Multiplying both sides by the $\mathcal{F}_{k-1}$-measurable indicator $\mathbf{1}_{\{T_\varepsilon > k\}}$ yields

$$\mathbb{E}\big[(\Phi_{k+1} - \Phi_k)\mathbf{1}_{\{T_\varepsilon>k\}} \mid \mathcal{F}_{k-1}\big] \le -M\Delta_k^2\mathbf{1}_{\{T_\varepsilon>k\}}.$$

The proof is complete. □

We have now verified all hypotheses of Theorem B.5. The probability condition follows from Lemma B.6, the stepsize comparison follows from Lemma B.11, and the potential decrease follows from Lemma B.12. Before completing the main theorem, we give the full renewal–reward proof.

B.6 Proof of the renewal–reward bound

Proof of Theorem B.5. Set $a := \log\beta > 0$,$b := \log\theta < 0$. We first bound the accumulated reward before the stopping time. For $k \ge 0$, define

$$R_k := \Phi_{k\wedge T_\varepsilon} + M\sum_{j=0}^{k\wedge T_\varepsilon - 1}\Delta_j^2.$$

We claim that $\{R_k\}_{k\ge0}$ is a nonnegative supermartingale with respect to the shifted filtration $\{\mathcal{F}_{k-1}\}_{k\ge0}$. Indeed, since T_ε is a stopping time with respect to this filtration, we have

$$\mathbb{E}[R_{k+1}\mathbf{1}_{\{T_\varepsilon\le k\}} \mid \mathcal{F}_{k-1}] = R_k\mathbf{1}_{\{T_\varepsilon\le k\}},$$

whereas, on $\{T_\varepsilon > k\}$, the drift assumption gives

$$\mathbb{E}[R_{k+1}\mathbf{1}_{\{T_\varepsilon>k\}} \mid \mathcal{F}_{k-1}] = \mathbb{E}\left[\left(\Phi_{k+1} + M\sum_{j=0}^{k}\Delta_j^2\right)\mathbf{1}_{\{T_\varepsilon>k\}} \mid \mathcal{F}_{k-1}\right]$$
$$\le \left(\Phi_k + M\sum_{j=0}^{k-1}\Delta_j^2\right)\mathbf{1}_{\{T_\varepsilon>k\}}.$$

Combining the two displays yields $\mathbb{E}[R_{k+1} \mid \mathcal{F}_{k-1}] \le R_k$.

Thus $\{R_k\}$ is a nonnegative supermartingale. Therefore,

$$M\mathbb{E}\left[\sum_{j=0}^{k\wedge T_\varepsilon-1}\Delta_j^2\right] \le \mathbb{E}[R_k] \le \Phi_0.$$

Letting $k \to \infty$ and using the monotone convergence theorem gives

$$M\mathbb{E}\left[\sum_{j=0}^{T_\varepsilon-1}\Delta_j^2\right] \le \Phi_0. \tag{36}$$

We next introduce the renewal process. Define the auxiliary process $\mathsf{Z}_0 := 0$, $\mathsf{Z}_{k+1} := \min\{\mathsf{Z}_k + W_k, 0\}$. Since $\Delta_0 \ge \delta_\varepsilon$, the stepsize condition implies, by induction, that on $\{T_\varepsilon > k\}$, $\Delta_k \ge \delta_\varepsilon e^{\mathsf{Z}_k}$. Consequently, whenever $\mathsf{Z}_k = 0$ and $T_\varepsilon > k$, one has $\Delta_k \ge \delta_\varepsilon$.

Define the renewal times

$$\mathsf{A}_0 := 0, \qquad \mathsf{A}_{n+1} := \inf\{k > \mathsf{A}_n : \mathsf{Z}_k = 0\}, \qquad n \ge 0,$$

and the corresponding counting process

$$N(T) := \max\{n \ge 0 : \mathsf{A}_n \le T\},$$

with $N(T) = -1$ if the set is empty and $N(\infty) = \infty$.

Since every renewal counted by $N(T_\varepsilon - 1)$ corresponds to an iteration $k < T_\varepsilon$ with $\Delta_k \ge \delta_\varepsilon$, we obtain

$$\sum_{j=0}^{T_\varepsilon-1}\Delta_j^2 \ge \big(N(T_\varepsilon - 1) + 1\big)\delta_\varepsilon^2.$$

Combining this with the reward bound gives

$$\mathbb{E}[N(T_\varepsilon - 1) + 1] \le \frac{\Phi_0}{M\delta_\varepsilon^2}.$$

Define the interarrival times

$$\tau_n := \mathsf{A}_{n+1} - \mathsf{A}_n, \qquad n \ge 0.$$

We claim that, for every $n \ge 0$,

$$\mathbb{E}[\tau_n \mid \mathcal{F}_{\mathsf{A}_n-1}] \le \frac{a}{pa + (1-p)b},$$

where $\mathcal{F}_{\mathsf{A}_n-1}$ denotes the sigma-algebra at the beginning of iteration A_n. If $W_{\mathsf{A}_n} = a$, then $\mathsf{Z}_{\mathsf{A}_n+1} = 0$, and hence $\tau_n = 1$. If $W_{\mathsf{A}_n} = b$, then $\mathsf{Z}_{\mathsf{A}_n+1} = b < 0$. Conditional on this event, define

$$S_0 := b, \qquad S_m := b + \sum_{\ell=1}^{m} W_{\mathsf{A}_n+\ell}, \qquad m \ge 1,$$

and let $\tau := \inf\{m \ge 0 : S_m \ge 0\}$. Then, on $\{W_{\mathsf{A}_n} = b\}$, the remaining time to renewal is τ, and therefore $\tau_n = 1 + \tau$.

We now bound the conditional expectation of τ. After time A_n, the process

$$S_m - [pa + (1-p)b]m, \qquad m \geq 0,$$

is a submartingale with respect to the shifted filtration $\{\mathcal{F}_{\mathsf{A}_n+m}\}_{m\geq 0}$. Hence, for each $m \geq 1$, optional sampling at the bounded stopping time $\tau \wedge m$ gives

$$\mathbb{E}[S_{\tau\wedge m} - [pa + (1-p)b](\tau \wedge m) \mid \mathcal{F}_{\mathsf{A}_n}] \geq S_0 = b.$$

Thus

$$[pa + (1-p)b]\,\mathbb{E}[\tau \wedge m \mid \mathcal{F}_{\mathsf{A}_n}] \leq \mathbb{E}[S_{\tau\wedge m} \mid \mathcal{F}_{\mathsf{A}_n}] - b.$$

Before the hitting time τ, the walk is negative, while at the hitting time the overshoot above the nonnegative half-line is at most a. Hence

$$\mathbb{E}[\tau \wedge m \mid \mathcal{F}_{\mathsf{A}_n}] \leq \frac{a-b}{pa + (1-p)b}.$$

Letting $m \to \infty$ and using the monotone convergence theorem yields

$$\mathbb{E}[\tau \mid \mathcal{F}_{\mathsf{A}_n}] \leq \frac{a-b}{pa + (1-p)b}.$$

Let $q_n := \mathbb{P}(W_{\mathsf{A}_n} = a \mid \mathcal{F}_{\mathsf{A}_n-1})$. Since $\{\mathsf{A}_n = k\} \in \mathcal{F}_{k-1}$, the standing conditional probability assumption gives

$$q_n = \sum_{k\geq 0} \mathbf{1}_{\{\mathsf{A}_n=k\}} \mathbb{P}(W_k = a \mid \mathcal{F}_{k-1}) \geq p.$$

Therefore,

$$\begin{aligned}
\mathbb{E}[\tau_n \mid \mathcal{F}_{\mathsf{A}_n-1}] &\leq q_n \cdot 1 + (1-q_n)\left(1 + \frac{a-b}{pa + (1-p)b}\right) \\
&= 1 + \frac{(1-q_n)(a-b)}{pa + (1-p)b} \\
&\leq 1 + \frac{(1-p)(a-b)}{pa + (1-p)b} \\
&= \frac{a}{pa + (1-p)b}.
\end{aligned}$$

This proves the claimed interarrival bound.

We now sum these interarrival bounds up to the random number of renewals. Since

$$T_\varepsilon - 1 \leq \mathsf{A}_{N(T_\varepsilon-1)+1},$$

we have

$$\begin{aligned}
\mathbb{E}[T_\varepsilon - 1] &\leq \mathbb{E}\left[\sum_{n=0}^{N(T_\varepsilon-1)} \tau_n\right] \\
&= \sum_{n=0}^{\infty} \mathbb{E}\left[\tau_n \mathbf{1}_{\{N(T_\varepsilon-1)+1>n\}}\right] \\
&= \sum_{n=0}^{\infty} \mathbb{E}\left[\mathbf{1}_{\{N(T_\varepsilon-1)+1>n\}} \mathbb{E}[\tau_n \mid \mathcal{F}_{\mathsf{A}_n-1}]\right] \\
&\leq \frac{a}{pa + (1-p)b} \sum_{n=0}^{\infty} \mathbb{P}(N(T_\varepsilon - 1) + 1 > n) \\
&= \frac{a}{pa + (1-p)b} \mathbb{E}[N(T_\varepsilon - 1) + 1] \\
&\leq \frac{a}{pa + (1-p)b} \frac{\Phi_0}{M\delta_\varepsilon^2}.
\end{aligned}$$

Finally, substituting $a = \log \beta$ and $p = p_d p_f$, gives

$$\mathbb{E}[T_\varepsilon] \le \frac{\log \beta}{p_d p_f \log \beta + (1 - p_d p_f) \log \theta} \frac{\Phi_0}{M \delta_\varepsilon^2} + 1.$$

This completes the proof.

□

B.7 Completion of the iteration-complexity theorem

Proof of Theorem 3.5. Lemmas B.6, B.11, and B.12 verify the three conditions of Theorem B.5 with $p = p_d p_f$.

Moreover, $\Delta_0 = \delta_0 \ge \delta_\varepsilon$. Therefore,

$$\mathbb{E}[T_\varepsilon] \le \frac{\log \beta}{p_d p_f \log \beta + (1 - p_d p_f) \log \theta} \frac{\Phi_0}{M \delta_\varepsilon^2} + 1 = O(\varepsilon^{-2}).$$

Since $\varepsilon = \epsilon / D^*$, this is $O(\epsilon^{-2})$. In particular, $T_\varepsilon < \infty$ almost surely.

By definition, $\pi(\boldsymbol{X}_{T_\varepsilon}) \le \varepsilon$. The three approximation guarantees now follow directly from Theorem A.3 with $\varepsilon = \epsilon / D^*$.

□

C Constructions and estimator accuracy

This section gives the constructions used in Section 4 and verifies Assumption 3.4. We first describe the randomized polling set in a nondegenerate state. We then analyze the original-value estimators separately for the three non-oblivious objectives.

C.1 Randomized polling in a nondegenerate state

We use the algebraic generators in Proposition 3.3. For $p_d = 1$, we simply use the full normalized generator set, which gives the deterministic directional-quality guarantee.

We now give a randomized implementation for $0 < p_d < 1$. On a nondegenerate iteration, let R_Q be a right inverse of $Q(\boldsymbol{X}_k, \Delta_k)^\top$, and let the columns of Z_Q form a basis of $\mathrm{Null}(Q(\boldsymbol{X}_k, \Delta_k)^\top)$, as in Proposition 3.3. The paired generators $\pm B(\boldsymbol{X}_k, \Delta_k) Z_Q$ span the lineality space

$$S_k := \mathrm{Range}\big(B(\boldsymbol{X}_k, \Delta_k) Z_Q\big),$$

where $\mathrm{Range}(\cdot)$ denotes the column space. The remaining pointed component is $T(\boldsymbol{X}_k, \Delta_k) \cap S_k^\perp$. To obtain generators for this component, we orthogonally project the columns of $-B(\boldsymbol{X}_k, \Delta_k) R_Q$ onto $S_k^\perp$, discard zero projections, and normalize the remaining vectors. Let V_k^c denote the resulting set. Then V_k^c positively generates $T(\boldsymbol{X}_k, \Delta_k) \cap S_k^\perp$.

Define

$$r_s := \left\lfloor \log_2 \left(\frac{1}{1 - p_d} \right) \right\rfloor + 1. \tag{37}$$

If $S_k \neq \{\mathbf{0}\}$, sample r_s independent unit directions $\boldsymbol{s}_{k,1}, \dots, \boldsymbol{s}_{k,r_s}$ uniformly from the unit sphere of S_k.

If $V_k^c \neq \emptyset$, select a subset $\widehat{V}_k^c \subseteq V_k^c$ uniformly without replacement with $|\widehat{V}_k^c| = \lceil p_d |V_k^c| \rceil$. If $V_k^c = \emptyset$, set $\widehat{V}_k^c = \emptyset$. The randomized polling set is

$$\mathfrak{D}_k := \{\boldsymbol{s}_{k,1}, \dots, \boldsymbol{s}_{k,r_s}\} \cup \widehat{V}_k^c, \tag{38}$$

where the first set is omitted when $S_k = \{\mathbf{0}\}$.

The result follows by applying the subspace-polling argument of Gratton et al. (2019, Proposition 7.1 and Appendix B) to the algebraic generators in Proposition 3.3, with the sampling sizes above chosen to attain the prescribed probability p_d.

Proposition C.1 (Subspace polling with prescribed reliability)**.** *Let $0 < p_d < 1$, and construct $\mathfrak{D}_k$ as above on every nondegenerate iteration with $T(\boldsymbol{X}_k, \Delta_k) \neq \{\mathbf{0}\}$. Then there exists a constant $\kappa \in (0,1)$, independent of k, such that*

$$\mathbb{P}(Z_k \mid \mathcal{F}_{k-1}) \geq p_d.$$

Moreover, $|\mathfrak{D}_k| = O(r)$. For $p_d = 1$, the full normalized generator set contains at most $2r$ directions and gives the corresponding deterministic guarantee.

For all finite-difference constructions below, we choose $0 < \mu \leq \min\left\{1, \frac{\bar{\mu}}{2e}\right\}$. This ensures that all transformed query points remain in $\mathcal{U}$. Let $R := \max_{\boldsymbol{x}\in\mathcal{P}} \|\boldsymbol{x}\|$, which is finite by compactness.

We next record two estimates used throughout the value-oracle constructions. First, let $h : [0,1] \to \mathbb{R}$ be K-Lipschitz, and let $z_j = (j - \frac{1}{2})/M$ be the midpoint of the j-th interval. Then

$$\begin{aligned}\left|\int_0^1 h(z)\,dz - \frac{1}{M}\sum_{j=1}^{M} h(z_j)\right| &\leq \sum_{j=1}^{M}\int_{(j-1)/M}^{j/M} |h(z) - h(z_j)|\,dz \\ &\leq K\sum_{j=1}^{M}\int_{(j-1)/M}^{j/M} |z - z_j|\,dz = \frac{K}{4M}.\end{aligned}$$

Thus a Lipschitz bound on the corresponding increment integrand gives the required deterministic quadrature accuracy directly.

