STOCHASTIC GRADIENT DESCENT ASCENT IS SUBOPTIMAL FOR NONCONVEX-PL MIN-MAX GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

How far can stochastic gradient descent ascent (SGDA) go by tuning its timescale ratio and step sizes in nonconvex min-max games? We answer this question for nonconvex–PL (NC-PL) games by establishing the first tight complexity of two-timescale SGDA with a fixed timescale ratio and non-increasing step sizes. For ℓ-smooth games with an inner μ-PL inequality and fixed initial-gap budgets, we prove a complexity lower bound $\Omega(\kappa^2\ell\varepsilon^{-2} + \kappa^4\ell\sigma^2\varepsilon^{-4})$, where $\kappa = \ell/\mu$ is the condition number, σ^2 is the gradient variance, and ε measures the outer gradient norm. This matches existing SGDA upper bounds and establishes a complexity separation from Smoothed-AGDA (Yang et al., 2022). In addition, we show that SGDA can fail to find a stationary point when its timescale ratio is as small as $o(\kappa^2)$. Our negative results highlight the fundamental limitation of SGDA in NC-PL games, and justify the development of alternative methods.

1 INTRODUCTION

Nonconvex min-max games have become a cornerstone of modern machine learning, underpinning generative adversarial networks (GANs) (Goodfellow et al., 2014), distributionally robust learning (Sinha et al., 2018), and language model alignment (Liu et al., 2026; Xu et al., 2025). Without additional structure, however, even finding a stationary point of a min-max game can be computationally intractable (Daskalakis et al., 2021; Bernasconi et al., 2026). For instance, games often exhibit limit cycles (Hsieh et al., 2021) and chaotic behavior (Sato et al., 2002). Such pathologies motivate the study of *structured* nonconvex min-max games (Lin et al., 2020; 2025; Ha, 2026). In particular, nonconvex–PL (NC-PL) games assume a Polyak–Łojasiewicz (PL) inequality (Polyak, 1963) on the inner maximization problem and have been actively studied as one of the weakest forms of structure (Nouiehed et al., 2019; Yang et al., 2022; Cho and Yun, 2023; Ha, 2026; Cai et al., 2026).

In NC-PL games, finding a Nash equilibrium (Nash, 1951) can be still intractable; hence, we seek a stationary point of the outer function $\Phi(x) := \max_y f(x, y)$. Two-timescale SGDA is a direct analog of SGD for such min-max games: it uses a faster ascent step to track the outer gradient $\nabla\Phi(x)$ and performs a slower inexact descent on Φ. For fixed initial-gap budgets, ℓ-smooth nonconvex–μ-PL games with condition number $\kappa = \ell/\mu$ and gradient variance σ^2 admit an SGDA complexity upper bound $O(\kappa^2\ell\varepsilon^{-2} + \kappa^4\ell\sigma^2\varepsilon^{-4})$ for finding an ε-stationary point (Yang et al., 2022; Ha, 2026). Meanwhile, recent studies have developed more sophisticated techniques to attain sharper bounds (Yang et al., 2022; Huang et al., 2025; Cai et al., 2026). This raises the following natural question:

Can SGDA obtain better complexity by tuning its timescale ratio and step sizes?

Contributions. In this work, we answer this question negatively. Our contributions are twofold.

(i) **Suboptimality of SGDA.** We prove matching lower bounds for two-timescale SGDA with a fixed timescale ratio and nonincreasing step sizes, and establish its suboptimality in NC-PL games.

(ii) **Optimal timescale ratio.** Additionally, we prove that SGDA with a small timescale ratio $o(\kappa^2)$ can fail to find a stationary point, and that existing $\Theta(\kappa^2)$ timescale ratio is optimal.

Sections 2–4 cover related work, the setup, and main results; Sections 5–6 prove the lower bounds; we discuss implications in Section 7 and conclude in Section 8.

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.

2 Related work

Nonconvex games are computationally intractable in general (Daskalakis et al., 2021), and even finding an approximate stationary point can be hard as well (Bernasconi et al., 2026). Such hardness results motivate structured games, including nonconvex–PL (NC-PL) (Nouiehed et al., 2019; Yang et al., 2020; 2022; Cho and Yun, 2023) and nonconvex–concave (NC-C) games (Thekumparampil et al., 2019; Lin et al., 2020; 2025; Ha, 2026), which assume a PL inequality and concavity on the inner maximization problem, respectively. Notably, these structures make finding stationary points computationally tractable and admit inexact-descent interpretations (Nouiehed et al., 2019; Lin et al., 2020; 2025; Ha, 2026). In particular, Nouiehed et al. (2019) showed that taking multiple ascent steps per descent step suffices to find an approximate stationary point in NC-PL games, and interpreted these updates as inexact descent on $\Phi(x) = \max_y f(x, y)$. Lin et al. (2020; 2025) extended this perspective to two-timescale SGDA for nonconvex–strongly-concave (NC–SC) and NC-C games. On the other hand, weakly-convex–weakly-concave games (Liu et al., 2021; Grimmer et al., 2023; Hajizadeh et al., 2024) and nonmonotone variational inequalities (Mertikopoulos et al., 2019; Diakonikolas et al., 2021; Lee and Kim, 2021; Pethick et al., 2022; Böhm, 2023) have been studied as well, focusing on extragradient and optimistic gradient (OGDA), rather than SGDA.

NC-PL games. Yang et al. (2020) gave the first convergence analysis of two-timescale alternating gradient descent ascent (GDA) for NC-PL games. Subsequently, Yang et al. (2022) extended the analysis to stochastic settings and introduced Smoothed-AGDA, attaining $\widetilde{O}(\kappa\ell\Delta\varepsilon^{-2} + \kappa^2\ell\Delta\sigma^2\varepsilon^{-4})$ under the comparison conditions in Section 4. Meanwhile, Huang et al. (2025) attained an improved $\widetilde{O}((\kappa^3\ell\Delta\sigma + \kappa^4\sigma^3)\varepsilon^{-3})$ leading stochastic complexity using momentum under Lipschitz sample gradient assumptions; Cai et al. (2026) obtained the same dependence on ε using Hessian-vector products to correct momentum bias. For finite-sum NC-PL games, Chen et al. (2022) and Cho and Yun (2023) obtained better complexities through variance reduction and sampling without replacement, respectively. Recent studies established high-probability guarantees for Smoothed-AGDA (Laguel et al., 2024) and SGDA (Ha, 2026) as well.

Two-timescale SGDA exposes a timescale ratio hyperparameter ρ that controls the size of ascent steps relative to descent steps. Fiez and Ratliff (2021) showed that a sufficiently large timescale ratio ρ yields local linear convergence of simultaneous GDA to differential Stackelberg equilibria. For NC–SC games, Li et al. (2022) proved that $\rho = \Theta(\kappa)$ suffices for local convergence, whereas GDA with $\rho = o(\kappa)$ can fail to converge locally. On the other hand, existing global convergence guarantees for GDA and SGDA in NC-PL games rely on a timescale ratio $\rho = \Theta(\kappa^2)$ (Fiez et al., 2021; Yang et al., 2022; Ha, 2026). Cho and Yun (2023) showed that an excessively large timescale ratio can worsen convergence rates of GDA in strongly-convex–strongly-concave (SC-SC) games, and raised the question of whether $\rho = \Theta(\kappa^2)$ is necessary for global convergence in NC-PL games.

Lower bounds for structured nonconvex min-max games have mainly addressed the minimax complexity of first-order methods. For instance, in NC–SC games, Zhang et al. (2021) proved a deterministic $\Omega(\sqrt{\kappa}\ell\Delta\varepsilon^{-2})$ lower bound, and Li et al. (2021) proved a stochastic $\Omega(\sqrt{\kappa}\ell\Delta\varepsilon^{-2} + \kappa^{1/3}\ell\Delta\sigma^2\varepsilon^{-4})$ bound. Pan and Li (2026) proved a sharper deterministic NC-PL lower bound $\Omega(\kappa\ell\Delta\varepsilon^{-2})$, establishing the minimax optimality of Smoothed-AGDA's leading complexity. In this work, we derive SGDA-specific lower bounds that match existing upper bounds (Yang et al., 2022; Ha, 2026) in leading κ, ε dependence and establish provable suboptimality of SGDA in NC-PL games.

3 Problem setup

We study the following min-max optimization problem:

$$\min_{x\in\mathbb{R}^{d_x}} \max_{y\in\mathbb{R}^{d_y}} f(x, y), \tag{3.1}$$

where f satisfies a smoothness assumption and $f(x, \cdot)$ satisfies a PL inequality (Karimi et al., 2016).

Assumption 3.1 (Smoothness)**.** The game $f : \mathbb{R}^{d_x} \times \mathbb{R}^{d_y} \to \mathbb{R}$ is differentiable and ℓ-smooth:

$$\|\nabla_x f(x_1, y_1) - \nabla_x f(x_2, y_2)\| \le \ell\big(\|x_1 - x_2\| + \|y_1 - y_2\|\big), \quad \forall x_1, x_2 \in \mathbb{R}^{d_x},\ y_1, y_2 \in \mathbb{R}^{d_y},$$
$$\|\nabla_y f(x_1, y_1) - \nabla_y f(x_2, y_2)\| \le \ell\big(\|x_1 - x_2\| + \|y_1 - y_2\|\big), \quad \forall x_1, x_2 \in \mathbb{R}^{d_x},\ y_1, y_2 \in \mathbb{R}^{d_y}.$$

We denote the outer objective of the game by $\Phi(x) := \max_{y\in\mathbb{R}^{d_y}} f(x,y)$, and let $\Phi^* := \inf_{x\in\mathbb{R}^{d_x}} \Phi(x)$. Define the condition number $\kappa := \ell/\mu \geq 1$, where $\mu > 0$ is the PL parameter.

Assumption 3.2 (NC-PL)**.** The outer objective satisfies $\Phi(x) < \infty$ for every $x \in \mathbb{R}^{d_x}$, and $Y^\star(x) := \arg\max_{y\in\mathbb{R}^{d_y}} f(x,y)$ is nonempty for every $x \in \mathbb{R}^{d_x}$. There exists $\mu > 0$ such that the following inequality holds for each $x \in \mathbb{R}^{d_x}$:

$$\frac{1}{2}\left\|\nabla_y f(x,y)\right\|^2 \geq \mu\big(\Phi(x) - f(x,y)\big), \qquad \forall y \in \mathbb{R}^{d_y}. \tag{3.2}$$

Assumption 3.3 (Initialization)**.** There exists $\Delta > 0$ and a universal constant $c_{\mathrm{init}} \geq 1$ such that

$$0 \leq \Phi(x_0) - \Phi^* \leq \Delta, \qquad 0 \leq \Phi(x_0) - f(x_0,y_0) \leq c_{\mathrm{init}}\Delta. \tag{3.3}$$

Throughout the paper, we make the following standard stochastic gradient estimator assumptions.

Assumption 3.4 (Unbiasedness)**.** The gradient estimators G_x, G_y are measurable and use fresh independent samples ξ from a fixed distribution. For every (x,y),

$$\mathbb{E}G_x(x,y;\xi) = \nabla_x f(x,y), \qquad \mathbb{E}G_y(x,y;\xi) = \nabla_y f(x,y). \tag{3.4}$$

Assumption 3.5 (Bounded variance)**.** For a fixed $\sigma \geq 0$ and every (x,y),

$$\mathbb{E}\left\|G_x(x,y;\xi) - \nabla_x f(x,y)\right\|^2 \leq \sigma^2, \qquad \mathbb{E}\left\|G_y(x,y;\xi) - \nabla_y f(x,y)\right\|^2 \leq \sigma^2. \tag{3.5}$$

Two-timescale SGDA. SGDA with a fixed timescale ratio $\rho \geq 1$ and step sizes $\eta_t \geq 0$ follows:

$$x_{t+1} = x_t - \frac{\eta_t}{\rho}G_x(x_t,y_t;\xi_t^x), \qquad y_{t+1} = y_t + \eta_t G_y(x_t',y_t;\xi_t^y), \tag{3.6}$$

where we use $x_t' = x_{t+1}$ for $\mathrm{SGDA}_{\mathrm{Alt}}$, and $x_t' = x_t$ for $\mathrm{SGDA}_{\mathrm{Sim}}$.

Minimax complexity. Let $\mathcal{C}(\ell,\kappa,\Delta,\sigma)$ collect all NC-PL game instances that satisfy Assumptions 3.1–3.5 with $d_x, d_y < \infty$. For each SGDA update order $\mathsf{o} \in \{\mathrm{Alt}, \mathrm{Sim}\}$, we define the minimax convergence rate of SGDA with a fixed timescale ratio ρ and nonincreasing learning rates η_t as

$$\mathcal{R}_{T,\mathsf{o}}(\rho) := \inf_{\eta_0\geq\cdots\geq\eta_{T-2}\geq 0} \sup_{I\in\mathcal{C}(\ell,\kappa,\Delta,\sigma)} \frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla\Phi(x_t)\right\|^2. \tag{3.7}$$

We write the timescale ratio tuned version as $\mathcal{R}^\star_{T,\mathsf{o}} := \inf_{\rho\geq 1}\mathcal{R}_{T,\mathsf{o}}(\rho)$. We emphasize that we allow parameter-dependent tuning of timescale ratio $\rho(\ell,\kappa,\Delta,\sigma,T)$ and learning rates $\eta_t(\ell,\kappa,\Delta,\sigma,T)$. We define the corresponding iteration complexities as

$$N_{\varepsilon,\mathsf{o}}(\rho) := \inf\left\{T \in \mathbb{N}_{\geq 1} : \mathcal{R}_{T,\mathsf{o}}(\rho) \leq \varepsilon^2\right\}, \qquad N^\star_{\varepsilon,\mathsf{o}} := \inf\left\{T \in \mathbb{N}_{\geq 1} : \mathcal{R}^\star_{T,\mathsf{o}} \leq \varepsilon^2\right\}. \tag{3.8}$$

4 Main Results

Our main result establishes a tight lower bound for two-timescale SGDA with a fixed timescale ratio $\rho \geq 1$ and non-increasing step size schedules η_t.