Second, for any admissible query points $\boldsymbol{x}, \boldsymbol{y}$, write $\widehat{f} = f + \bar{f}$. Since $\mathbb{E}_\xi[\bar{f}(\boldsymbol{x},\xi)] = \mathbb{E}_\xi[\bar{f}(\boldsymbol{y},\xi)] = 0$, the cross term vanishes and

$$\begin{aligned}\mathbb{E}_\xi|\widehat{f}(\boldsymbol{x},\xi) - \widehat{f}(\boldsymbol{y},\xi)|^2 &= |f(\boldsymbol{x}) - f(\boldsymbol{y})|^2 + \mathbb{E}_\xi|\bar{f}(\boldsymbol{x},\xi) - \bar{f}(\boldsymbol{y},\xi)|^2 \\ &\leq (L^2 + D_2)\|\boldsymbol{x} - \boldsymbol{y}\|^2. \end{aligned} \tag{39}$$

Here the last inequality follows from the L-Lipschitz continuity of f and the paired-noise condition equation 3. This estimate is the basic second-moment bound for all paired stochastic differences below.

We now show how a per-direction mean-square error bound verifies both conditions involving J_k in Assumption 3.4. Condition on $\mathcal{F}_{k-1/2}$, so that $\mathfrak{D}_k$, $\boldsymbol{X}_k$, and Δ_k are fixed. For a nonempty poll, suppose that every $\boldsymbol{d} \in \mathfrak{D}_k$ satisfies

$$\mathbb{E}\Big[|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \,\Big|\, \mathcal{F}_{k-1/2}\Big] \leq \frac{1-p_f}{|\mathfrak{D}_k|}\varepsilon_f^2\Delta_k^4. \tag{40}$$

By the definition of J_k,

$$J_k^c = \left\{\max_{\boldsymbol{d}\in\mathfrak{D}_k} |\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| > \varepsilon_f\Delta_k^2\right\}.$$

Applying Markov's inequality to the squared errors and then taking a union bound gives

$$\begin{aligned}\mathbb{P}(J_k^c \mid \mathcal{F}_{k-1/2}) &\leq \sum_{\boldsymbol{d}\in\mathfrak{D}_k} \mathbb{P}\Big(|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| > \varepsilon_f\Delta_k^2 \,\Big|\, \mathcal{F}_{k-1/2}\Big) \\ &\leq \sum_{\boldsymbol{d}\in\mathfrak{D}_k} \frac{\mathbb{E}[|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \mid \mathcal{F}_{k-1/2}]}{\varepsilon_f^2\Delta_k^4} \\ &\leq 1 - p_f.\end{aligned}$$

Hence $\mathbb{P}(J_k \mid \mathcal{F}_{k-1/2}) \geq p_f$. The same mean-square bound also gives the conditional tail estimate required in Assumption 3.4. By conditional Cauchy–Schwarz,

$$\begin{aligned}&\mathbb{E}\left[\mathbf{1}_{J_k^c}\max_{\boldsymbol{d}\in\mathfrak{D}_k}|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \,\middle|\, \mathcal{F}_{k-1/2}\right] \\ &\quad\leq \mathbb{P}(J_k^c \mid \mathcal{F}_{k-1/2})^{1/2}\left(\mathbb{E}\left[\max_{\boldsymbol{d}\in\mathfrak{D}_k}|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \,\middle|\, \mathcal{F}_{k-1/2}\right]\right)^{1/2} \\ &\quad\leq (1-p_f)^{1/2}\left(\sum_{\boldsymbol{d}\in\mathfrak{D}_k}\mathbb{E}\Big[|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \,\Big|\, \mathcal{F}_{k-1/2}\Big]\right)^{1/2} \\ &\quad\leq (1-p_f)\varepsilon_f\Delta_k^2.\end{aligned}$$

Therefore, for each of the three objectives, it remains only to verify the per-direction bound equation 40 and the corresponding accuracy bound for the reduced-gradient estimator.

C.2 MONOTONE OBJECTIVE

For the monotone objective,

$$F(\boldsymbol{x}) = \int_0^1 \frac{e^{z-1}}{z} f(z\boldsymbol{x})\, dz, \qquad \nabla F(\boldsymbol{x}) = \int_0^1 e^{z-1} \nabla f(z\boldsymbol{x})\, dz.$$

Exact directional increments. Fix $\boldsymbol{d} \in \mathfrak{D}_k$. Since $\boldsymbol{X}_k$ and $\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}$ are feasible, convexity gives $\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} \in \mathcal{P}$ for every $t \in [0,1]$. By the fundamental theorem of calculus,

$$\frac{e^{z-1}}{z}\big[f(z(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d})) - f(z\boldsymbol{X}_k)\big] = e^{z-1} \int_0^1 \big\langle \nabla f\big(z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big), \Delta_k W\boldsymbol{d} \big\rangle\, dt.$$

Thus the apparent singularity at $z = 0$ cancels.

We next bound the variation of this integrand with respect to z. For $z, z' \in [0,1]$ and $t \in [0,1]$, adding and subtracting the corresponding term with coefficient $e^{z'-1}$ gives

$$\begin{aligned}
&\Big| e^{z-1} \big\langle \nabla f\big(z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big), \Delta_k W\boldsymbol{d} \big\rangle - e^{z'-1} \big\langle \nabla f\big(z'(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big), \Delta_k W\boldsymbol{d} \big\rangle \Big| \\
&\quad \le |e^{z-1} - e^{z'-1}|\, \|\nabla f\big(z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big)\| \Delta_k \|W\boldsymbol{d}\| \\
&\qquad + e^{z'-1} \big\|\nabla f\big(z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big) - \nabla f\big(z'(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big)\big\|\, \Delta_k \|W\boldsymbol{d}\|.
\end{aligned}$$

All polling directions are normalized, so $\|W\boldsymbol{d}\| = \|\boldsymbol{d}\| = 1$. Moreover, $|e^{z-1} - e^{z'-1}| \le |z - z'|$, $e^{z'-1} \le 1$, and $\|\nabla f\| \le L$. By L_1-smoothness,

$$\big\|\nabla f\big(z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big) - \nabla f\big(z'(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})\big)\big\| \le L_1 |z - z'|\, \|\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d}\|.$$

Since $\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} \in \mathcal{P}$, we have $\|\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d}\| \le R$. Hence the preceding difference is bounded by

$$(L + L_1 R)\Delta_k |z - z'|.$$

Integrating over $t \in [0,1]$ shows that the increment integrand is $(L + L_1 R)\Delta_k$-Lipschitz in z. The midpoint estimate established above therefore gives

$$|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \le \frac{(L + L_1 R)\Delta_k}{4N}.$$

Hence J_k holds deterministically whenever

$$N \ge \frac{L + L_1 R}{4\varepsilon_f \Delta_k}. \tag{41}$$

Thus $N = O(\Delta_k^{-1})$, and a complete exact poll uses $O(|\mathfrak{D}_k|\Delta_k^{-1})$ original-value calls.

Exact reduced-gradient estimator. We next bound the error of the exact fallback estimator. For $\ell = 1, \dots, r$,

$$(\boldsymbol{G}_k)_\ell = \frac{1}{N} \sum_{j=1}^{N} \frac{e^{z_j - 1}}{2\mu z_j} \left[f\big(z_j(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)\big) - f\big(z_j(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)\big) \right].$$

Recall that

$$\boldsymbol{g}(\boldsymbol{X}_k) = W^\top \nabla F(\boldsymbol{X}_k) = \int_0^1 e^{z-1} W^\top \nabla f(z\boldsymbol{X}_k)\, dz.$$

We control the quadrature and finite-difference errors separately.

First consider the quadrature error. For $z, z' \in [0,1]$,

$$\begin{aligned}
&\left\| e^{z-1} W^\top \nabla f(z\boldsymbol{X}_k) - e^{z'-1} W^\top \nabla f(z'\boldsymbol{X}_k) \right\| \\
&\quad \le |e^{z-1} - e^{z'-1}| \, \|W^\top \nabla f(z\boldsymbol{X}_k)\| \\
&\qquad + e^{z'-1} \|W^\top (\nabla f(z\boldsymbol{X}_k) - \nabla f(z'\boldsymbol{X}_k))\| \\
&\quad \le L|z - z'| + L_1 \|\boldsymbol{X}_k\| |z - z'| \\
&\quad \le (L + L_1 R)|z - z'|.
\end{aligned}$$

Here we used $|e^{z-1} - e^{z'-1}| \le |z - z'|$, $e^{z'-1} \le 1$, $\|W^\top\| = 1$, $\|\nabla f\| \le L$, L_1-smoothness, and $\|\boldsymbol{X}_k\| \le R$. Thus the vector integrand is $(L + L_1 R)$-Lipschitz. Applying the midpoint estimate,

$$\left\| \frac{1}{N} \sum_{j=1}^{N} e^{z_j - 1} W^\top \nabla f(z_j \boldsymbol{X}_k) - \boldsymbol{g}(\boldsymbol{X}_k) \right\| \le \frac{L + L_1 R}{4N}.$$

We next bound the finite-difference error at each midpoint. For a fixed z and coordinate ℓ, the fundamental theorem of calculus gives

$$\frac{f(z(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)) - f(z(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell))}{2\mu z} = \frac{1}{2\mu z} \int_{-z\mu}^{z\mu} \langle \nabla f(z\boldsymbol{X}_k + s\boldsymbol{w}_\ell), \boldsymbol{w}_\ell \rangle \, ds.$$

Subtracting $\langle \nabla f(z\boldsymbol{X}_k), \boldsymbol{w}_\ell \rangle$ and using L_1-smoothness and $\|\boldsymbol{w}_\ell\| = 1$ gives

$$\begin{aligned}
&\left| \frac{f(z(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)) - f(z(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell))}{2\mu z} - \langle \nabla f(z\boldsymbol{X}_k), \boldsymbol{w}_\ell \rangle \right| \\
&\quad \le \frac{1}{2\mu z} \int_{-z\mu}^{z\mu} L_1 |s| \, ds = \frac{L_1 z \mu}{2}.
\end{aligned}$$

After multiplication by e^{z-1}, this is at most $L_1\mu/2$. Therefore, at each midpoint, the Euclidean error over the r reduced coordinates is at most $L_1\mu\sqrt{r}/2$, and averaging over the midpoint nodes does not increase this bound.

Combining the finite-difference and quadrature errors yields

$$\begin{aligned}
\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| &\le \left\| \boldsymbol{G}_k - \frac{1}{N} \sum_{j=1}^{N} e^{z_j - 1} W^\top \nabla f(z_j \boldsymbol{X}_k) \right\| \\
&\quad + \left\| \frac{1}{N} \sum_{j=1}^{N} e^{z_j - 1} W^\top \nabla f(z_j \boldsymbol{X}_k) - \boldsymbol{g}(\boldsymbol{X}_k) \right\| \\
&\le \frac{L_1 \mu \sqrt{r}}{2} + \frac{L + L_1 R}{4N}.
\end{aligned}$$

It is therefore sufficient to take

$$0 < \mu \le \min\left\{ 1, \frac{\bar{\mu}}{2e}, \frac{\varepsilon_g}{L_1\sqrt{r}} \right\}, \qquad N \ge \frac{L + L_1 R}{2\varepsilon_g}. \tag{42}$$

Hence $\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \varepsilon_g$, so H_k holds deterministically.

Finally, each of the r reduced coordinates requires two original function values at each of N midpoint nodes. Since $N = O(\varepsilon_g^{-1}) = O(\varepsilon^{-1})$, one exact fallback uses $O(r\varepsilon^{-1})$ original-value calls.

Stochastic directional increments. Condition on $\mathcal{F}_{k-1/2}$, so that $\boldsymbol{X}_k$, Δ_k, and $\mathfrak{D}_k$ are fixed. For $\boldsymbol{d} \in \mathfrak{D}_k$, recall that

$$\widetilde{A}_{k,\boldsymbol{d}} = \frac{1 - e^{-1}}{m} \sum_{i=1}^{m} \frac{\widehat{f}(Z_i(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d}), \xi_i) - \widehat{f}(Z_i \boldsymbol{X}_k, \xi_i)}{Z_i},$$

where Z_i has density $e^{z-1}/(1-e^{-1})$ on $(0,1)$.

We first verify conditional unbiasedness. Taking expectation over the oracle noise and then over Z_i,

$$\begin{aligned}\mathbb{E}[\widetilde{A}_{k,\boldsymbol{d}} \mid \mathcal{F}_{k-1/2}] &= \int_0^1 \frac{e^{z-1}}{z}\big[f(z(\boldsymbol{X}_k+\Delta_k W\boldsymbol{d}))-f(z\boldsymbol{X}_k)\big]\,dz \\ &= F(\boldsymbol{X}_k+\Delta_k W\boldsymbol{d})-F(\boldsymbol{X}_k)=A_{k,\boldsymbol{d}}.\end{aligned}$$

We next bound the mean-square error. Conditional on $Z_i = z$, the two paired query points are $z(\boldsymbol{X}_k+\Delta_k W\boldsymbol{d})$ and $z\boldsymbol{X}_k$, whose distance is $z\Delta_k$. Therefore equation 39 gives

$$\begin{aligned}&\mathbb{E}\left[\left|\frac{1-e^{-1}}{z}\big[\widehat{f}(z(\boldsymbol{X}_k+\Delta_k W\boldsymbol{d}),\xi_i)-\widehat{f}(z\boldsymbol{X}_k,\xi_i)\big]\right|^2\,\middle|\,\mathcal{F}_{k-1/2}, Z_i=z\right] \\ &\quad\le \frac{(1-e^{-1})^2}{z^2}(L^2+D_2)z^2\Delta_k^2 \\ &\quad= (1-e^{-1})^2(L^2+D_2)\Delta_k^2.\end{aligned}$$

The bound is independent of z, so it remains valid after averaging over Z_i. Since the m replications are independent and $\widetilde{A}_{k,\boldsymbol{d}}$ is conditionally unbiased,

$$\begin{aligned}\mathbb{E}\Big[|\widetilde{A}_{k,\boldsymbol{d}}-A_{k,\boldsymbol{d}}|^2\,\Big|\,\mathcal{F}_{k-1/2}\Big] &= \mathrm{Var}(\widetilde{A}_{k,\boldsymbol{d}} \mid \mathcal{F}_{k-1/2}) \\ &\le \frac{(1-e^{-1})^2(L^2+D_2)\Delta_k^2}{m}.\end{aligned}$$

Hence the common bound equation 40 is satisfied whenever

$$m \ge \frac{|\mathfrak{D}_k|(1-e^{-1})^2(L^2+D_2)}{(1-p_f)\varepsilon_f^2\Delta_k^2}. \tag{43}$$

The argument preceding this subsection then gives both $\mathbb{P}(J_k \mid \mathcal{F}_{k-1/2}) \ge p_f$ and the conditional tail bound required in Assumption 3.4. Since each replication uses two original-value calls and $m = O(|\mathfrak{D}_k|\Delta_k^{-2})$ samples are taken for each tested direction, a complete stochastic poll uses $O(|\mathfrak{D}_k|^2\Delta_k^{-2})$ calls.

Stochastic reduced-gradient estimator. For the stochastic constructions below, we take $p_d, p_f \in (0,1)$. We now consider the stochastic fallback estimator. Throughout the finite-difference constructions, μ is chosen small enough that all transformed query points lie in $\mathcal{U}$.