Theorem 4.1 (A lower bound for two-timescale SGDA in NC-PL games)**.** *There are universal constants $c > 0$ and $\kappa_0 \geq 2$, depending on c_{init}, such that, for every $\kappa \geq \kappa_0$, $\ell > 0$, $\Delta > 0$, $\sigma \geq 0$, and integer $T \geq 1$,*

$$\mathcal{R}^\star_{T,\mathrm{Sim}}, \mathcal{R}^\star_{T,\mathrm{Alt}} \geq c\min\left\{\ell\Delta, \max\left\{\frac{\kappa^2\ell\Delta}{T}, \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right\}\right\}. \tag{4.1}$$

Intuitively, the first term $\ell\Delta$ inside the minimum accounts for short horizons. Hence, in the usual small-error regime $\varepsilon^2 \leq c\ell\Delta$, Theorem 4.1 and the matching upper bounds for $\mathrm{SGDA}_{\mathrm{Alt}}$ from Yang et al. (2022) and for $\mathrm{SGDA}_{\mathrm{Sim}}$ in Theorem A.1 yield a tight complexity:

$$N^\star_{\varepsilon,\mathrm{Sim}}, N^\star_{\varepsilon,\mathrm{Alt}} \asymp \frac{\kappa^2\ell\Delta}{\varepsilon^2} + \frac{\kappa^4\ell\Delta\sigma^2}{\varepsilon^4}. \tag{4.2}$$

Table 1: **Upper and lower bounds in NC-PL games.** Our lower bounds match SGDA's upper bounds after optimizing the fixed timescale ratio and predetermined nonincreasing step sizes. Theorem 6.3 further rules out a uniform convergence guarantee for fixed ratios $\rho = o(\kappa^2)$, even with exact gradients; $\Omega(1)$ denotes the resulting nonvanishing stationarity error. Appendix F details the differing assumptions, stationarity criteria, and regimes of the displayed complexities.

Algorithm	Citation	Feedback	Complexity / obstruction
SGDA_{Alt}	Yang et al. (2022)	Bounded σ^2	$O(\kappa^2\ell\varepsilon^{-2} + \kappa^4\ell\sigma^2\varepsilon^{-4})$
SGDA_{Sim}	Theorem A.1	Bounded σ^2	$O(\kappa^2\ell\varepsilon^{-2} + \kappa^4\ell\sigma^2\varepsilon^{-4})$
Smoothed-AGDA	Yang et al. (2022)	Bounded σ^2	$\widetilde{O}(\boldsymbol{\kappa}\ell\varepsilon^{-2} + \boldsymbol{\kappa^2}\ell\sigma^2\varepsilon^{-4})$
SPIDER-GDA	Chen et al. (2022)	Finite-sum	$O(n + \sqrt{n}\kappa^2\ell\varepsilon^{-2})$
SGDA-RR	Cho and Yun (2023)	Finite-sum	$O(n\kappa^2\ell\varepsilon^{-2} + \sqrt{n}\kappa^3\ell\sigma\varepsilon^{-3})$
MSGDA	Huang et al. (2025)	Smooth samples	$\widetilde{O}((\kappa^3\ell\sigma + \kappa^4\sigma^3)\boldsymbol{\varepsilon^{-3}})$
AdaMSGDA	Huang et al. (2025)	Smooth samples	$\widetilde{O}((\kappa^3\ell\sigma + \kappa^4\sigma^3)\boldsymbol{\varepsilon^{-3}})$
HCMM-2	Cai et al. (2026)	Hessian products	$O_\ell(\kappa^3\sigma^3\boldsymbol{\varepsilon^{-3}})$
SGDA_{Alt}, SGDA_{Sim}	**Theorem 4.1**	**Bounded $\boldsymbol{\sigma^2}$**	$\Theta(\boldsymbol{\kappa^2}\ell\varepsilon^{-2} + \boldsymbol{\kappa^4}\ell\sigma^2\boldsymbol{\varepsilon^{-4}})$
Linear-span (NC–SC)	Zhang et al. (2021)	Deterministic	$\Omega(\sqrt{\kappa}\ell\varepsilon^{-2})$
Linear-span (NC–SC)	Zhang et al. (2021)	Finite-sum	$\Omega(n + \sqrt{n\kappa}\ell\varepsilon^{-2})$
First-order (NC–SC)	Li et al. (2021)	Bounded σ^2	$\Omega(\sqrt{\kappa}\ell\varepsilon^{-2} + \kappa^{1/3}\ell\sigma^2\varepsilon^{-4})$
First-order (NC-PL)	Pan and Li (2026)	Deterministic	$\Omega(\kappa\ell\varepsilon^{-2})$
SGDA_{Sim}	Li et al. (2022)	Deterministic	Local failure if $\rho = o(\kappa)$
SGDA_{Alt}, SGDA_{Sim}	**Theorem 6.3**	**Deterministic**	$\Omega(1)$ if $\rho = o(\kappa^2)$
SGDA_{Alt}, SGDA_{Sim}	**Theorem 5.1**	**Bounded $\boldsymbol{\sigma^2}$**	$\Omega(\rho\ell\varepsilon^{-2} + \kappa^4\ell\sigma^2\varepsilon^{-4})$

Suboptimality of SGDA. Crucially, Theorem 4.1 establishes the provable suboptimality of SGDA in NC-PL games: it yields a tight sample complexity $\Theta(\kappa^2\ell\Delta\varepsilon^{-2} + \kappa^4\ell\Delta\sigma^2\varepsilon^{-4})$ for SGDA, which is suboptimal compared with the complexity upper bound $O(\kappa\ell\Delta\varepsilon^{-2} + \kappa^2\ell\Delta\sigma^2\varepsilon^{-4})$ of Smoothed-AGDA (Yang et al., 2022).[1] Consequently, our negative result justifies the recent development of alternative methods such as Smoothed-AGDA (Yang et al., 2022) or momentum methods (Huang et al., 2025; Cai et al., 2026).

Optimal timescale ratio and step-size schedule. Notably, every existing analysis of two-timescale SGDA in NC-PL games relies on the choice $\rho = \Theta(\kappa^2)$ (Fiez et al., 2021; Yang et al., 2022; Cho and Yun, 2023; Ha, 2026), but the question whether it is actually optimal remained open. Together with the matching upper bounds (Yang et al., 2022; Ha, 2026), Theorem 4.1 shows that the choice of the ratio $\rho = \Theta(\kappa^2)$ and the constant schedule $\eta_t \equiv c_\eta \min\{1/\ell, \sqrt{\Delta/(\ell\sigma^2 T)}\}$ (with second term $+\infty$ for $\sigma = 0$) is optimal for SGDA with a fixed timescale ratio and non-increasing step sizes.

In the following sections, we study two distinctive regimes of the timescale ratio, namely $\rho = \Omega(\kappa^2)$ and $\rho = O(\kappa^2)$, and explain the intuition behind the hard instances that establish Theorem 4.1. We defer the proof of Theorem 4.1 to Appendix A.1.

5 The High-Ratio Lower Bound

In this section, we establish the lower bound for large timescale ratios $\rho = \Omega(\kappa^2)$. In this high-ratio regime, smaller descent steps slow down the outer progress, whereas larger ascent steps increase noise-induced errors that accumulate on the outer variable. These complementary obstructions prevent SGDA from obtaining sharper convergence rates.

[1] The Smoothed-AGDA bound includes the output conversion in Appendix E.

Theorem 5.1 (High-ratio lower bound). *There are universal $c, c_{\text{high}} > 0$ and $\kappa_{\text{high}} \geq 2$ such that, for every $\ell, \Delta > 0$, $\sigma \geq 0$, integer $T \geq 1$, $\kappa \geq \kappa_{\text{high}}$, and $\rho \geq c_{\text{high}}\kappa^2$, both SGDA$_{Sim}$ and SGDA$_{Alt}$ satisfy*

$$\mathcal{R}_{T,\text{Sim}}(\rho), \mathcal{R}_{T,\text{Alt}}(\rho) \geq c \min\left\{\ell\Delta, \max\left\{\frac{\rho\ell\Delta}{T}, \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right\}\right\}. \tag{5.1}$$

Together with the matching upper bounds from Yang et al. (2022) and Theorem A.1, Theorem 5.1 shows that the smallest order $\rho = \Theta(\kappa^2)$ is optimal in the high-ratio regime, as larger ratios only increase the lower bound. This suggests that the standard choice $\rho = \Theta(\kappa^2)$ (Yang et al., 2022) may be actually optimal among all fixed ratios. We establish this optimality in Section 6 by ruling out a uniform convergence guarantee for timescale ratios of smaller order $\rho = o(\kappa^2)$.

5.1 The instance for unstable step sizes

Consider a simple linear quadratic game $f_{\text{unstable}} : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ and its outer objective:

$$f_{\text{unstable}}(x; y) = \frac{\ell + \mu}{4}(xy - y^2), \qquad \Phi_{\text{unstable}}(x) = \frac{\ell + \mu}{16}x^2. \tag{5.2}$$

This game is ℓ-smooth and μ-strongly concave in y. For a fixed x, an inner ascent step becomes unstable for $\eta_t \geq 8/\ell$. Let m count such initial steps. Since the outer step is proportional to $-y$, inner instability can push x away from the minimizer at zero, despite the strongly convex outer objective.

Lemma 5.2 (Quadratic lower bound for large dual steps). *There is a universal $c > 0$ such that, for $\ell, \Delta > 0$, $\kappa \geq 2$, $\sigma \geq 0$, $\rho \geq 1$, integers $T \geq 1$ and $0 \leq m < T$, and predetermined nonnegative non-increasing step sizes with $\eta_t \geq 8/\ell$ for $t < m$, each of SGDA$_{Sim}$ and SGDA$_{Alt}$ admits an initialization (x_0, y_0) making f_{unstable} with $\mu = \ell/\kappa$ an exact-gradient instance in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with*

$$\frac{1}{T}\sum_{t=0}^{T-1} |\Phi'_{\text{unstable}}(x_t)|^2 \geq c\frac{m+1}{T}\ell\Delta. \tag{5.3}$$

For $m + 1 \geq T/2$, Lemma 5.2 immediately proves Theorem 5.1 with constant lower bound $c\ell\Delta$. We construct a distinct hard instance f_{stable} for $m + 1 < T/2$ in the next subsection.

5.2 The instance for stable step sizes

For the remaining case $m+1 < T/2$, we design an instance f_{stable} with outer coordinates $x = (u, v)$:

$$f_{\text{stable}}(u, v; y) = g_{\text{corridor}}(u) + g_\sigma(v, y), \qquad \Phi_{\text{stable}}(u, v) = g_{\text{corridor}}(u) + \Phi_\sigma(v), \tag{5.4}$$

where we use a scalar coordinate $u \in \mathbb{R}$ for the deterministic scalar function g_{corridor}, and define the outer objective $\Phi_\sigma(v) := \max_y g_\sigma(v, y)$ for the noisy function $g_\sigma(v, y)$. Intuitively, the corridor function g_{corridor} limits progress from small descent steps, and the noisy function g_σ transfers the inner stochastic perturbations to the outer variable. The stable ascent steps sum to $S = \sum_{t=m}^{T-2} \eta_t$.

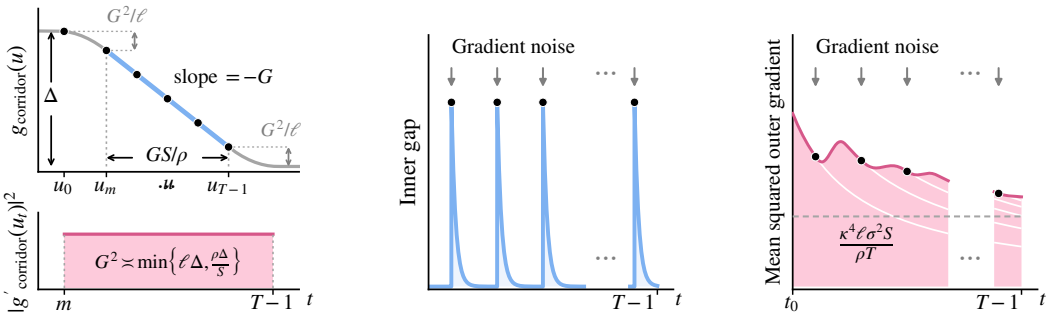
5.2.1 The corridor function g_{corridor}

The following lower bound shows that the corridor function g_{corridor} keeps the remaining primal gradients constant and nonzero.

Lemma 5.3 (Corridor lower bound). *For $\ell, \Delta > 0$, $\rho \geq 1$, integers $T \geq 1$, $0 \leq m < T$, and nonnegative steps with $S = \sum_{t=m}^{T-2} \eta_t$, there is an ℓ-smooth corridor $g_{\text{corridor}} : \mathbb{R} \to \mathbb{R}$ with gap Δ at zero and slope $G > 0$, where $G^2 \asymp \min\{\ell\Delta, \rho\Delta/S\}$. Gradient descent from zero with steps η_t/ρ satisfies $g'_{\text{corridor}}(u_t) < 0$ for all $t < T$ and $g'_{\text{corridor}}(u_t) = -G$ for $m \leq t < T$. Consequently,*

$$\frac{1}{T}\sum_{t<T} |g'_{\text{corridor}}(u_t)|^2 \gtrsim \frac{T-m}{T}\min\left\{\ell\Delta, \frac{\rho\Delta}{S}\right\}. \tag{5.5}$$

We illustrate the corridor function g_{corridor} in Figure 1(a). Intuitively, it consists of an entry curve and a constant-slope corridor, followed by an exit curve. The entry curve brings the first m updates



(a) The corridor function (b) Inner suboptimality (c) Mean squared gradient

Figure 1: **The competing bounds in f_{stable} (schematic).** Let m count the initial large steps and $S = \sum_{t=m}^{T-2} \eta_t$. (a) The corridor keeps gradient magnitude G at $u_m, \dots, u_{T-1}$. (b) Gradient noise creates inner gap $\Phi_\sigma(v_t) - g_\sigma(v_t, y_t)$. (c) Pink bands show contributions to $t^{-1}\sum_{i<t} \|\nabla\Phi_\sigma(v_i)\|^2$; the initial increases are omitted. The dashed level is the bound's scale $\kappa^4\ell\sigma^2 S/(\rho T)$ for $S \geq C\rho/(\ell\kappa)$.

to the corridor, and the remaining iterates travel its length GS/ρ at slope $-G$, using gap $G^2 S/\rho$. The exit curve brings the slope to zero; with curvature of magnitude ℓ, the entry and exit curves together use gap of order G^2/ℓ. Hence $G^2(S/\rho + 1/\ell) \asymp \Delta$ gives the stated lower bound (Appendix C.2).

While Lemma 5.3 provides a lower bound for small cumulative ascent steps S, it weakens as S grows. We complement this by constructing the function g_σ that induces a lower bound that grows with S.