For coordinate ℓ,

$$(\boldsymbol{G}_k)_\ell = \frac{1-e^{-1}}{m}\sum_{i=1}^{m}\frac{\widehat{f}(Z_i(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell),\xi_{i,\ell})-\widehat{f}(Z_i(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell),\xi_{i,\ell})}{2\mu Z_i}.$$

Taking expectation over the oracle noise and Z_i gives

$$\begin{aligned}\mathbb{E}[(\boldsymbol{G}_k)_\ell \mid \mathcal{F}_{k-1}] &= \int_0^1 \frac{e^{z-1}}{2\mu z}\big[f(z(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell))-f(z(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell))\big]\,dz \\ &= \frac{F(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell)-F(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell)}{2\mu}.\end{aligned}$$

Thus the conditional mean of $\boldsymbol{G}_k$ is the centered finite-difference estimate of $\boldsymbol{g}(\boldsymbol{X}_k) = W^\top\nabla F(\boldsymbol{X}_k)$.

We first bound its deterministic bias. By the fundamental theorem of calculus,

$$\frac{F(\boldsymbol{X}_k+\mu\boldsymbol{w}_\ell)-F(\boldsymbol{X}_k-\mu\boldsymbol{w}_\ell)}{2\mu} = \frac{1}{2\mu}\int_{-\mu}^{\mu}\langle\nabla F(\boldsymbol{X}_k+s\boldsymbol{w}_\ell),\boldsymbol{w}_\ell\rangle\,ds.$$

Since F is L_g-smooth and $\|\boldsymbol{w}_\ell\|=1$,

$$\begin{aligned}&|\mathbb{E}[(\boldsymbol{G}_k)_\ell \mid \mathcal{F}_{k-1}]-\langle\nabla F(\boldsymbol{X}_k),\boldsymbol{w}_\ell\rangle| \\ &\quad\le \frac{1}{2\mu}\int_{-\mu}^{\mu}L_g|s|\,ds = \frac{L_g\mu}{2}.\end{aligned}$$

Summing over the r reduced coordinates gives

$$\|\mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}] - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \frac{L_g \mu \sqrt{r}}{2}.$$

With the choice of μ in equation 42 and $L_g = e^{-1}L_1$, this is at most $\varepsilon_g/2$.

It remains to control the sampling error around this conditional mean. Conditional on $Z_i = z$, the paired query points $z(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)$ and $z(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)$ are separated by $2z\mu$. Applying equation 39 and including the multiplier $(1 - e^{-1})/(2\mu z)$ gives

$$\begin{aligned}
&\mathbb{E}\left[\left|\frac{1-e^{-1}}{2\mu z}\big[\widehat{f}(z(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell), \xi_{i,\ell}) - \widehat{f}(z(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell), \xi_{i,\ell})\big]\right|^2 \,\middle|\, \mathcal{F}_{k-1}, Z_i = z\right] \\
&\quad \le \frac{(1-e^{-1})^2}{4\mu^2 z^2}(L^2 + D_2)(2z\mu)^2 \\
&\quad = (1-e^{-1})^2(L^2 + D_2).
\end{aligned}$$

Again the bound is independent of z. Averaging m independent replications in each coordinate and summing the coordinatewise variances therefore gives

$$\mathbb{E}\big[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \,\big|\, \mathcal{F}_{k-1}\big] \le \frac{r(1-e^{-1})^2(L^2 + D_2)}{m}.$$

Choose

$$m \ge \frac{4r(1-e^{-1})^2(L^2 + D_2)}{(1-p_d)\varepsilon_g^2}. \tag{44}$$

Then

$$\mathbb{E}\big[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \,\big|\, \mathcal{F}_{k-1}\big] \le \frac{1-p_d}{4}\varepsilon_g^2.$$

Applying Markov's inequality to the squared norm yields

$$\begin{aligned}
&\mathbb{P}\Big(\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\| > \frac{\varepsilon_g}{2} \,\Big|\, \mathcal{F}_{k-1}\Big) \\
&\quad \le \frac{4\,\mathbb{E}[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \mid \mathcal{F}_{k-1}]}{\varepsilon_g^2} \le 1 - p_d.
\end{aligned}$$

Hence, with conditional probability at least p_d, the sampling error is at most $\varepsilon_g/2$. On this event, the triangle inequality and the bias bound give

$$\begin{aligned}
\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| &\le \|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\| + \|\mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}] - \boldsymbol{g}(\boldsymbol{X}_k)\| \\
&\le \varepsilon_g.
\end{aligned}$$

This is precisely the event H_k, and therefore

$$\mathbb{P}(H_k \mid \mathcal{F}_{k-1}) \ge p_d.$$

Finally, each of the r reduced coordinates uses m paired replications, with two original-value calls per replication. Since $m = O(r\varepsilon_g^{-2}) = O(r\varepsilon^{-2})$, one stochastic fallback uses $O(r^2\varepsilon^{-2})$ original-value calls.

C.3 General nonmonotone objective

For the general nonmonotone objective,

$$F(\boldsymbol{x}) = \int_0^1 \frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big) - f(\underline{\boldsymbol{x}})}{4z(1 - z/2)^3}\, dz, \qquad \nabla F(\boldsymbol{x}) = \int_0^1 \frac{\nabla f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{x} - \underline{\boldsymbol{x}})\big)}{8(1 - z/2)^3}\, dz.$$

Exact directional increments. Fix $\boldsymbol{d} \in \mathfrak{D}_k$. Since $\boldsymbol{X}_k$ and $\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}$ are feasible, convexity gives

$$\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} \in \mathcal{P}, \qquad t \in [0,1].$$

Moreover, for $z \in [0,1]$,

$$\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}}) = \left(1 - \frac{z}{2}\right)\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d}) \in \mathcal{P},$$

so all points appearing below lie in the region where the regularity assumptions on f apply.

By the fundamental theorem of calculus,

$$\frac{f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}})\right) - f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\right)}{4z(1-z/2)^3}$$
$$= \frac{1}{8(1-z/2)^3}\int_0^1 \left\langle \nabla f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}})\right), \Delta_k W\boldsymbol{d}\right\rangle dt.$$

Thus the apparent singularity at $z = 0$ cancels.

We next bound the variation of the integrand with respect to z. For $z, z' \in [0,1]$ and $t \in [0,1]$, together with $\|\nabla f\| \le L$, L_1-smoothness, $\|W\boldsymbol{d}\| = 1$, and $\|\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}}\| \le D$, we obtain

$$\left| \frac{\left\langle \nabla f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}})\right), \Delta_k W\boldsymbol{d}\right\rangle}{8(1-z/2)^3} - \frac{\left\langle \nabla f\left(\underline{\boldsymbol{x}} + \frac{z'}{2}(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} - \underline{\boldsymbol{x}})\right), \Delta_k W\boldsymbol{d}\right\rangle}{8(1-z'/2)^3} \right|$$
$$\le \left(3L + \frac{L_1 D}{2}\right)\Delta_k |z - z'|.$$

Hence the increment integrand is $(3L + L_1 D/2)\Delta_k$-Lipschitz in z. The midpoint estimate therefore gives

$$|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \le \frac{(3L + L_1 D/2)\Delta_k}{4N}.$$

Hence J_k holds deterministically whenever

$$N \ge \frac{3L + L_1 D/2}{4\varepsilon_f \Delta_k}. \tag{45}$$

Thus $N = O(\Delta_k^{-1})$, and a complete exact poll uses $O(|\mathfrak{D}_k|\Delta_k^{-1})$ original-value calls.

Exact reduced-gradient estimator. For $\ell = 1, \dots, r$,

$$(\boldsymbol{G}_k)_\ell = \frac{1}{N}\sum_{j=1}^{N} \frac{f\left(\underline{\boldsymbol{x}} + \frac{z_j}{2}(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\right) - f\left(\underline{\boldsymbol{x}} + \frac{z_j}{2}(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\right)}{8\mu z_j (1 - z_j/2)^3}.$$

Recall that

$$\boldsymbol{g}(\boldsymbol{X}_k) = \int_0^1 \frac{W^\top \nabla f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\right)}{8(1-z/2)^3}\, dz.$$

As in the monotone case, we control the quadrature and finite-difference errors separately.

For the quadrature error, the scalar weight $1/[8(1-z/2)^3]$ is at most 1 and is 3-Lipschitz. Moreover,

$$\left\| \underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}) - \underline{\boldsymbol{x}} - \frac{z'}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}) \right\| \le \frac{D}{2}|z - z'|.$$

Using $\|W^\top\| = 1$, $\|\nabla f\| \le L$, and L_1-smoothness therefore gives

$$\left\| \frac{W^\top \nabla f\left(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\right)}{8(1-z/2)^3} - \frac{W^\top \nabla f\left(\underline{\boldsymbol{x}} + \frac{z'}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\right)}{8(1-z'/2)^3} \right\|$$
$$\le \left(3L + \frac{L_1 D}{2}\right)|z - z'|.$$

The midpoint quadrature error is thus bounded by

$$\left\| \frac{1}{N} \sum_{j=1}^{N} \frac{W^\top \nabla f\big(\underline{\boldsymbol{x}} + \frac{z_j}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\big)}{8(1 - z_j/2)^3} - \boldsymbol{g}(\boldsymbol{X}_k) \right\| \le \frac{3L + L_1 D/2}{4N}.$$

We next bound the finite-difference error at a fixed midpoint z. The two query points are centered at $\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})$ and displaced by $\pm z\mu \boldsymbol{w}_\ell/2$. Hence

$$\frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\big) - f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\big)}{\mu z}$$
$$= \frac{1}{\mu z} \int_{-z\mu/2}^{z\mu/2} \left\langle \nabla f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}) + s\boldsymbol{w}_\ell\Big), \boldsymbol{w}_\ell \right\rangle ds.$$

Subtracting the directional derivative at the center and using L_1-smoothness gives

$$\left| \frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\big) - f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell - \underline{\boldsymbol{x}})\big)}{\mu z} - \left\langle \nabla f\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\Big), \boldsymbol{w}_\ell \right\rangle \right|$$
$$\le \frac{1}{\mu z} \int_{-z\mu/2}^{z\mu/2} L_1 |s| \, ds = \frac{L_1 z \mu}{4}.$$

Multiplication by $1/[8(1 - z/2)^3] \le 1$ leaves an error at most $L_1\mu/4$ in each reduced coordinate. Hence the Euclidean finite-difference error is at most $L_1\mu\sqrt{r}/4$.

Combining the two errors yields

$$\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \frac{L_1 \mu \sqrt{r}}{4} + \frac{3L + L_1 D/2}{4N}.$$

It is therefore sufficient to take

$$0 < \mu \le \min\left\{1, \frac{\bar{\mu}}{2e}, \frac{2\varepsilon_g}{L_1\sqrt{r}}\right\}, \qquad N \ge \frac{3L + L_1 D/2}{2\varepsilon_g}. \tag{46}$$

Then $\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \varepsilon_g$, so H_k holds deterministically. Since $N = O(\varepsilon_g^{-1}) = O(\varepsilon^{-1})$, one exact fallback uses $O(r\varepsilon^{-1})$ original-value calls.

Stochastic directional increments. Condition on $\mathcal{F}_{k-1/2}$. For $\boldsymbol{d} \in \mathfrak{D}_k$,

$$\widetilde{A}_{k,\boldsymbol{d}} = \frac{3}{4m} \sum_{i=1}^{m} \frac{\widehat{f}\big(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d} - \underline{\boldsymbol{x}}), \xi_i\big) - \widehat{f}\big(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}), \xi_i\big)}{Z_i},$$

where Z_i has density $1/[3(1 - z/2)^3]$ on $(0, 1)$.

Taking expectation over the oracle noise and Z_i,

$$\mathbb{E}[\widetilde{A}_{k,\boldsymbol{d}} \mid \mathcal{F}_{k-1/2}] = \int_0^1 \frac{f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d} - \underline{\boldsymbol{x}})\big) - f\big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}})\big)}{4z(1 - z/2)^3} \, dz$$
$$= A_{k,\boldsymbol{d}}.$$

Thus the estimator is conditionally unbiased.

Hence equation 39 gives

$$\mathbb{E}\left[\left| \frac{3}{4z} \left[\widehat{f}\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d} - \underline{\boldsymbol{x}}), \xi_i\Big) - \widehat{f}\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \underline{\boldsymbol{x}}), \xi_i\Big) \right] \right|^2 \,\middle|\, \mathcal{F}_{k-1/2}, Z_i = z \right]$$
$$\le \frac{9}{16z^2}(L^2 + D_2)\frac{z^2\Delta_k^2}{4} = \frac{9}{64}(L^2 + D_2)\Delta_k^2.$$

The bound is independent of z, and hence also bounds the conditional second moment after averaging over Z_i. Since conditional variance is bounded by the corresponding second moment, each replication has conditional variance at most $9(L^2 + D_2)\Delta_k^2/64$.

The m replications are conditionally independent. Therefore

$$\mathbb{E}\Big[|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \,\Big|\, \mathcal{F}_{k-1/2}\Big] = \mathrm{Var}(\widetilde{A}_{k,\boldsymbol{d}} \mid \mathcal{F}_{k-1/2}) \le \frac{9(L^2+D_2)\Delta_k^2}{64m}.$$

Consequently, equation 40 holds whenever

$$m \ge \frac{9|\mathfrak{D}_k|(L^2+D_2)}{64(1-p_f)\varepsilon_f^2\Delta_k^2}. \tag{47}$$

The common argument then gives both reliability conditions involving J_k. Since $m = O(|\mathfrak{D}_k|\Delta_k^{-2})$ samples are used for each tested direction, a complete stochastic poll uses $O(|\mathfrak{D}_k|^2\Delta_k^{-2})$ original-value calls.

Stochastic reduced-gradient estimator. For coordinate ℓ,

$$(\boldsymbol{G}_k)_\ell = \frac{3}{4m}\sum_{i=1}^{m}\frac{\widehat{f}\big(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell}\big) - \widehat{f}\big(\underline{\boldsymbol{x}} + \frac{Z_i}{2}(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell}\big)}{2\mu Z_i}.$$

Taking expectation over the oracle noise and Z_i gives

$$\mathbb{E}[(\boldsymbol{G}_k)_\ell \mid \mathcal{F}_{k-1}] = \frac{F(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell) - F(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell)}{2\mu}.$$

Thus the conditional mean is the centered finite-difference estimate of the corresponding component of $\boldsymbol{g}(\boldsymbol{X}_k)$.