5.2.2 The lasting effect of inner perturbation

To construct a noisy function g_σ whose lower bound grows with the cumulative ascent steps S, we study an inner perturbation that causes a long-lasting outer error. To make exposition clear, we represent the inner perturbation by a deterministic shift $\delta > 0$ in a simple linear-quadratic game $g_\delta : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}$:

$$g_\delta(z; a, b) = \ell z a - \frac{\ell}{2\kappa}\left[a^2 + (a + b + \delta)^2\right] \tag{5.6}$$

with an outer variable z and inner variables a, b. Here, the inner maximizer $(\kappa z, -\kappa z - \delta)$ gives outer objective $\Phi_{g_\delta}(z) = \ell\kappa z^2/2$; hence, the inner perturbation δ does not affect the outer objective. From the origin, however, an ascent step follows the negative derivative $\partial_a g_\delta = -\ell\delta/\kappa$ and makes the inner variable a negative. A descent step then follows the derivative $-\partial_z g_\delta = -\ell a > 0$ and moves the outer variable z away from zero, which is the minimizer of the outer objective Φ_{g_δ}. Consequentially, the inner perturbation δ induces an outer error through GDA dynamics.

In the high-ratio regime $\rho = \Omega(\kappa^2)$, such an outer error takes longer to remove than to create. To see this clearly, let us measure the time by cumulative ascent step $\sum_{i<t} \eta_i$. The outer variable z moves slowly, so we may approximate its effect by fixing $z = 0$. Then, the inner gradients follow $\partial_a g_\delta = -(\ell/\kappa)(2a + b + \delta)$ and $\partial_b g_\delta = -(\ell/\kappa)(a + b + \delta)$. In particular, the gradients start at $-\ell\delta/\kappa$ and remain negative while $|a| + |b|$ is small relative to δ. Therefore, a reaches a negative constant fraction of δ on the cumulative ascent timescale κ/ℓ; then it returns to zero as b approaches $-\delta$. Substituting this response into $z_{i+1} - z_i = -\ell\eta_i a_i/\rho$ predicts the scale

$$\underbrace{\frac{\ell\delta}{\rho}}_{\text{outer speed}} \times \underbrace{\frac{\kappa}{\ell}}_{\text{formation time}} = \underbrace{\frac{\kappa\delta}{\rho}}_{\text{error size}}.$$

After the inner variables adapt enough to the perturbation δ, the inner maximizer gives the approximation $a \approx \kappa z$, so we obtain an outer contraction $z_{i+1} \approx (1 - \ell\kappa\eta_i/\rho) z_i$. This suggests that removing a fixed fraction of the error induced by an inner perturbation takes cumulative ascent step of order $\rho/(\ell\kappa)$, a factor ρ/κ^2 longer than its formation time. The following lemma captures this intuition and bounds the error on the outer variable z_t, and specifies the interval over which this error persists.

Lemma 5.4 (Persistent displacement under a fixed shift)**.** *There are universal constants $c, C, C_0 > 0$ such that, for $\ell, \delta > 0$, $\kappa \geq 1$, $\rho \geq C_0\kappa^2$, and any finite step-size sequence with $0 \leq \ell\eta_t \leq c\kappa$, both SGDA$_{Sim}$ and SGDA$_{Alt}$ with exact gradients on the linear quadratic game g_δ given by (5.6), initialized at $(0,0,0)$, satisfy $|z_t| \leq C\kappa\delta/\rho$ and $|a_t| + |b_t| \leq C\delta$ whenever $\sum_{i<t}\eta_i \leq c\rho/(\ell\kappa)$, and*

$$z_t \geq c\frac{\kappa\delta}{\rho} \quad \text{if} \quad C\frac{\kappa}{\ell} \leq \sum_{i<t}\eta_i \leq c\frac{\rho}{\ell\kappa}. \tag{5.7}$$

5.2.3 The noisy function g_σ

We now construct the noisy function g_σ with its coordinate-wise components $g_{\sigma,j}$, each of which implements g_δ through a stochastic noise oracle. We replace δ in (5.6) by the perturbation $\delta(s)$ controlled by a third inner coordinate s:

$$\delta(s) = \sqrt{\kappa}\,\eta_j\sigma\sin^2\left(\frac{\pi s}{2\eta_j\sigma}\right), \qquad \delta(0) = 0, \qquad \delta(\pm\eta_j\sigma) = \sqrt{\kappa}\,\eta_j\sigma. \tag{5.8}$$

A mean-zero s-gradient error $\pm\sigma$ moves s from zero to one of the symmetric peaks of $\delta(s)$, creating the same positive shift for either sign. At these peaks, $\delta'(s) = 0$, so later exact-gradient updates preserve the shift and reproduce the persistent outer error of g_δ.

Writing $v_t[j]$ for the outer iterate of component $g_{\sigma,j}$, Lemma 5.4 shows that the perturbation of order $\sqrt{\kappa}\eta_j\sigma$ produces an outer error $v_t[j] \gtrsim \kappa^{3/2}\eta_j\sigma/\rho$. Its outer objective $\Phi_{\sigma,j}$ has curvature of order $\ell\kappa$, so its gradient has magnitude at least order $\ell\kappa^{5/2}\eta_j\sigma/\rho$. If enough cumulative step remains after update j, the interval of length $\asymp \rho/(\ell\kappa)$ contains at least order $\rho/(\ell\kappa\eta_j)$ iterates because the nonincreasing step sizes are at most η_j. Hence

$$\sum_{t<T}\mathbb{E}|\Phi'_{\sigma,j}(v_t[j])|^2 \gtrsim \underbrace{\left(\frac{\ell\kappa^{5/2}\eta_j\sigma}{\rho}\right)^2}_{\text{squared outer gradient}} \times \underbrace{\frac{\rho}{\ell\kappa\eta_j}}_{\text{number of iterates}} = \frac{\ell\kappa^4\sigma^2\eta_j}{\rho}. \tag{5.9}$$

The next lemma captures this idea and produces a lower bound for the sum of expected squared outer gradients induced by one inner-gradient perturbation.

Lemma 5.5 (Contribution of one inner-gradient perturbation)**.** *There are universal $c, C, c_{\text{step}} > 0$ such that, under Theorem 5.1's parameter conditions with $\sigma > 0$, $T \geq 2$, any predetermined steps $\eta_0 \geq \cdots \geq \eta_{T-2} \geq 0$ and index $0 \leq j \leq T-2$ with $0 < \ell\eta_j \leq c_{\text{step}}\kappa$ admit an ℓ-smooth NC-(ℓ/κ)-PL component $g_{\sigma,j} : \mathbb{R}\times\mathbb{R}^3 \to \mathbb{R}$. Its quadratic outer objective $\Phi_{\sigma,j}$ has curvature $\Phi''_{\sigma,j} \asymp \ell\kappa$, and both gaps vanish at the origin. Started there with one equiprobable $\pm\sigma$ error in the third inner gradient at update j and exact gradients otherwise, both SGDA$_{Sim}$ and SGDA$_{Alt}$ generate outer iterates $v_t[j]$ satisfying*

$$\sum_{t<T}\mathbb{E}|\Phi'_{\sigma,j}(v_t[j])|^2 \geq c\frac{\ell\kappa^4\sigma^2\eta_j}{\rho} \quad \text{if} \quad \sum_{t=j+1}^{T-2}\eta_t \geq C\frac{\rho}{\ell\kappa}. \tag{5.10}$$

Next we sum these contributions over disjoint coordinates j, and construct a lower bound for g_σ.

5.2.4 Combining the two components

We give each step $j \in J := \{j : m \leq j \leq T-2\}$ coordinates $v[j] \in \mathbb{R}$, $y[j] \in \mathbb{R}^3$ and define

$$g_\sigma(v,y) = \sum_{j\in J} g_{\sigma,j}(v[j];y[j]), \qquad \|\nabla\Phi_\sigma(v)\|^2 = \sum_{j\in J}|\Phi'_{\sigma,j}(v[j])|^2. \tag{5.11}$$

Combining Lemma 5.3 and Lemma 5.5 produces the following lower bound for the instance f_{stable}.

Lemma 5.6 (Stable-step lower bound and admissibility)**.** *Under Theorem 5.1's conditions, let m count the steps $\eta_t \geq 8/\ell$ in a predetermined nonnegative non-increasing sequence, and assume $m+1 < T/2$. For each of SGDA$_{\text{Sim}}$ and SGDA$_{\text{Alt}}$, the game f_{stable} given by (5.4) and a fixed oracle can be chosen as an instance in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with initial outer gap Δ and inner gap zero at the origin. Writing $S = \sum_{t=m}^{T-2} \eta_t$ and $\rho\Delta/0 := +\infty$, its iterates satisfy, for a universal $c > 0$,*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla \Phi_{\text{stable}}(u_t, v_t)\right\|^2 \geq c \min\left\{\ell\Delta, \max\left\{\frac{\rho\Delta}{S}, \frac{\kappa^4 \ell \sigma^2 S}{\rho T}\right\}\right\}. \tag{5.12}$$

Figure 1 illustrates the tradeoff in Lemma 5.6: increasing S weakens the corridor bound $\rho\Delta/S$ but strengthens the stochastic bound $\kappa^4\ell\sigma^2 S/(\rho T)$. For $S > 0$, applying the AM–GM inequality and $S < 8T/\ell$ to Lemma 5.6 proves Theorem 5.1 as follows:

$$\max\left\{\frac{\rho\Delta}{S}, \frac{\ell\kappa^4\sigma^2 S}{\rho T}\right\} \geq \max\left\{\frac{\rho\ell\Delta}{8T}, \sqrt{\frac{\rho\Delta}{S}\frac{\ell\kappa^4\sigma^2 S}{\rho T}}\right\} = \max\left\{\frac{\rho\ell\Delta}{8T}, \sqrt{\frac{\ell\kappa^4\Delta\sigma^2}{T}}\right\}. \tag{5.13}$$

6 Low-ratio lower bound

In the low ratio regime $\rho = O(\kappa^2)$, inner variables may fail to follow their moving maximizer even with exact gradients. In particular, Mourrat (2026) constructed a PL-PL game that exhibits a periodic orbit of a continuous-time GDA flow, suggesting that a PL inequality can fail to prevent limit cycles of GDA. In this section, we realize a similar construction in NC-PL games.

6.1 Inner PL does not prevent a limit cycle

We start with a simple bilinear game that creates a periodic orbit for GDA:

$$g(u, a) = \sqrt{2}\, u^{\mathsf{T}} a, \qquad u, a \in \mathbb{R}^2. \tag{6.1}$$

For timescale ratio $\rho = 1$, GDA flow oscillates in this game (Hsieh et al., 2021), but the infinite inner maximum for $u \neq 0$ violates NC-PL. Adding $-\|a\|^2$ gives the outer objective $\|u\|^2/2$ yet damps the flow: $\ddot{u} + 2\dot{u} + 2u = 0$. Hence, we introduce $b \in \mathbb{R}^2$ and a residual $R : \mathbb{R}^2 \times \mathbb{R}^2 \to \mathbb{R}^2$ and define

$$g_{\text{cycle}}(u, a, b) = \sqrt{2}\, u^{\mathsf{T}} a - \|a\|^2 - \tfrac{1}{2}\|R(a, b)\|^2. \tag{6.2}$$

Designing the residual $R(a, b)$. We want the residual $R(a, b)$ to restore periodic motion while preserving an inner PL inequality and the outer objective $\|u\|^2/2$. If every a admits b that yields $R(a, b) = 0$, maximization cancels the residual and leaves the outer objective unchanged. On the other hand, the inner PL inequality follows when the map $(a, b) \mapsto (\sqrt{2}a, R(a, b))$ has a global smooth inverse with a uniformly bounded derivative (c.f. Lemma B.1). To specify the periodic motion we want, let J denote counterclockwise rotation by $\pi/2$, and define

$$u_\star(t) = (\cos t, \sin t)^{\mathsf{T}}, \qquad a_\star(t) = -Ju_\star(t)/\sqrt{2}, \qquad b_\star(t) = u_\star(t)/\sqrt{2}. \tag{6.3}$$

Our goal is to design the residual $R(a, b)$ for GDA to follow this path. Outer descent already has the required velocity, since $\dot{u}_\star = Ju_\star = -\sqrt{2}a_\star$. The inner variables must have velocities $\dot{a}_\star = Ja_\star$ and $\dot{b}_\star = -a_\star$. Hence, to obtain the desired inner motion, we require

$$\nabla_a\left(\tfrac{1}{2}\|R(a, b)\|^2\right) = -2a + Ja, \qquad \nabla_b\left(\tfrac{1}{2}\|R(a, b)\|^2\right) = a, \tag{6.4}$$

on the circle $b = Ja$, $\|a\| = 1/\sqrt{2}$ traced by the motion $(a_\star, b_\star)$. With these identities, the GDA equations agree with the target path. We choose further derivatives so that nearby trajectories approach the orbit, and use cutoffs for global smoothness without changing these dynamics (c.f. Appendix D.1).

Lemma 6.1 (An attracting periodic orbit and its discrete dynamics)**.** *There are universal $C, \mu, h_\star > 0$ and a smooth residual R such that the base game $g_{\text{cycle}} : \mathbb{R}^2 \times \mathbb{R}^4 \to \mathbb{R}$ in (6.2) is jointly C-smooth and affine in u, has outer objective $\Phi_{g_{\text{cycle}}}(u) = \frac{1}{2}\|u\|^2$, and satisfies the inner μ-PL inequality. The path (6.3), with $y_\star = (a_\star, b_\star)$, traces a locally stable 2π-periodic orbit Γ of GDA flow. There is a compact neighborhood $\mathcal{U}$ of Γ such that both SGDA$_{\text{Sim}}$ and SGDA$_{\text{Alt}}$, started in $\mathcal{U}$ with exact gradients and equal primal and dual steps $h_t \in [0, h_\star]$, satisfy $(u_t, y_t) \in \mathcal{U}$ and $\left\|\nabla \Phi_{g_{\text{cycle}}}(u_t)\right\| = \|u_t\| \geq \frac{1}{2}$ for all $t \geq 0$, where $y_t = (a_t, b_t)$.*

6.2 From the cycle to the low-ratio lower bound

Lemma 6.1 keeps the outer gradient bounded away from zero for sufficiently small steps. However, steps exceeding $h_\star$ can occur initially. The next lemma accommodates these initial steps by modifying the game away from the cycle, while keeping the outer variable fixed until the trajectory reaches it.