Since F is L_g-smooth,

$$|\mathbb{E}[(\boldsymbol{G}_k)_\ell \mid \mathcal{F}_{k-1}] - \langle \nabla F(\boldsymbol{X}_k), \boldsymbol{w}_\ell\rangle| \le \frac{1}{2\mu}\int_{-\mu}^{\mu} L_g|s|\,ds = \frac{L_g\mu}{2}.$$

Hence

$$\|\mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}] - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \frac{L_g\mu\sqrt{r}}{2}.$$

It remains to control the sampling error. Conditional on $Z_i = z$, the two transformed query points are separated by $z\mu$, while the multiplier is $3/(8\mu z)$. Hence

$$\mathbb{E}\left[\left|\frac{3}{8\mu z}\left[\widehat{f}\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k + \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell}\Big) - \widehat{f}\Big(\underline{\boldsymbol{x}} + \frac{z}{2}(\boldsymbol{X}_k - \mu\boldsymbol{w}_\ell - \underline{\boldsymbol{x}}), \xi_{i,\ell}\Big)\right]\right|^2 \,\middle|\, \mathcal{F}_{k-1}, Z_i = z\right] \le \frac{9}{64\mu^2 z^2}(L^2+D_2)z^2\mu^2 = \frac{9}{64}(L^2+D_2).$$

Therefore, after averaging m independent replications in each coordinate and summing over the r coordinates,

$$\mathbb{E}\big[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \,\big|\, \mathcal{F}_{k-1}\big] \le \frac{9r(L^2+D_2)}{64m}.$$

Choose

$$m \ge \frac{9r(L^2+D_2)}{16(1-p_d)\varepsilon_g^2}. \tag{48}$$

Then

$$\mathbb{E}\big[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \,\big|\, \mathcal{F}_{k-1}\big] \le \frac{1-p_d}{4}\varepsilon_g^2.$$

Markov's inequality therefore gives, with conditional probability at least p_d,

$$\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\| \le \frac{\varepsilon_g}{2}.$$

Together with the deterministic bias bound and the triangle inequality, this yields

$$\mathbb{P}(H_k \mid \mathcal{F}_{k-1}) \ge p_d.$$

Since $m = O(r\varepsilon^{-2})$ replications are used for each of the r reduced coordinates, one stochastic fallback uses $O(r^2\varepsilon^{-2})$ original-value calls.

C.4 DOWN-CLOSED NONMONOTONE OBJECTIVE

For the down-closed nonmonotone objective,

$$F(\boldsymbol{x}) = \int_0^1 \frac{e^{z-1}}{(1-e^{-1})z} \left[f(\mathbf{1} - e^{-z\boldsymbol{x}}) - f(\mathbf{0})\right] \, dz,$$

and

$$\nabla F(\boldsymbol{x}) = \int_0^1 \frac{e^{z-1}}{1-e^{-1}} \operatorname{Diag}(e^{-z\boldsymbol{x}}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{x}}) \, dz.$$

Exact directional increments. Fix $\boldsymbol{d} \in \mathfrak{D}_k$. Since $\boldsymbol{X}_k$ and $\boldsymbol{X}_k + \Delta_k W\boldsymbol{d}$ are feasible, $\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d} \in \mathcal{P} \subseteq [0,1]^n$ for every $t \in [0,1]$. By the fundamental theorem of calculus,

$$\begin{aligned}
&\frac{e^{z-1}}{(1-e^{-1})z} \left[f\Big(\mathbf{1} - e^{-z(\boldsymbol{X}_k + \Delta_k W\boldsymbol{d})}\Big) - f\big(\mathbf{1} - e^{-z\boldsymbol{X}_k}\big)\right] \\
&\quad = \frac{e^{z-1}}{1-e^{-1}} \int_0^1 \left\langle \operatorname{Diag}\Big(e^{-z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big) \nabla f\Big(\mathbf{1} - e^{-z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big), \Delta_k W\boldsymbol{d} \right\rangle dt.
\end{aligned}$$

Hence the factor $1/z$ causes no singularity at $z = 0$.

We next bound the variation of the integrand with respect to z. For $\boldsymbol{x} \in \mathcal{P} \subseteq [0,1]^n$, the diagonal entries of $\frac{e^{z-1}}{1-e^{-1}} \operatorname{Diag}(e^{-z\boldsymbol{x}})$ are $\frac{e^{z(1-x_i)-1}}{1-e^{-1}}$. Since $0 \le x_i \le 1$ and $z \in [0,1]$, both these entries and their derivatives with respect to z are bounded by $1/(1-e^{-1})$. Consequently, for $z, z' \in [0,1]$,

$$\left\| \frac{e^{z-1}}{1-e^{-1}} \operatorname{Diag}(e^{-z\boldsymbol{x}}) - \frac{e^{z'-1}}{1-e^{-1}} \operatorname{Diag}(e^{-z'\boldsymbol{x}}) \right\| \le \frac{|z-z'|}{1-e^{-1}}.$$

Moreover,

$$\left\| (\mathbf{1} - e^{-z\boldsymbol{x}}) - (\mathbf{1} - e^{-z'\boldsymbol{x}}) \right\| \le R|z-z'|.$$

Applying these two bounds with $\boldsymbol{x} = \boldsymbol{X}_k + t\Delta_k W\boldsymbol{d}$, and using $\|\nabla f\| \le L$, L_1-smoothness, and $\|W\boldsymbol{d}\| = 1$, gives

$$\begin{aligned}
&\left| \frac{e^{z-1}}{1-e^{-1}} \left\langle \operatorname{Diag}\Big(e^{-z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big) \nabla f\Big(\mathbf{1} - e^{-z(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big), \Delta_k W\boldsymbol{d} \right\rangle \right. \\
&\qquad \left. - \frac{e^{z'-1}}{1-e^{-1}} \left\langle \operatorname{Diag}\Big(e^{-z'(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big) \nabla f\Big(\mathbf{1} - e^{-z'(\boldsymbol{X}_k + t\Delta_k W\boldsymbol{d})}\Big), \Delta_k W\boldsymbol{d} \right\rangle \right| \\
&\qquad \le \frac{L + L_1 R}{1-e^{-1}} \Delta_k |z - z'|.
\end{aligned}$$

Thus, after integrating over $t \in [0,1]$, the increment integrand is $(L + L_1 R)\Delta_k/(1-e^{-1})$-Lipschitz in z. The midpoint estimate therefore gives

$$|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}| \le \frac{(L + L_1 R)\Delta_k}{4(1-e^{-1})N}.$$

Hence J_k holds deterministically whenever

$$N \ge \frac{L + L_1 R}{4(1-e^{-1})\varepsilon_f \Delta_k}. \tag{49}$$

Thus $N = O(\Delta_k^{-1})$, and a complete exact poll uses $O(|\mathfrak{D}_k|\Delta_k^{-1})$ original-value calls.

Exact reduced-gradient estimator. For $\ell = 1, \ldots, r$,

$$(\boldsymbol{G}_k)_\ell = \frac{1}{N} \sum_{j=1}^{N} \frac{e^{z_j - 1}}{2(1-e^{-1})\mu z_j} \left[f\Big(\mathbf{1} - e^{-z_j(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)}\Big) - f\Big(\mathbf{1} - e^{-z_j(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)}\Big)\right].$$

Recall that the vector integrand in $\boldsymbol{g}(\boldsymbol{X}_k)$ is

$$\frac{e^{z-1}}{1-e^{-1}} W^\top \operatorname{Diag}(e^{-z\boldsymbol{X}_k}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{X}_k}).$$

We first control the quadrature error. Since $\boldsymbol{X}_k \in \mathcal{P}$, $\|\boldsymbol{X}_k\| \le R$. As above, the matrix factor $e^{z-1}\operatorname{Diag}(e^{-z\boldsymbol{X}_k})/(1-e^{-1})$ has norm at most $1/(1-e^{-1})$ and is $1/(1-e^{-1})$-Lipschitz in z. In addition,

$$\|\mathbf{1} - e^{-z\boldsymbol{X}_k} - (\mathbf{1} - e^{-z'\boldsymbol{X}_k})\| \le R|z - z'|.$$

Using $\|W^\top\| = 1$, $\|\nabla f\| \le L$, and L_1-smoothness therefore gives

$$\left\| \frac{e^{z-1}}{1-e^{-1}} W^\top \operatorname{Diag}(e^{-z\boldsymbol{X}_k}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{X}_k}) - \frac{e^{z'-1}}{1-e^{-1}} W^\top \operatorname{Diag}(e^{-z'\boldsymbol{X}_k}) \nabla f(\mathbf{1} - e^{-z'\boldsymbol{X}_k}) \right\|$$
$$\le \frac{L + L_1 R}{1-e^{-1}} |z - z'|.$$

Thus the midpoint quadrature error is at most $\frac{L+L_1R}{4(1-e^{-1})N}$.

We next control the finite-difference error. At fixed z, consider the map $\boldsymbol{u} \mapsto f(\mathbf{1} - e^{-z\boldsymbol{u}})$. Its gradient is

$$z \operatorname{Diag}(e^{-z\boldsymbol{u}}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{u}}).$$

Under the standing restriction $0 < \mu \le \min\{1, \bar{\mu}/(2e)\}$, all finite-difference points lie in the query neighborhood, and the intermediate arguments $\boldsymbol{X}_k + s\boldsymbol{w}_\ell$, $|s| \le \mu$, have coordinates in $[-1, 2]$. Hence

$$\|\operatorname{Diag}(e^{-z\boldsymbol{u}})\| \le e,$$

and, for any two such points $\boldsymbol{u}, \boldsymbol{v}$,

$$\|\mathbf{1} - e^{-z\boldsymbol{u}} - (\mathbf{1} - e^{-z\boldsymbol{v}})\| \le ez\|\boldsymbol{u} - \boldsymbol{v}\|, \qquad \|\operatorname{Diag}(e^{-z\boldsymbol{u}}) - \operatorname{Diag}(e^{-z\boldsymbol{v}})\| \le ez\|\boldsymbol{u} - \boldsymbol{v}\|.$$

It follows that

$$\left\| z \operatorname{Diag}(e^{-z\boldsymbol{u}}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{u}}) - z \operatorname{Diag}(e^{-z\boldsymbol{v}}) \nabla f(\mathbf{1} - e^{-z\boldsymbol{v}}) \right\|$$
$$\le z^2(e^2 L_1 + eL)\|\boldsymbol{u} - \boldsymbol{v}\|.$$

Therefore the centered difference of $\boldsymbol{u} \mapsto f(\mathbf{1} - e^{-z\boldsymbol{u}})$ in direction $\boldsymbol{w}_\ell$ has error at most $z^2(e^2L_1 + eL)\mu/2$. Multiplying by $e^{z-1}/[(1-e^{-1})z]$ gives an error at most $\frac{(e^2L_1+eL)\mu}{2(1-e^{-1})}$ in each reduced coordinate. Thus

$$\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \frac{L + L_1 R}{4(1-e^{-1})N} + \frac{(e^2 L_1 + eL)\mu\sqrt{r}}{2(1-e^{-1})}.$$

It is sufficient to take

$$0 < \mu \le \min\left\{1, \frac{\bar{\mu}}{2e}, \frac{(1-e^{-1})\varepsilon_g}{(e^2L_1 + eL)\sqrt{r}}\right\}, \qquad N \ge \frac{L + L_1 R}{2(1-e^{-1})\varepsilon_g}. \tag{50}$$

Then $\|\boldsymbol{G}_k - \boldsymbol{g}(\boldsymbol{X}_k)\| \le \varepsilon_g$, so H_k holds deterministically. Since $N = O(\varepsilon^{-1})$, one exact fallback uses $O(r\varepsilon^{-1})$ original-value calls.

Stochastic directional increments. Condition on $\mathcal{F}_{k-1/2}$. For $\boldsymbol{d} \in \mathfrak{D}_k$,

$$\widetilde{A}_{k,\boldsymbol{d}} = \frac{1}{m} \sum_{i=1}^{m} \frac{\widehat{f}\big(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d})}, \xi_i\big) - \widehat{f}\big(\mathbf{1} - e^{-Z_i \boldsymbol{X}_k}, \xi_i\big)}{Z_i},$$

where Z_i has density $e^{z-1}/(1-e^{-1})$ on $(0,1)$.

Taking expectation over the oracle noise and Z_i gives

$$\mathbb{E}[\widetilde{A}_{k,\boldsymbol{d}} \mid \mathcal{F}_{k-1/2}] = \int_0^1 \frac{e^{z-1}}{(1-e^{-1})z} \left[f\Big(\mathbf{1} - e^{-z(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d})}\Big) - f\big(\mathbf{1} - e^{-z\boldsymbol{X}_k}\big) \right] dz$$
$$= A_{k,\boldsymbol{d}}.$$

Thus the estimator is conditionally unbiased.

For $\boldsymbol{x}, \boldsymbol{y} \in [0,1]^n$, the map $\boldsymbol{x} \mapsto \mathbf{1} - e^{-z\boldsymbol{x}}$ is z-Lipschitz, because its Jacobian has norm at most z. Hence, conditional on $Z_i = z$, the two paired query points are separated by at most $z\Delta_k$. Since the multiplier is $1/z$, equation 39 gives

$$\mathbb{E}\left[\left|\frac{\widehat{f}(\mathbf{1} - e^{-z(\boldsymbol{X}_k + \Delta_k W \boldsymbol{d})}, \xi_i) - \widehat{f}(\mathbf{1} - e^{-z\boldsymbol{X}_k}, \xi_i)}{z}\right|^2 \,\middle|\, \mathcal{F}_{k-1/2}, Z_i = z\right]$$
$$\leq \frac{1}{z^2}(L^2 + D_2) z^2 \Delta_k^2 = (L^2 + D_2)\Delta_k^2.$$

The bound is independent of z. As before, the conditional variance of one replication is bounded by this second moment, and independence of the m replications gives

$$\mathbb{E}\Big[|\widetilde{A}_{k,\boldsymbol{d}} - A_{k,\boldsymbol{d}}|^2 \,\Big|\, \mathcal{F}_{k-1/2}\Big] \leq \frac{(L^2 + D_2)\Delta_k^2}{m}.$$

Thus equation 40 holds whenever

$$m \geq \frac{|\mathfrak{D}_k|(L^2 + D_2)}{(1 - p_f)\varepsilon_f^2 \Delta_k^2}. \tag{51}$$

This verifies both reliability conditions involving J_k. A complete stochastic poll uses $O(|\mathfrak{D}_k|^2 \Delta_k^{-2})$ original-value calls.

Stochastic reduced-gradient estimator. For coordinate ℓ,

$$(\boldsymbol{G}_k)_\ell = \frac{1}{m}\sum_{i=1}^{m} \frac{\widehat{f}\big(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)}, \xi_{i,\ell}\big) - \widehat{f}\big(\mathbf{1} - e^{-Z_i(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)}, \xi_{i,\ell}\big)}{2\mu Z_i}.$$

Taking expectation over the oracle noise and Z_i gives

$$\mathbb{E}[(\boldsymbol{G}_k)_\ell \mid \mathcal{F}_{k-1}] = \frac{F(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell) - F(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)}{2\mu}.$$

Since F is L_g-smooth,

$$\|\mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}] - \boldsymbol{g}(\boldsymbol{X}_k)\| \leq \frac{L_g \mu \sqrt{r}}{2}.$$

With $L_g = \frac{e^2 L_1 + eL}{e-1}$ and the choice in equation 50, this is at most $\varepsilon_g/2$.