Lemma 6.2 (Large initial steps reach the cycle)**.** *There are universal $C, \mu > 0$ with the following property. For any finite sequence of steps $h_0, \dots, h_{m-1} > h_\star$, including $m = 0$, we can modify the game g_{cycle} in Lemma 6.1 to obtain $\tilde{g}_{\mathrm{cycle}} : \mathbb{R}^2 \times \mathbb{R}^5 \to \mathbb{R}$. This game is jointly C-smooth and affine in u, satisfies the inner μ-PL inequality, and has outer objective $\Phi_{\tilde{g}_{\mathrm{cycle}}}(u) = \|u\|^2/2$. With $y = (a, b, r)$ and $e_1 = (1, 0)^\mathsf{T}$, there is an initialization (e_1, y_0) with inner gap at most C. From this initialization, both SGDA$_{Sim}$ and SGDA$_{Alt}$, using exact gradients and equal primal and dual steps, keep $u_t = e_1$ for $0 \le t \le m$ and reach $y_m = (y_\star(0), 0)$. For every continuation with $h_t \in [0, h_\star]$, each algorithm follows its g_{cycle} dynamics in (u, a, b), while r remains zero.*

We now rescale $\tilde{g}_{\mathrm{cycle}}$ to satisfy ℓ-smoothness, the inner (ℓ/κ)-PL inequality, and the initial-gap budgets. We choose $f_{\mathrm{cycle}}(x, y) = A\tilde{g}_{\mathrm{cycle}}(x/s_x, y/s_y)$ with positive scales satisfying $s_y^2 = \rho s_x^2$, which converts both updates at ratio ρ to equal-step dynamics. The scales in Appendix A.3 make the game ℓ-smooth with inner PL constant of order $\ell/\sqrt{\rho}$, so $\rho \le c_{\mathrm{low}}\kappa^2$ ensures inner PL at least ℓ/κ. They also meet the initial-gap budgets and give $\|\nabla\Phi_{\mathrm{cycle}}(x_t)\|^2 \gtrsim \ell\Delta$ at every iterate, proving the following theorem.

Theorem 6.3 (Necessity of an order-κ^2 fixed ratio)**.** *There are universal $c, c_{\mathrm{low}} > 0$ and $\kappa_{\mathrm{low}} \ge 2$ such that, for every $\kappa \ge \kappa_{\mathrm{low}}$, $\ell, \Delta > 0$, $\sigma \ge 0$, integer $T \ge 1$, $1 \le \rho \le c_{\mathrm{low}}\kappa^2$, and predetermined nonnegative non-increasing step sizes, each of SGDA$_{Sim}$ and SGDA$_{Alt}$ admits a deterministic instance f_{cycle} in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with $d_x = 2$, $d_y = 5$, a strongly convex quadratic outer objective, and*

$$\min_{0 \le t < T} \|\nabla\Phi(x_t)\|^2 \ge c\ell\Delta. \tag{6.5}$$

The theorem answers the question raised by Cho and Yun (2023): is timescale ratio $\rho = \Omega(\kappa^2)$ necessary for SGDA to converge in NC-PL games? The theorem shows that smaller timescale ratios $\rho = o(\kappa^2)$ cannot provide such a guarantee, even with deterministic gradients and non-increasing step sizes. Together with Theorem 5.1, it yields our tight main lower bound in Theorem 4.1.

7 Discussion

Suboptimality of SGDA. In this work, we established the first provable suboptimality of SGDA in NC-PL games. Our high-ratio construction combines a corridor that limits outer progress with inner-gradient perturbations that create persistent outer errors. Meanwhile, our low-ratio lower bound rules out uniform convergence at ratios $\rho = o(\kappa^2)$ under the permitted step-size schedules. The lower bounds suggest a provable advantage of Smoothed-AGDA (Yang et al., 2022) over SGDA in NC-PL games.

Minimum timescale ratio. Notably, we also prove that uniform convergence within this tuning family requires $\rho = \Omega(\kappa^2)$, answering the question of Cho and Yun (2023). For $\rho = o(\kappa^2)$ and each prescribed horizon and permitted schedule, we construct a deterministic instance whose outer gradients stay bounded away from zero by rescaling a game with an attracting periodic orbit. This occurs even with a strongly convex quadratic outer objective: inner PL does not ensure tracking as the outer variable moves. Together with matching upper bounds (Yang et al., 2022), our results make $\rho = \Theta(\kappa^2)$ and constant steps optimal within this family.

8 Conclusion

We characterized SGDA's optimal complexity and necessary timescale separation in NC-PL games. Open directions include nonmonotone or signed schedules and extensions to nonconvex–concave games.

AI Use Statement

Generative AI tools were used to brainstorm and refine theoretical constructions, formulate and check proof strategies, survey relevant literature, draft and edit mathematical exposition, and prepare and check the LaTeX source. The authors are responsible for the final content, including all AI-assisted text, claims, proofs, and artifacts.

Reproducibility Statement

All assumptions and algorithms are stated in Section 3. Appendix A contains the proofs of the main theorem, the high- and low-ratio lower bounds, and the matching $\mathrm{SGDA}_{\mathrm{Sim}}$ upper bound. Appendices C and D prove the construction lemmas. Together, these proofs give the construction formulas, parameter choices, oracle definitions, initialization checks, and arguments for $\mathrm{SGDA}_{\mathrm{Sim}}$ and $\mathrm{SGDA}_{\mathrm{Alt}}$.

References

Martino Bernasconi, Matteo Castiglioni, Andrea Celli, and Alexandros Hollender. The complexity of min-max optimization for quadratic polynomials. *arXiv preprint arXiv:2606.17000*, 2026. URL https://arxiv.org/abs/2606.17000.

Axel Böhm. Solving nonconvex-nonconcave min-max problems exhibiting weak Minty solutions. *Transactions on Machine Learning Research*, 2023. URL https://openreview.net/forum?id=Gp0pHyUyrb.

Haoyuan Cai, Sulaiman A. Alghunaim, and Ali H. Sayed. Accelerated stochastic min-max optimization based on bias-corrected momentum. *arXiv preprint arXiv:2406.13041v3*, 2026. URL https://arxiv.org/abs/2406.13041v3.

Lesi Chen, Boyuan Yao, and Luo Luo. Faster stochastic algorithms for minimax optimization under Polyak–Łojasiewicz condition. In *Advances in Neural Information Processing Systems*, volume 35, 2022. URL https://arxiv.org/abs/2307.15868.

Hanseul Cho and Chulhee Yun. SGDA with shuffling: Faster convergence for nonconvex-PL minimax optimization. In *International Conference on Learning Representations*, 2023. URL https://iclr.cc/virtual/2023/poster/11501.

Constantinos Daskalakis, Stratis Skoulakis, and Manolis Zampetakis. The complexity of constrained min-max optimization. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, 2021. doi: 10.1145/3406325.3451125. URL https://arxiv.org/abs/2009.09623.

Jelena Diakonikolas, Constantinos Daskalakis, and Michael I. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *Proceedings of the 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 2746–2754. PMLR, 2021. URL https://proceedings.mlr.press/v130/diakonikolas21a.html.

Tanner Fiez and Lillian J. Ratliff. Local convergence analysis of gradient descent ascent with finite timescale separation. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=AWOSz_mMAPx.