We finally control the sampling error. Conditional on $Z_i = z$, the points $\boldsymbol{X}_k \pm \mu \boldsymbol{w}_\ell$ lie in the finite-difference neighborhood. On this neighborhood, the Jacobian of $\boldsymbol{x} \mapsto \mathbf{1} - e^{-z\boldsymbol{x}}$ has norm at most ez. Hence the two transformed query points are separated by at most $2ez\mu$. Including the multiplier $1/(2\mu z)$ in equation 39 gives

$$\mathbb{E}\left[\left|\frac{\widehat{f}(\mathbf{1} - e^{-z(\boldsymbol{X}_k + \mu \boldsymbol{w}_\ell)}, \xi_{i,\ell}) - \widehat{f}(\mathbf{1} - e^{-z(\boldsymbol{X}_k - \mu \boldsymbol{w}_\ell)}, \xi_{i,\ell})}{2\mu z}\right|^2 \,\middle|\, \mathcal{F}_{k-1}, Z_i = z\right]$$
$$\leq \frac{1}{4\mu^2 z^2}(L^2 + D_2)(2ez\mu)^2 = e^2(L^2 + D_2).$$

Therefore

$$\mathbb{E}\big[\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\|^2 \,\big|\, \mathcal{F}_{k-1}\big] \leq \frac{re^2(L^2 + D_2)}{m}.$$

Choose

$$m \geq \frac{4re^2(L^2 + D_2)}{(1 - p_d)\varepsilon_g^2}. \tag{52}$$

Then the conditional mean-square sampling error is at most $(1 - p_d)\varepsilon_g^2/4$. Markov's inequality therefore gives, with conditional probability at least p_d,

$$\|\boldsymbol{G}_k - \mathbb{E}[\boldsymbol{G}_k \mid \mathcal{F}_{k-1}]\| \leq \frac{\varepsilon_g}{2}.$$

Combining this with the deterministic bias bound yields

$$\mathbb{P}(H_k \mid \mathcal{F}_{k-1}) \geq p_d.$$

Since $m = O(r\varepsilon^{-2})$ replications are used for each of the r reduced coordinates, one stochastic fallback uses $O(r^2\varepsilon^{-2})$ original-value calls.

D PROOF OF THE VALUE-ORACLE COMPLEXITY BOUNDS

This section proves Theorem 4.1. The first subsection establishes the exact-oracle bound by controlling the cumulative inverse polling radius, while the second establishes the stochastic-oracle bound by controlling the cumulative squared inverse polling radius.

D.1 EXACT VALUE ORACLE

For the exact oracle, $p_f = 1$, so J_k holds deterministically. Before T_ε, if $\Delta_k \le \delta_\varepsilon/\theta$, the small-step analysis in Appendix B.4 shows that a good iteration cannot fail. Hence, on $\{T_\varepsilon > k,\ \Delta_k \le \delta_\varepsilon/\theta\}$, $\mathcal{G}_k$ implies success. Since $\mathbb{P}(\mathcal{G}_k \mid \mathcal{F}_{k-1}) \ge p_d$,

$$\mathbb{P}(\text{iteration } k \text{ is successful} \mid \mathcal{F}_{k-1}) \ge p_d.$$

Assume

$$p_d\beta^{-1} + (1-p_d)\theta^{-1} < 1.$$

If $\Delta_k > \delta_\varepsilon/\theta$, failure gives $\Delta_{k+1} = \theta\Delta_k > \delta_\varepsilon$, while success gives $\Delta_{k+1} = \min\{\beta\Delta_k, \delta_{\max}\} \ge \delta_\varepsilon$. Thus $\Delta_{k+1}^{-1} \le \delta_\varepsilon^{-1}$.

Now suppose $\Delta_k \le \delta_\varepsilon/\theta$. On success,

$$\Delta_{k+1}^{-1} = \max\{\beta^{-1}\Delta_k^{-1}, \delta_{\max}^{-1}\} \le \beta^{-1}\Delta_k^{-1} + \delta_\varepsilon^{-1},$$

where $\delta_{\max} \ge \delta_\varepsilon$. On failure, $\Delta_{k+1}^{-1} = \theta^{-1}\Delta_k^{-1}$. Since $\beta^{-1} < \theta^{-1}$ and the conditional success probability is at least p_d,

$$\mathbb{E}[\Delta_{k+1}^{-1} \mid \mathcal{F}_{k-1}] \le [p_d\beta^{-1} + (1-p_d)\theta^{-1}]\Delta_k^{-1} + \delta_\varepsilon^{-1}.$$

The same bound holds when $\Delta_k > \delta_\varepsilon/\theta$, because then $\Delta_{k+1}^{-1} \le \delta_\varepsilon^{-1}$. Therefore, before the hitting time,

$$\mathbf{1}_{\{T_\varepsilon > k\}}\mathbb{E}[\Delta_{k+1}^{-1} \mid \mathcal{F}_{k-1}] \le \mathbf{1}_{\{T_\varepsilon > k\}}\left([p_d\beta^{-1} + (1-p_d)\theta^{-1}]\Delta_k^{-1} + \delta_\varepsilon^{-1}\right).$$

Since $\{T_\varepsilon > k+1\} \subseteq \{T_\varepsilon > k\}$, taking expectations yields

$$\mathbb{E}[\Delta_{k+1}^{-1}\mathbf{1}_{\{T_\varepsilon > k+1\}}] \le [p_d\beta^{-1} + (1-p_d)\theta^{-1}]\mathbb{E}[\Delta_k^{-1}\mathbf{1}_{\{T_\varepsilon > k\}}] + \delta_\varepsilon^{-1}\mathbb{P}(T_\varepsilon > k).$$

Summing over $k = 0, \dots, K-1$ and rearranging gives

$$[1 - p_d\beta^{-1} - (1-p_d)\theta^{-1}]\sum_{k=0}^{K-1}\mathbb{E}[\Delta_k^{-1}\mathbf{1}_{\{T_\varepsilon > k\}}] \le \delta_0^{-1} + \delta_\varepsilon^{-1}\sum_{k=0}^{K-1}\mathbb{P}(T_\varepsilon > k),$$

where the nonnegative terminal term has been dropped.

Letting $K \to \infty$, monotone convergence and $\sum_{k\ge 0}\mathbb{P}(T_\varepsilon > k) = \mathbb{E}[T_\varepsilon]$ give

$$\mathbb{E}\left[\sum_{k<T_\varepsilon}\Delta_k^{-1}\right] \le \frac{\delta_0^{-1} + \delta_\varepsilon^{-1}\mathbb{E}[T_\varepsilon]}{1 - p_d\beta^{-1} - (1-p_d)\theta^{-1}}. \tag{53}$$

Since $\delta_\varepsilon = \Theta(\varepsilon)$ and $\mathbb{E}[T_\varepsilon] = O(\varepsilon^{-2})$,

$$\mathbb{E}\left[\sum_{k<T_\varepsilon}\Delta_k^{-1}\right] = O(\varepsilon^{-3}).$$

It remains to count value-oracle calls. Therefore

$$\begin{aligned}\mathbb{E}[\text{exact value-oracle calls before } T_\varepsilon] &\le O(r)\,\mathbb{E}\left[\sum_{k<T_\varepsilon}\Delta_k^{-1}\right] + O(r\varepsilon^{-1})\mathbb{E}[T_\varepsilon] \\ &= O(r\varepsilon^{-3}).\end{aligned}$$

This proves the exact-oracle part of Theorem 4.1.

D.2 STOCHASTIC VALUE ORACLE

Before T_ε, if $\Delta_k \le \delta_\varepsilon/\theta$, the small-step analysis in Appendix B.4 shows that a good iteration cannot fail. Hence, on $\{T_\varepsilon > k,\ \Delta_k \le \delta_\varepsilon/\theta\}$, $\mathcal{G}_k$ implies success. Since $\mathbb{P}(\mathcal{G}_k \mid \mathcal{F}_{k-1}) \ge p_d p_f$,

$$\mathbb{P}(\text{iteration } k \text{ is successful} \mid \mathcal{F}_{k-1}) \ge p_d p_f.$$

Assume

$$p_d p_f \beta^{-2} + (1 - p_d p_f)\theta^{-2} < 1.$$

If $\Delta_k > \delta_\varepsilon/\theta$, failure gives $\Delta_{k+1} = \theta\Delta_k > \delta_\varepsilon$, while success gives $\Delta_{k+1} = \min\{\beta\Delta_k, \delta_{\max}\} \ge \delta_\varepsilon$. Thus $\Delta_{k+1}^{-2} \le \delta_\varepsilon^{-2}$.

Now suppose $\Delta_k \le \delta_\varepsilon/\theta$. On success,

$$\Delta_{k+1}^{-2} = \max\{\beta^{-2}\Delta_k^{-2}, \delta_{\max}^{-2}\} \le \beta^{-2}\Delta_k^{-2} + \delta_\varepsilon^{-2},$$

where $\delta_{\max} \ge \delta_\varepsilon$. On failure, $\Delta_{k+1}^{-2} = \theta^{-2}\Delta_k^{-2}$. Since $\beta^{-2} < \theta^{-2}$ and the conditional success probability is at least $p_d p_f$,

$$\mathbb{E}[\Delta_{k+1}^{-2} \mid \mathcal{F}_{k-1}] \le [p_d p_f \beta^{-2} + (1 - p_d p_f)\theta^{-2}]\Delta_k^{-2} + \delta_\varepsilon^{-2}.$$

The same bound holds when $\Delta_k > \delta_\varepsilon/\theta$, because then $\Delta_{k+1}^{-2} \le \delta_\varepsilon^{-2}$. Therefore, before the hitting time,

$$\mathbf{1}_{\{T_\varepsilon > k\}}\mathbb{E}[\Delta_{k+1}^{-2} \mid \mathcal{F}_{k-1}] \le \mathbf{1}_{\{T_\varepsilon > k\}}\left([p_d p_f \beta^{-2} + (1 - p_d p_f)\theta^{-2}]\Delta_k^{-2} + \delta_\varepsilon^{-2}\right).$$

Since $\{T_\varepsilon > k+1\} \subseteq \{T_\varepsilon > k\}$, taking expectations yields

$$\mathbb{E}[\Delta_{k+1}^{-2}\mathbf{1}_{\{T_\varepsilon > k+1\}}] \le [p_d p_f \beta^{-2} + (1 - p_d p_f)\theta^{-2}]\mathbb{E}[\Delta_k^{-2}\mathbf{1}_{\{T_\varepsilon > k\}}] + \delta_\varepsilon^{-2}\mathbb{P}(T_\varepsilon > k).$$

Summing over $k = 0, \dots, K-1$ and rearranging gives

$$[1 - p_d p_f \beta^{-2} - (1 - p_d p_f)\theta^{-2}] \sum_{k=0}^{K-1} \mathbb{E}[\Delta_k^{-2}\mathbf{1}_{\{T_\varepsilon > k\}}] \le \delta_0^{-2} + \delta_\varepsilon^{-2} \sum_{k=0}^{K-1} \mathbb{P}(T_\varepsilon > k),$$

where the nonnegative terminal term has been dropped.

Letting $K \to \infty$, monotone convergence and $\sum_{k\ge 0} \mathbb{P}(T_\varepsilon > k) = \mathbb{E}[T_\varepsilon]$ give

$$\mathbb{E}\left[\sum_{k<T_\varepsilon} \Delta_k^{-2}\right] \le \frac{\delta_0^{-2} + \delta_\varepsilon^{-2}\mathbb{E}[T_\varepsilon]}{1 - p_d p_f \beta^{-2} - (1 - p_d p_f)\theta^{-2}}. \tag{54}$$

Since $\delta_\varepsilon = \Theta(\varepsilon)$ and $\mathbb{E}[T_\varepsilon] = O(\varepsilon^{-2})$,

$$\mathbb{E}\left[\sum_{k<T_\varepsilon} \Delta_k^{-2}\right] = O(\varepsilon^{-4}).$$

It remains to count value-oracle calls. Therefore

$$\begin{aligned}\mathbb{E}[\text{stochastic value-oracle calls before } T_\varepsilon] &\le O(r^2)\,\mathbb{E}\left[\sum_{k<T_\varepsilon} \Delta_k^{-2}\right] + O(r^2\varepsilon^{-2})\mathbb{E}[T_\varepsilon] \\ &= O(r^2\varepsilon^{-4}).\end{aligned}$$

This proves the stochastic-oracle part of Theorem 4.1.

E EXPERIMENTAL DETAILS AND ADDITIONAL RESULTS

This section provides the full experimental setup and additional results. We first describe the benchmark objectives and constraints, the exact and stochastic oracle models, the compared methods, and the tuning protocol. We then report iteration and value-oracle results across all benchmark families and dimensions, together with the selected parameters and the observed frequency of degenerate SDS states.

E.1 Experimental Settings

Throughout the experiments, we use the following six settings: S1 is monotone maximization over a general polytope with an exact oracle, S2 is the corresponding stochastic-oracle setting, S3 is nonmonotone maximization over a general polytope with an exact oracle, S4 is the corresponding stochastic-oracle setting, S5 is nonmonotone maximization over a down-closed polytope with an exact oracle, and S6 is the corresponding stochastic-oracle setting. Each figure has exact-oracle results in the top row (S1, S3, and S5) and stochastic-oracle results in the bottom row (S2, S4, and S6).

We use three benchmark families: quadratic DR-submodular objectives, finite sums of quadratic objectives, and regular coverage objectives (Staib & Jegelka, 2020; Lian et al., 2024; Chen et al., 2023). For a fixed problem, dimension, and seed, each exact–stochastic pair uses the same objective and constraints. The general and down-closed nonmonotone settings also share the same objective, differing only in the lower-bound constraint.

- **Quadratic objective.** We use
$$f(\boldsymbol{x}) = \frac{1}{2}\boldsymbol{x}^\top H\boldsymbol{x} + \boldsymbol{h}^\top \boldsymbol{x} + c$$
with $n = 100$. The upper-triangular entries of H, including the diagonal, are sampled independently from $\mathrm{Unif}[-1, 0]$ and mirrored to form a symmetric matrix. For the monotone case, we set $\boldsymbol{h} = -H\mathbf{1}$ and $c = 0$. For the nonmonotone case, the entries of $\boldsymbol{h}$ are sampled independently from $\mathrm{Unif}[0, 1]$, independently of H, and we set $c = \sum_i [(-H\mathbf{1}/2)_i - h_i]_+$. This makes f nonnegative on $[0, 1]^n$.
- **Finite-sum quadratic objective.** Following the finite-sum setting in Lian et al. (2024), we consider
$$f(\boldsymbol{x}) = \frac{1}{50}\sum_{t=1}^{50}\left(\frac{1}{2}\boldsymbol{x}^\top H_t\boldsymbol{x} + \boldsymbol{h}_t^\top \boldsymbol{x} + c_t\right), \qquad n \in \{50, 100, 300\}.$$
Each $H_t, \boldsymbol{h}_t, c_t$ is generated by the same rule as above. Every oracle call returns the complete objective f, not a randomly sampled component.
- **Regular coverage objective.** Following the regular coverage model in Chen et al. (2023), we use a cyclic set construction with $n = 100$. Let $V = \{1, \dots, n\}$, and let π be a random permutation of V. Set $S_{\pi_j} = \{\pi_j, \pi_{j+1}, \pi_{j+2}\}$, with cyclic indexing. Every set contains three elements, and every element belongs to three sets. For $C_u = \{i : u \in S_i\}$, the multilinear coverage function is
$$g(\boldsymbol{x}) = \sum_{u=1}^{n}\left[1 - \prod_{i \in C_u}(1 - x_i)\right].$$
We use $f(\boldsymbol{x}) = g(\boldsymbol{x})$ for the monotone problems and $f(\boldsymbol{x}) = g(\boldsymbol{x}) - \sum_i x_i$ for the nonmonotone problems. The latter is the multilinear extension of $\left|\bigcup_{i \in X} S_i\right| - |X|$. Both objectives are nonnegative and DR-submodular on $[0, 1]^n$.

Constraint sets. For the quadratic problems, we draw $B_{\mathrm{res}} \in \mathbb{R}^{30 \times n}$ with independent $\mathrm{Unif}[0, 1]$ entries and use
$$\mathcal{P}_G = \{\boldsymbol{x} \in [0, 1]^n : B_{\mathrm{res}}\boldsymbol{x} \le \mathbf{1},\ x_1 \ge 0.25\},$$
$$\mathcal{P}_D = \{\boldsymbol{x} \in [0, 1]^n : B_{\mathrm{res}}\boldsymbol{x} \le \mathbf{1}\}.$$
The positive lower bound makes $\mathcal{P}_G$ non-down-closed, whereas $\mathcal{P}_D$ is down-closed because $B_{\mathrm{res}} \ge 0$. For the coverage problems, we draw $B_{\mathrm{raw}} \in \mathbb{R}^{30 \times 100}$ in the same way and use
$$\mathcal{P}_G^{\mathrm{cov}} = \{\boldsymbol{x} \in [0, 1]^{100} : B_{\mathrm{raw}}\boldsymbol{x} \le 0.65 B_{\mathrm{raw}}\mathbf{1},\ x_1 \ge 0.25\},$$
$$\mathcal{P}_D^{\mathrm{cov}} = \{\boldsymbol{x} \in [0, 1]^{100} : B_{\mathrm{raw}}\boldsymbol{x} \le 0.65 B_{\mathrm{raw}}\mathbf{1}\}.$$
All experiments have 30 resource inequalities and no equality constraints, thus the reduced dimension in the main text is $r = n$ and $W = I_n$.

Value-oracle models. The exact oracle returns $f(\boldsymbol{x})$. For the stochastic experiments, we use

$$\widehat{f}(\boldsymbol{x}, \boldsymbol{\xi}) = f(\boldsymbol{x}) + \frac{1}{\sqrt{n}} \boldsymbol{\xi}^\top \boldsymbol{x},$$

where the entries of $\boldsymbol{\xi}$ are independent Rademacher random variables. This oracle satisfies both the stochastic condition used in our analysis and the stochastic oracle assumptions required by the baseline methods considered below. For a fair comparison, we provide the stochastic baselines that use paired differences with the same information available to SDS: the two queried points in each difference share the same realization of $\boldsymbol{\xi}$, while different pairs use independent realizations. All methods use the fixed finite-difference radius $\mu = 10^{-3}$. Each scalar oracle evaluation counts as one value-oracle call.

Compared methods. All compared methods make their optimization decisions using only function values. For methods without a standard algorithm name, we use PA24, CSYZZ26, and LJPSA26 for Pedramfar & Aggarwal (2024), Chen et al. (2026), and Lu et al. (2026), respectively. Let $\boldsymbol{l} = 0.25\boldsymbol{e}_1$ in S1–S4 and $\boldsymbol{l} = \boldsymbol{0}$ in S5–S6. The common starting point is $\boldsymbol{q} = \boldsymbol{l} + 10^{-3}\boldsymbol{1}$. FW, ZO-GA, RG-ZOSA, and BCG start directly from $\boldsymbol{q}$. PA24 uses $\boldsymbol{q}$ as both its anchor and initial internal point. For SDS and LJPSA26, we initialize the internal variables through their output maps so that the initial reported output is $\boldsymbol{q}$. CSYZZ26 starts preprocessing from $\boldsymbol{q}$ and restarts its main stage from the origin, as detailed below. For methods that estimate a gradient from random directions, we use

$$\widehat{\boldsymbol{g}}(\boldsymbol{x}) = \frac{n}{2\mu BT} \sum_{b=1}^{B} \sum_{t=1}^{T} \left[\widehat{f}(\boldsymbol{x} + \mu \boldsymbol{u}_b, \boldsymbol{\xi}_{bt}) - \widehat{f}(\boldsymbol{x} - \mu \boldsymbol{u}_b, \boldsymbol{\xi}_{bt}) \right] \boldsymbol{u}_b,$$

where the $\boldsymbol{u}_b$ are independent and uniform on the unit sphere. Here B is the number of directions and T is the number of replications per direction. For the exact oracle, $\widehat{f}$ is replaced by f. A two-point estimate costs $2BT$ calls. When a one-point estimate is used, it averages samples of $n\widehat{f}(\boldsymbol{x} + \mu \boldsymbol{u}, \boldsymbol{\xi})\boldsymbol{u}/\mu$ and costs BT calls.

- **FW.** We use the value-only implementation of the appropriate branch of Pedramfar et al. (2023) in all six settings. At iteration k, a two-point gradient estimate is combined with the running estimate $\boldsymbol{d}_k = (1-\rho_k)\boldsymbol{d}_{k-1} + \rho_k \widehat{\boldsymbol{g}}_k$, where $\rho_k = \min\{1, \rho_{\text{scale}}/(k+3)^{\rho_{\text{power}}}\}$. A linear optimization problem then gives the Frank–Wolfe direction $\boldsymbol{v}_k$. For S1–S2, we use $\boldsymbol{x}_{k+1} = (1-\eta)\boldsymbol{x}_k + \eta \boldsymbol{v}_k$ with $\eta = \log(K)/(2K)$. For S3–S4, we use the same update with $\eta = \log(2)/K$. For S5–S6, the update is $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \boldsymbol{v}_k/K$. The horizon is fixed before each run as $K = \max\{500, \lceil 49999/(2BT+1) \rceil\}$, so the same run covers both the 500-iteration and 50,000-call checkpoints.
- **ZO-GA.** We use the zeroth-order projected gradient ascent method in Lian et al. (2024) for S1. It estimates the gradient of f with the two-point estimator and projects the resulting ascent step onto $\mathcal{P}$. At update $k \geq 1$, the stepsize is $\eta_{\text{scale}}/k^{\eta_{\text{power}}}$.
- **RG-ZOSA.** We use a full-objective of RG-ZOSA from Lian et al. (2024) for S1. At each iteration, we sample R values of Z with density $e^{z-1}/(1-e^{-1})$ and estimate $\nabla f(Z\boldsymbol{x}_k)$ using B two-point directions per sample. The update uses

  $$\boldsymbol{G}_k = \frac{1-e^{-1}}{R} \sum_{j=1}^{R} \widehat{\boldsymbol{g}}(Z_j \boldsymbol{x}_k), \qquad \boldsymbol{x}_{k+1} = \text{Proj}_{\mathcal{P}}(\boldsymbol{x}_k + \eta_k \boldsymbol{G}_k),$$

  with the same power-law stepsize convention as ZO-GA.
- **BCG.** We use Black-box Continuous Greedy (BCG) of Chen et al. (2020a) for S2. Starting from $\boldsymbol{q}$, BCG uses the shared two-point CRN estimator and the same momentum update as FW, $\boldsymbol{d}_k = (1-\rho_k)\boldsymbol{d}_{k-1} + \rho_k \widehat{\boldsymbol{g}}_k$, with $\rho_k = \min\{1, \rho_{\text{scale}}/(k+3)^{\rho_{\text{power}}}\}$. A linear optimization problem gives the direction $\boldsymbol{v}_k$, and the iterate is updated as $\boldsymbol{x}_{k+1} = \boldsymbol{x}_k + \boldsymbol{v}_k/K$.
- **PA24.** For S3–S4, we combine the query transformation of Pedramfar & Aggarwal (2024) with projected gradient ascent and zeroth-order gradient estimates. Let $\boldsymbol{q}$ be the anchor. At each iteration, we sample R values of Z with density $1/[3(1-z/2)^3]$ and form

  $$\boldsymbol{G}_k = \frac{1}{R} \sum_{j=1}^{R} \widehat{\boldsymbol{g}} \left(\boldsymbol{q} + \frac{Z_j}{2} (\boldsymbol{x}_k - \boldsymbol{q}) \right).$$

We then update $\boldsymbol{x}_{k+1} = \mathrm{Proj}_{\mathcal{P}}(\boldsymbol{x}_k + \eta_k \boldsymbol{G}_k)$, using the power-law stepsize specified above. The reported point is $(\boldsymbol{x}_k + \boldsymbol{q})/2$, and the plotted value is $f((\boldsymbol{x}_k + \boldsymbol{q})/2)$.

- **CSYZZ26.** For S5, we use the classical aided Frank–Wolfe method of Chen et al. (2026). The method first computes an approximate stationary point $\boldsymbol{y}$ from the common initial point $\boldsymbol{q}$, using coordinatewise centered differences and stopping when the Frank–Wolfe gap is at most 0.1, with at most 250 preprocessing iterations. The main stage restarts from $\boldsymbol{x}_0 = \boldsymbol{0}$. With $s = 0.372$, it uses the two prescribed Frank–Wolfe phases before and after iteration $\lfloor sK \rfloor$, with stepsizes $1 - e^{-1/K}$ and $e^{-1/K}/K$, respectively. Gradient estimates use B two-point directions and momentum $\rho_j = \min\{1, \rho_{\mathrm{scale}}/(j + \rho_{\mathrm{offset}})^{2/3}\}$. The method returns $\boldsymbol{y}$ or the main-stage iterate at random, so we plot the corresponding expected value $pf(\boldsymbol{y}) + (1-p)f(\boldsymbol{x}_j)$, where $p = [e^s(3-s) - 3]/[e + e^s(3-s) - 3]$. For the iteration plots, preprocessing is completed before iteration zero; for the value-oracle plots, all preprocessing calls are included.
- **LJPSA26.** For S6, we instantiate the query transformation of Lu et al. (2026) for a fixed objective, using projected gradient ascent and zeroth-order estimates. At each iteration, we sample R values of Z with density $e^{z-1}/(1 - e^{-1})$ and form

$$\boldsymbol{G}_k = \frac{1}{R} \sum_{j=1}^{R} e^{-Z_j \boldsymbol{x}_k} \odot \widehat{\boldsymbol{g}}(\boldsymbol{1} - e^{-Z_j \boldsymbol{x}_k}).$$

  We use $\boldsymbol{x}_{k+1} = \mathrm{Proj}_{\mathcal{P}}(\boldsymbol{x}_k + \eta_k \boldsymbol{G}_k)$, with the same power-law stepsize convention. The reported point is $\boldsymbol{1} - e^{-\boldsymbol{x}_k}$.
- **SDS.** We use Algorithm 1 with the three non-oblivious objectives in equation 4. We fix $\gamma = 2$ and $\mu = 10^{-3}$. For $p_d = 1$, the full generator set is used; for $p_d < 1$, we use the randomized polling construction in Appendix C.1. The quadrature and Monte Carlo sample sizes are fixed within each experiment and tuned together with $\delta_0, \delta_{\max}, \theta, \beta, \varepsilon_f$, and p_d. For the monotone setting, the reported point is $\boldsymbol{x}_k$. For the general nonmonotone setting, we take $\underline{\boldsymbol{x}} = 0.25\boldsymbol{e}_1$, initialize $\boldsymbol{x}_0 = 2\boldsymbol{q} - \underline{\boldsymbol{x}}$, and report $(\boldsymbol{x}_k + \underline{\boldsymbol{x}})/2$. For the down-closed setting, we initialize $\boldsymbol{x}_0 = -\log(\boldsymbol{1} - \boldsymbol{q})$ and report $\boldsymbol{1} - e^{-\boldsymbol{x}_k}$.

Parameter tuning. We use 10 tuning instances and 10 test instances with disjoint seed sets, and select all parameters using only the tuning instances. Candidate configurations are evaluated by grid search. The tuning score gives equal weight to the final and average current objective values on both the iteration and value-oracle axes.

E.2 Results

Finite-sum quadratic objective. Figures 2 and 3 report the 50-component quadratic problem with $n = 100$. SDS reaches high objective values rapidly across all six settings and attains the highest mean value at the 50,000-call checkpoint in each case. On the iteration axis, PA24 is slightly higher in the stochastic general-nonmonotone setting, but SDS overtakes it on the value-oracle axis. This contrast is consistent with the state-aware design of SDS: on nondegenerate iterations, search directions are obtained directly from the local polyhedral geometry, avoiding a full zeroth-order gradient estimate and reducing the number of oracle calls per iteration. In the exact down-closed setting, the reported objective peaks early and then decreases slightly. This can occur because SDS accepts steps using the non-oblivious objective F, while the figures report the original objective f.

The selected parameters for the $n = 100$ problem are as follows.

For S1, FW uses $B = 50$, $T = 1$, $\rho_{\mathrm{scale}} = 1$, $\rho_{\mathrm{power}} = 1$, and $K = 500$. RG-ZOSA uses $R = 4$, $B = 100$, $\eta_{\mathrm{scale}} = 0.01$, and $\eta_{\mathrm{power}} = 1$. ZO-GA uses $B = 100$, $\eta_{\mathrm{scale}} = 0.01$, and $\eta_{\mathrm{power}} = 1$. SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.6$, $\beta = 1.015$, $\varepsilon_f = 100$, and $p_d = 1$, with one polling sample and one gradient sample.

For S2, FW uses $B = 50$, $T = 1$, $\rho_{\mathrm{scale}} = 0.5$, $\rho_{\mathrm{power}} = 1$, and $K = 500$. BCG uses $B = 50$, $T = 1$, $\rho_{\mathrm{scale}} = 1.5$, $\rho_{\mathrm{power}} = 1.1$, and $K = 500$. SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.48$, $\beta = 1.015$, $\varepsilon_f = 100$, and $p_d = 1$, with one polling sample and one gradient sample.

For S3, FW uses $B = 50$, $T = 1$, $\rho_{\mathrm{scale}} = 0.5$, $\rho_{\mathrm{power}} = 0.5$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\mathrm{scale}} = 1$, and $\eta_{\mathrm{power}} = 1$, with centered two-point estimates. SDS uses

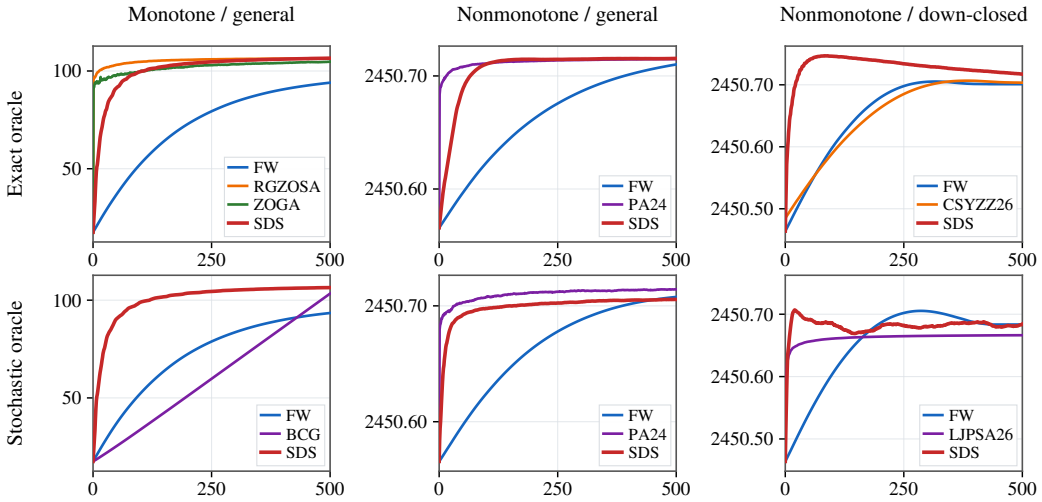


Figure 2: Average current objective value versus iterations for the 50-component quadratic problem with $n = 100$.

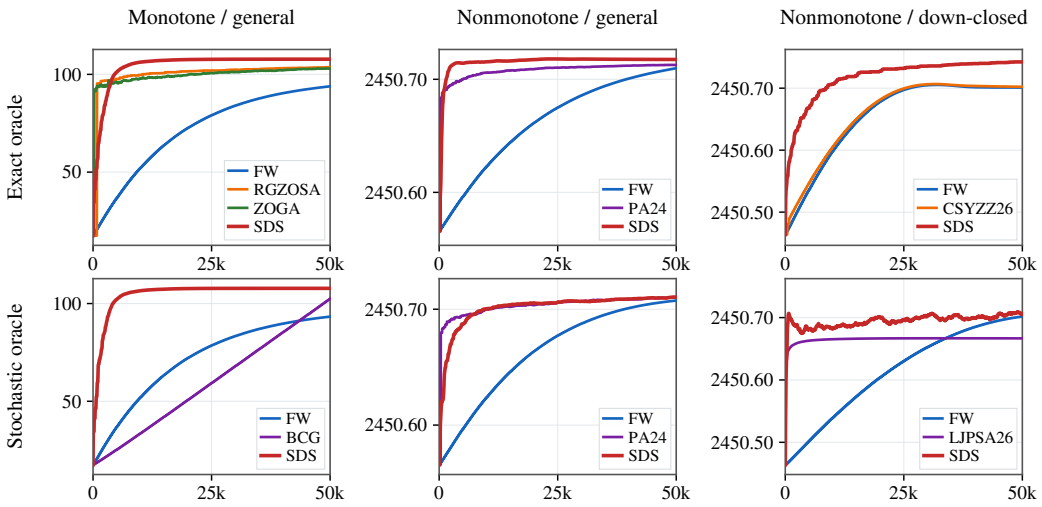


Figure 3: Average current objective value versus value-oracle calls for the 50-component quadratic problem with $n = 100$.

$\delta_0 = 0.0195$, $\delta_{\max} = 0.03$, $\theta = 0.66$, $\beta = 1.01$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and one gradient sample.

For S4, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 2$, $\rho_{\text{power}} = 1$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses $\delta_0 = 0.094$, $\delta_{\max} = 0.1$, $\theta = 0.6$, $\beta = 1.01$, $\varepsilon_f = 1$, and $p_d = 1$, with one polling sample and four gradient samples.

For S5, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. CSYZZ26 uses $B = 40$, $\rho_{\text{scale}} = 2$, and $\rho_{\text{offset}} = 32$. SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.875$, $\beta = 4$, $\varepsilon_f = 3$, and $p_d = 1$, with three polling samples and 16 gradient samples.

For S6, FW uses $B = 100$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 500$. LJPSA26 uses $R = 4$, $B = 1$, $T = 4$, $\eta_{\text{scale}} = 0.003$, and $\eta_{\text{power}} = 1$, with two-point CRN estimates. SDS uses

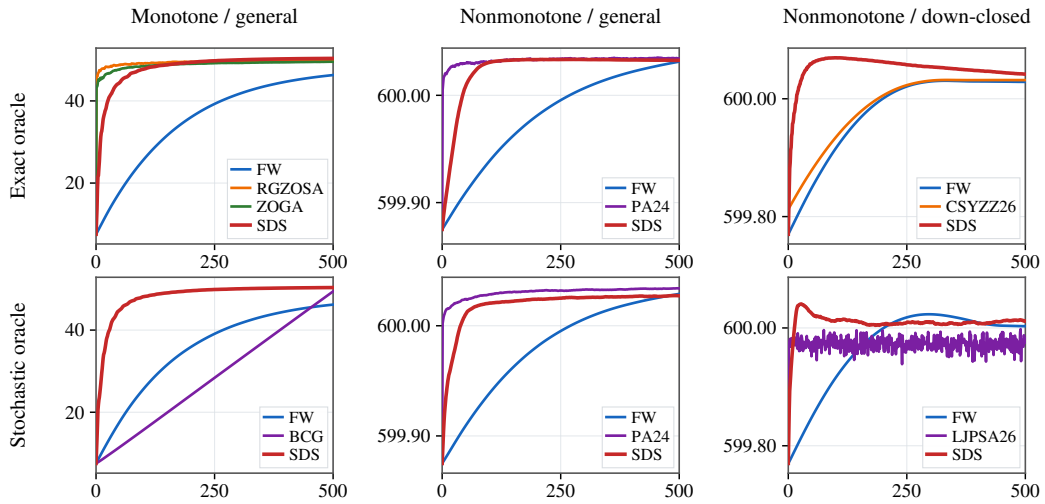


Figure 4: Average current objective value versus iterations for the 50-component quadratic problem with $n = 50$.

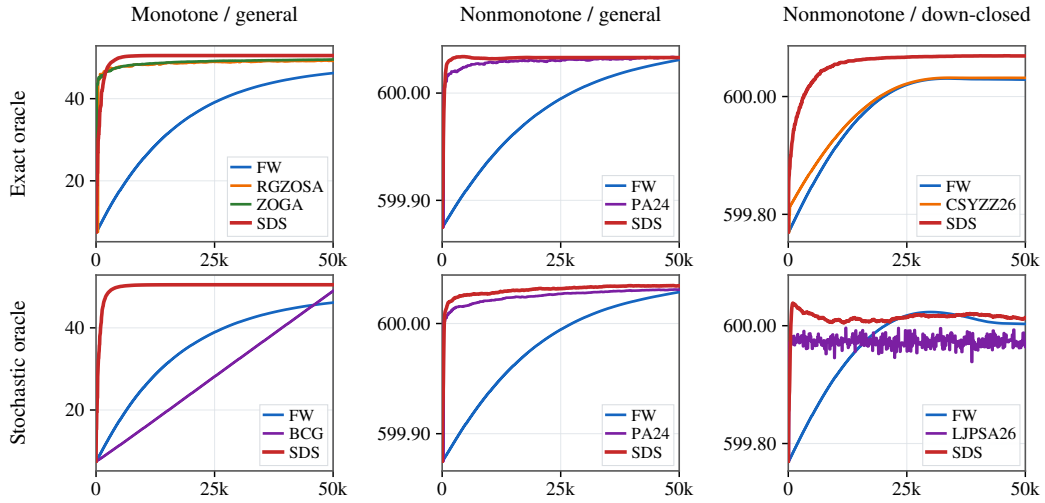


Figure 5: Average current objective value versus value-oracle calls for the 50-component quadratic problem with $n = 50$.

$\delta_0 = 0.095$, $\delta_{\max} = 0.19$, $\theta = 0.4$, $\beta = 1.005$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and one gradient sample.

Figures 4 and 5 report the finite-sum quadratic problem with $n = 50$. The overall behavior is consistent with the $n = 100$ results: SDS reaches competitive objective values rapidly and is particularly effective on the value-oracle axis. Under the stochastic oracle, SDS attains the highest mean value at the 50,000-call checkpoint in all three settings. In S4, PA24 remains slightly higher at 500 iterations, whereas SDS is higher at the query checkpoint, again illustrating the lower per-iteration oracle cost enabled by state-aware direct polling. The mild decrease of SDS in the exact down-closed setting has the same non-oblivious-objective explanation as in the $n = 100$ case.

The selected parameters for the $n = 50$ problem are as follows.

For S1, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1$, $\rho_{\text{power}} = 1$, and $K = 500$. RG-ZOSA uses $R = 4$, $B = 50$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 0.5$. ZO-GA uses $B = 50$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 1$.

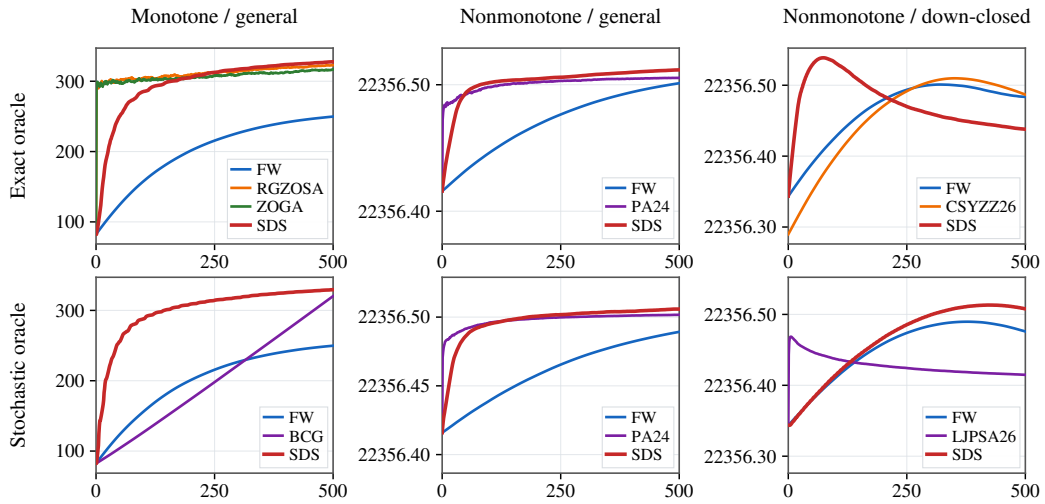


Figure 6: Average current objective value versus iterations for the 50-component quadratic problem with $n = 300$.

SDS uses $\delta_0 = 0.12$, $\delta_{\max} = 0.2$, $\theta = 0.55$, $\beta = 1.02$, $\varepsilon_f = 10$, and $p_d = 1$, with one polling sample and one gradient sample.

For S2, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 2$, $\rho_{\text{power}} = 1$, and $K = 500$. BCG uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1.5$, $\rho_{\text{power}} = 1.1$, and $K = 500$. SDS uses $\delta_0 = 0.12$, $\delta_{\max} = 0.12$, $\theta = 0.32$, $\beta = 1.075$, $\varepsilon_f = 20$, and $p_d = 1$, with one polling sample and one gradient sample.

For S3, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 1$, and $\eta_{\text{power}} = 0.5$, with centered two-point estimates. SDS uses $\delta_0 = 0.021$, $\delta_{\max} = 0.03$, $\theta = 0.66$, $\beta = 1.01$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and four gradient samples.

For S4, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses $\delta_0 = 0.05$, $\delta_{\max} = 0.05$, $\theta = 0.2$, $\beta = 1.02$, $\varepsilon_f = 2$, and $p_d = 0.999$, with one polling sample and four gradient samples.

For S5, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. CSYZZ26 uses $B = 50$, $\rho_{\text{scale}} = 2$, and $\rho_{\text{offset}} = 32$. SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.875$, $\beta = 4$, $\varepsilon_f = 3$, and $p_d = 1$, with three polling samples and 16 gradient samples.

For S6, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 500$. LJPSA26 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.003$, and $\eta_{\text{power}} = 0.5$, with one-point estimates. SDS uses $\delta_0 = 0.1$, $\delta_{\max} = 0.1$, $\theta = 0.4$, $\beta = 1.02$, $\varepsilon_f = 3$, and $p_d = 1$, with two polling samples and 16 gradient samples.

Figures 6 and 7 report the finite-sum quadratic problem with $n = 300$. Under the stochastic oracle, SDS attains the highest mean value at both the 500-iteration and 50,000-call checkpoints in all three settings. It also performs strongly in the exact monotone and general nonmonotone settings, with a particularly clear advantage on the value-oracle axis. This persistence at the larger dimension again reflects the state-aware design: on nondegenerate iterations, SDS obtains search directions from the local geometry without reconstructing a full zeroth-order gradient. In the exact down-closed setting, SDS peaks early and then decreases, showing the same nonoblivious objective behavior observed above.

The selected parameters for the $n = 300$ problem are as follows.

For S1, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1$, $\rho_{\text{power}} = 1$, and $K = 500$. RG-ZOSA uses $R = 1$, $B = 300$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 1$. ZO-GA uses $B = 300$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 1$.

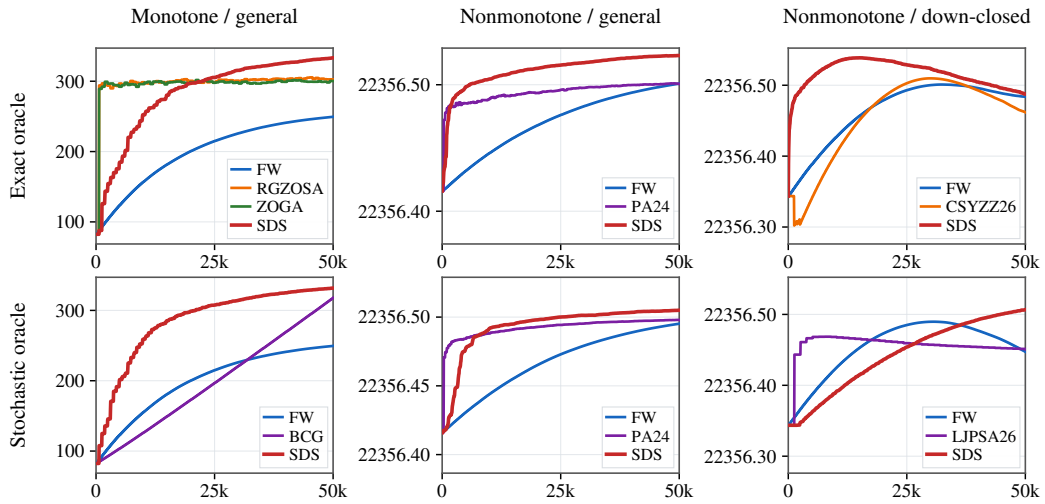


Figure 7: Average current objective value versus value-oracle calls for the 50-component quadratic problem with $n = 300$.

SDS uses $\delta_0 = 0.12$, $\delta_{\max} = 0.2$, $\theta = 0.5$, $\beta = 1.02$, $\varepsilon_f = 10$, and $p_d = 1$, with one polling sample and one gradient sample.

For S2, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 1$, and $K = 500$. BCG uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1.5$, $\rho_{\text{power}} = 1.1$, and $K = 500$. SDS uses $\delta_0 = 0.1$, $\delta_{\max} = 0.12$, $\theta = 0.4$, $\beta = 1.05$, $\varepsilon_f = 10$, and $p_d = 1$, with one polling sample and one gradient sample.

For S3, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 1$, and $\eta_{\text{power}} = 1$, with centered two-point estimates. SDS uses $\delta_0 = 0.0195$, $\delta_{\max} = 0.03$, $\theta = 0.66$, $\beta = 1.005$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and one gradient sample.

For S4, FW uses $B = 10$, $T = 4$, $\rho_{\text{scale}} = 2$, $\rho_{\text{power}} = 1$, and $K = 618$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 1$, with two-point CRN estimates. SDS uses $\delta_0 = 0.028$, $\delta_{\max} = 0.08$, $\theta = 0.5$, $\beta = 1.01$, $\varepsilon_f = 3$, and $p_d = 1$, with two polling samples and one gradient sample.

For S5, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. CSYZZ26 uses $B = 40$, $\rho_{\text{scale}} = 2$, and $\rho_{\text{offset}} = 32$. SDS uses $\delta_0 = 0.02$, $\delta_{\max} = 0.02$, $\theta = 0.97$, $\beta = 1.005$, $\varepsilon_f = 10$, and $p_d = 1$, with one polling sample and eight gradient samples.

For S6, FW uses $B = 10$, $T = 4$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 618$. LJPSA26 uses $R = 4$, $B = 40$, $T = 4$, $\eta_{\text{scale}} = 0.003$, and $\eta_{\text{power}} = 1$, with two-point CRN estimates. SDS uses $\delta_0 = 0.015$, $\delta_{\max} = 0.015$, $\theta = 0.6$, $\beta = 1.1$, $\varepsilon_f = 100$, and $p_d = 1$, with one polling sample and 16 gradient samples.

Single quadratic objective. Figures 8 and 9 report the single-quadratic problem with $n = 100$. SDS reaches the highest final values in all three exact-oracle settings and shows the same strong value-oracle efficiency observed for the finite-sum objectives. Under the stochastic oracle, SDS performs best in the monotone and general nonmonotone settings, with particularly rapid progress on the value-oracle axis. This shows that the benefit of state-aware direct polling is not tied to the finite-sum structure. In the stochastic down-closed setting, LJPSA26 attains a slightly higher final value than SDS.

The selected parameters are as follows.

For S1, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1$, $\rho_{\text{power}} = 1$, and $K = 500$. RG-ZOSA uses $R = 4$, $B = 100$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 1$. ZO-GA uses $B = 100$, $\eta_{\text{scale}} = 0.01$, and $\eta_{\text{power}} = 1$.

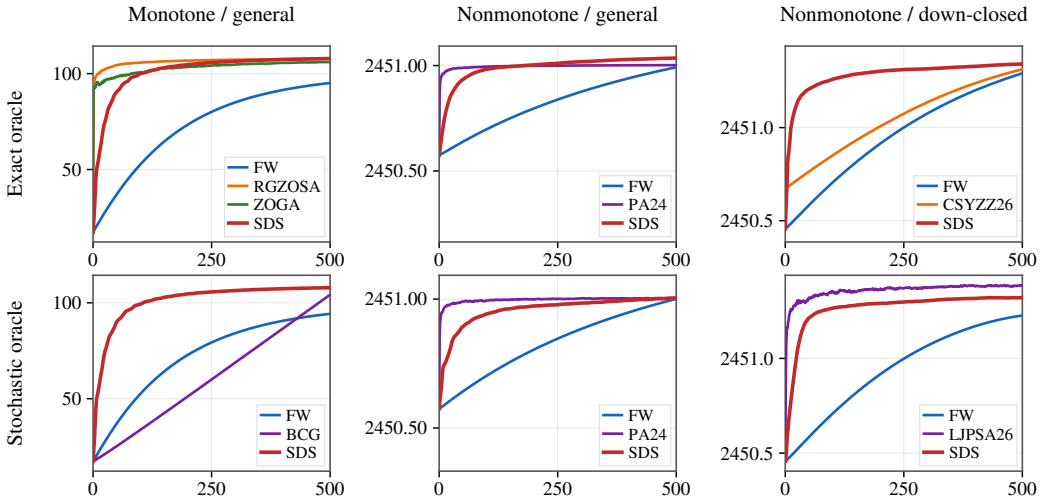


Figure 8: Average current objective value versus iterations for the single quadratic problem with $n = 100$.

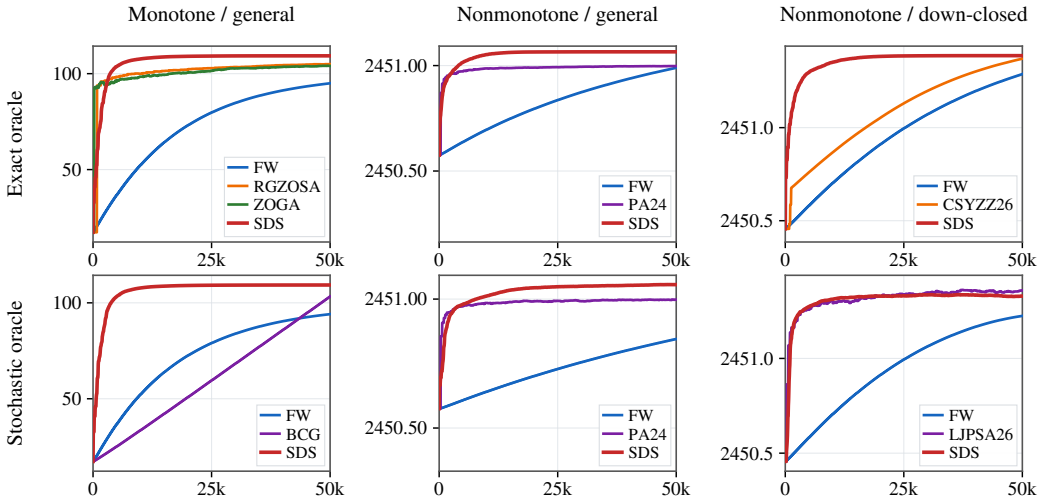


Figure 9: Average current objective value versus value-oracle calls for the single quadratic problem with $n = 100$.

SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.33$, $\beta = 1.05$, $\varepsilon_f = 100$, and $p_d = 1$, with one polling sample and one gradient sample.

For S2, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 1$, and $K = 500$. BCG uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1.5$, $\rho_{\text{power}} = 1.1$, and $K = 500$. SDS uses $\delta_0 = 0.08$, $\delta_{\max} = 0.08$, $\theta = 0.3$, $\beta = 1.05$, $\varepsilon_f = 100$, and $p_d = 1$, with one polling sample and one gradient sample.

For S3, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 1$, and $\eta_{\text{power}} = 1$, with centered two-point estimates. SDS uses $\delta_0 = 0.04$, $\delta_{\max} = 0.08$, $\theta = 0.66$, $\beta = 1.0025$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and two gradient samples.

For S4, FW uses $B = 100$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses

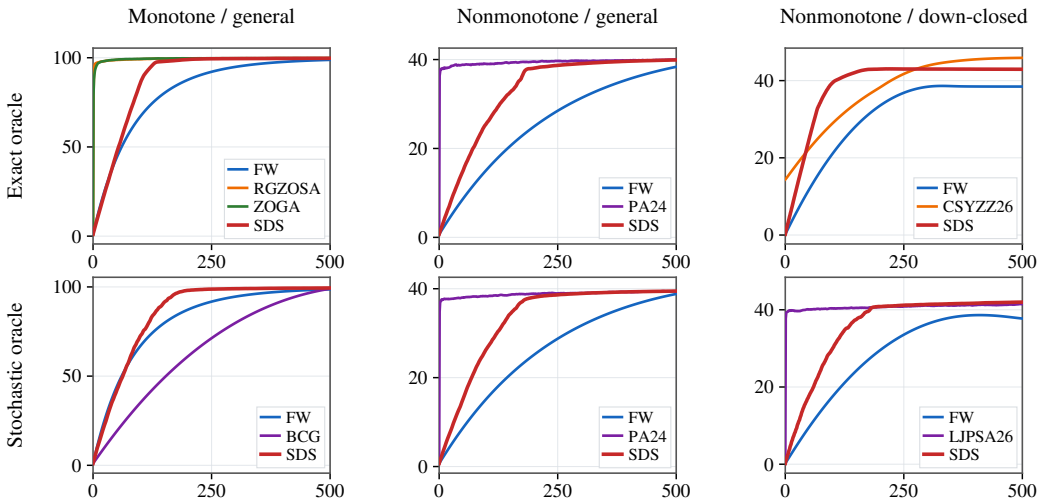


Figure 10: Average current objective value versus iterations for the regular coverage problem with $n = 100$.

$\delta_0 = 0.07$, $\delta_{\max} = 0.07$, $\theta = 0.2$, $\beta = 1.05$, $\varepsilon_f = 1$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S5, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. CSYZZ26 uses $B = 40$, $\rho_{\text{scale}} = 2$, and $\rho_{\text{offset}} = 32$. SDS uses $\delta_0 = 0.1$, $\delta_{\max} = 0.2$, $\theta = 0.66$, $\beta = 1.005$, $\varepsilon_f = 3$, and $p_d = 1$, with one polling sample and one gradient sample.

For S6, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 1$, and $K = 500$. LJPSA26 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses $\delta_0 = 0.069$, $\delta_{\max} = 0.11$, $\theta = 0.7$, $\beta = 1.005$, $\varepsilon_f = 10$, and $p_d = 1$, with one polling sample and one gradient sample.

Regular coverage. Figures 10 and 11 report the regular coverage problem with $n = 100$. Despite the different objective structure, SDS again reaches high objective values rapidly and remains particularly effective on the value-oracle axis. Under the stochastic oracle, SDS attains the highest mean value at both the 500-iteration and 50,000-call checkpoints in all three settings, consistent with the oracle savings from state-aware direct polling. In the exact down-closed setting, CSYZZ26 eventually attains a higher objective value than SDS, which is consistent with its stronger $0.385 - O(\epsilon)$ approximation guarantee compared with the e^{-1} guarantee of SDS. Overall, the coverage results show that the value-oracle efficiency of SDS extends beyond the quadratic benchmarks.

The selected parameters are as follows.

For S1, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1$, $\rho_{\text{power}} = 1$, and $K = 500$. RG-ZOSA uses $R = 1$, $B = 100$, $\eta_{\text{scale}} = 0.2$, and $\eta_{\text{power}} = 0.5$. ZO-GA uses $B = 10$, $\eta_{\text{scale}} = 0.1$, and $\eta_{\text{power}} = 0.5$. SDS uses $\delta_0 = 0.325$, $\delta_{\max} = 0.325$, $\theta = 0.9$, $\beta = 1.00125$, $\varepsilon_f = 2$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S2, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 1$, and $K = 500$. BCG uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 1.5$, $\rho_{\text{power}} = 1.1$, and $K = 500$. SDS uses $\delta_0 = 0.325$, $\delta_{\max} = 0.325$, $\theta = 0.35$, $\beta = 1.02$, $\varepsilon_f = 1$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S3, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. PA24 uses $R = 4$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 1$, and $\eta_{\text{power}} = 0.5$, with centered two-point estimates. SDS uses $\delta_0 = 0.325$, $\delta_{\max} = 0.325$, $\theta = 0.65$, $\beta = 1.005$, $\varepsilon_f = 0.05$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S4, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 2$, $\rho_{\text{power}} = 1$, and $K = 500$. PA24 uses $R = 1$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses $\delta_0 = 0.325$,

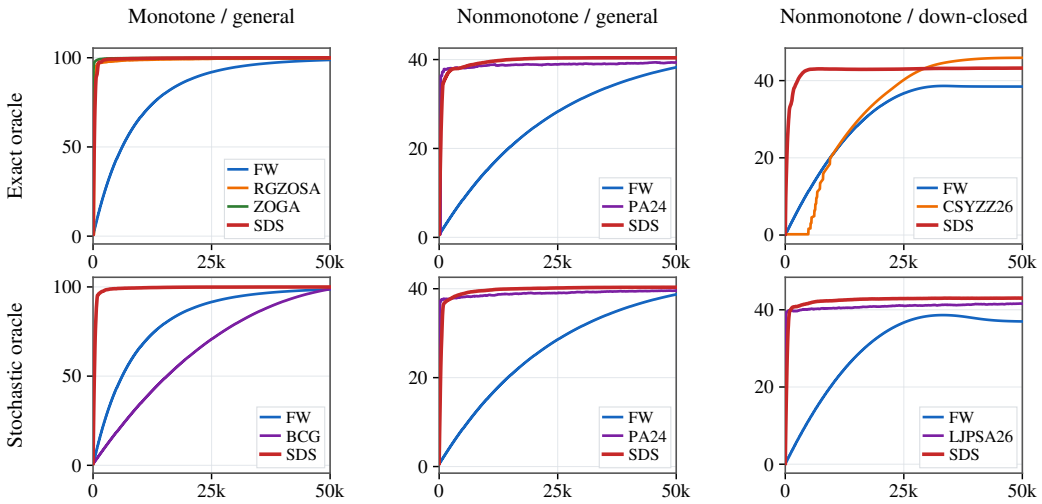


Figure 11: Average current objective value versus value-oracle calls for the regular coverage problem with $n = 100$.

Table 2: Percentage of completed SDS iterations that enter the degenerate state before 50,000 value-oracle calls. Entries are averages of the per-instance percentages over 10 test instances.

Problem	S1	S2	S3	S4	S5	S6
Finite-sum quadratic, $n = 50$	0.8522	0.6785	0.0358	1.4123	16.8237	1.8167
Finite-sum quadratic, $n = 100$	0.5269	0.4191	0.0137	3.0752	16.7303	0.7158
Finite-sum quadratic, $n = 300$	10.0604	12.9983	0.0993	0.3340	0.0000	0.0000
Single quadratic, $n = 100$	0.4141	0.5419	0.0495	3.5290	0.1625	0.1195
Regular coverage, $n = 100$	0.6942	1.8291	1.7821	1.6756	0.0000	1.6445

$\delta_{\max} = 0.325$, $\theta = 0.7$, $\beta = 1.005$, $\varepsilon_f = 0.3$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S5, FW uses $B = 50$, $T = 1$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 0.5$, and $K = 500$. CSYZZ26 uses $B = 40$, $\rho_{\text{scale}} = 2$, and $\rho_{\text{offset}} = 32$. SDS uses $\delta_0 = 0.35$, $\delta_{\max} = 0.35$, $\theta = 0.9$, $\beta = 1.00125$, $\varepsilon_f = 2$, and $p_d = 0.999$, with one polling sample and one gradient sample.

For S6, FW uses $B = 10$, $T = 4$, $\rho_{\text{scale}} = 0.5$, $\rho_{\text{power}} = 2/3$, and $K = 618$. LJPSA26 uses $R = 1$, $B = 40$, $T = 1$, $\eta_{\text{scale}} = 0.3$, and $\eta_{\text{power}} = 0.5$, with two-point CRN estimates. SDS uses $\delta_0 = 0.325$, $\delta_{\max} = 0.325$, $\theta = 0.5$, $\beta = 1.01$, $\varepsilon_f = 0.25$, and $p_d = 0.999$, with one polling sample and one gradient sample.

Degenerate states. Table 2 reports the fraction of completed SDS iterations that enter the degenerate branch before 50,000 value-oracle calls. For each setting, we compute the fraction separately for each test instance and average the 10 fractions. Degenerate states occur in less than 1% of the iterations in 18 of the 30 settings and in less than 5% in 26 settings. The larger fractions occur in a few finite-sum quadratic settings.

Overall observations. Overall, the experiments support the main idea of SDS: useful search directions can often be obtained from local geometry without estimating a full zeroth-order gradient, leading to strong value-oracle efficiency. The gains are especially clear in the monotone settings. One possible reason is that the monotone case reports the internal iterate directly, while the nonmonotone cases apply an output transformation before evaluating f. Finally, degenerate states occur rarely in most settings, so SDS operates mainly as a direct-search method in practice.