Tanner Fiez, Lillian J. Ratliff, Eric Mazumdar, Evan Faulkner, and Adhyyan Narang. Global convergence to local minmax equilibrium in classes of nonconvex zero-sum games. In *Advances in Neural Information Processing Systems*, volume 34, 2021.

Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, 2014. URL https://arxiv.org/abs/1406.2661.

Benjamin Grimmer, Haihao Lu, Pratik Worah, and Vahab Mirrokni. The landscape of the proximal point method for nonconvex–nonconcave minimax optimization. *Mathematical Programming*, 201(1–2):373–407, 2023. doi: 10.1007/s10107-022-01910-8. URL https://doi.org/10.1007/s10107-022-01910-8.

Junsoo Ha. High probability convergence guarantees of stochastic gradient descent ascent in structured nonconvex min-max games. In *Proceedings of Thirty Ninth Conference on Learning Theory*, volume 336 of *Proceedings of Machine Learning Research*, pages 3005–3075, 2026. URL https://proceedings.mlr.press/v336/ha26a.html.

Saeed Hajizadeh, Haihao Lu, and Benjamin Grimmer. On the linear convergence of extragradient methods for nonconvex–nonconcave minimax problems. *INFORMS Journal on Optimization*, 6 (1):19–31, 2024. doi: 10.1287/ijoo.2022.0004. URL https://doi.org/10.1287/ijoo.2022.0004.

Ya-Ping Hsieh, Panayotis Mertikopoulos, and Volkan Cevher. The limits of min-max optimization algorithms: Convergence to spurious non-critical sets. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 4337–4348, 2021. URL https://proceedings.mlr.press/v139/hsieh21a.html.

Feihu Huang, Chunyu Xuan, Xinrui Wang, Siqi Zhang, and Songcan Chen. Enhanced adaptive gradient algorithms for nonconvex-PL minimax optimization. In *Proceedings of the 28th International Conference on Artificial Intelligence and Statistics*, volume 258 of *Proceedings of Machine Learning Research*, pages 3439–3447, 2025. URL https://proceedings.mlr.press/v258/huang25d.html.

Hamed Karimi, Julie Nutini, and Mark Schmidt. Linear convergence of gradient and proximal-gradient methods under the polyak–łojasiewicz condition. In *Machine Learning and Knowledge Discovery in Databases*, volume 9851 of *Lecture Notes in Computer Science*, pages 795–811. Springer, 2016. doi: 10.1007/978-3-319-46128-1_50.

Yassine Laguel, Yasa Syed, Necdet Serhat Aybat, and Mert Gürbüzbalaban. High-probability complexity bounds for stochastic non-convex minimax optimization. In *Advances in Neural Information Processing Systems*, volume 37, 2024. URL https://arxiv.org/abs/2405.14130.

Sucheol Lee and Donghwan Kim. Fast extra gradient methods for smooth structured nonconvex-nonconcave minimax problems. In *Advances in Neural Information Processing Systems*, volume 34, pages 22588–22600, 2021. URL https://proceedings.neurips.cc/paper/2021/hash/be767243ca8f574c740fb4c26cc6dceb-Abstract.html.

Haochuan Li, Yi Tian, Jingzhao Zhang, and Ali Jadbabaie. Complexity lower bounds for nonconvex-strongly-concave min-max optimization. In *Advances in Neural Information Processing Systems*, volume 34, 2021. URL https://proceedings.neurips.cc/paper_files/paper/2021/hash/0e105949d99a32ca1751703e94ece601-Abstract.html.

Haochuan Li, Farzan Farnia, Subhro Das, and Ali Jadbabaie. On convergence of gradient descent ascent: A tight local analysis. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pages 12717–12740, 2022. URL https://proceedings.mlr.press/v162/li22e.html.

Tianyi Lin, Chi Jin, and Michael I. Jordan. On gradient descent ascent for nonconvex-concave minimax problems. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 6083–6093, 2020. URL https://proceedings.mlr.press/v119/lin20a.html.

Tianyi Lin, Chi Jin, and Michael I. Jordan. Two-timescale gradient descent ascent algorithms for nonconvex minimax optimization. *Journal of Machine Learning Research*, 26(11):1–45, 2025. URL https://jmlr.org/papers/v26/22-0863.html.

Mingrui Liu, Hassan Rafique, Qihang Lin, and Tianbao Yang. First-order convergence theory for weakly-convex-weakly-concave min-max problems. *Journal of Machine Learning Research*, 22 (169):1–34, 2021. URL https://jmlr.org/papers/v22/20-533.html.

Qihao Liu, Luoxin Ye, Wufei Ma, Yu-Cheng Chou, and Alan Yuille. Generative adversarial reasoner: Enhancing LLM reasoning with adversarial reinforcement learning. In *International Conference on Learning Representations*, 2026. URL https://arxiv.org/abs/2512.16917.

Panayotis Mertikopoulos, Bruno Lecouat, Houssam Zenati, Chuan-Sheng Foo, Vijay Chandrasekhar, and Georgios Piliouras. Optimistic mirror descent in saddle-point problems: Going the extra (gradient) mile. In *International Conference on Learning Representations*, 2019. URL https://openreview.net/forum?id=Bkg8jjC9KQ.

Jean-Christophe Mourrat. PL conditions do not guarantee convergence of gradient descent-ascent dynamics. *arXiv preprint arXiv:2602.16517*, 2026. URL https://arxiv.org/abs/2602.16517.

John Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951. doi: 10.2307/1969529.

Maher Nouiehed, Maziar Sanjabi, Tianjian Huang, Jason D. Lee, and Meisam Razaviyayn. Solving a class of non-convex min-max games using iterative first order methods. In *Advances in Neural Information Processing Systems*, volume 32, 2019.

Siyu Pan and Jiajin Li. Lower bounds for nonconvex–PŁ minimax optimization. *arXiv preprint arXiv:2608.26799*, 2026. URL https://arxiv.org/abs/2608.26799.

Thomas Pethick, Puya Latafat, Panagiotis Patrinos, Olivier Fercoq, and Volkan Cevher. Escaping limit cycles: Global convergence for constrained nonconvex-nonconcave minimax problems. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=2_vhkAMARk.

Boris T. Polyak. Gradient methods for minimizing functionals. *USSR Computational Mathematics and Mathematical Physics*, 3(4):864–878, 1963. doi: 10.1016/0041-5553(63)90382-3.

Yuzuru Sato, Eizo Akiyama, and J. Doyne Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences*, 99(7):4748–4751, 2002. doi: 10.1073/pnas.032086299. URL https://pmc.ncbi.nlm.nih.gov/articles/PMC123719/.

Aman Sinha, Hongseok Namkoong, and John Duchi. Certifying some distributional robustness with principled adversarial training. In *International Conference on Learning Representations*, 2018. URL https://arxiv.org/abs/1710.10571v4.

Kiran K. Thekumparampil, Prateek Jain, Praneeth Netrapalli, and Sewoong Oh. Efficient algorithms for smooth minimax optimization. In *Advances in Neural Information Processing Systems*, volume 32, 2019. URL https://arxiv.org/abs/1907.01543.

Zaiyan Xu, Sushil Vemuri, Kishan Panaganti, Dileep Kalathil, Rahul Jain, and Deepak Ramachandran. Robust LLM alignment via distributionally robust direct preference optimization. In *Advances in Neural Information Processing Systems*, volume 38, 2025. URL https://arxiv.org/abs/2502.01930.

Junchi Yang, Negar Kiyavash, and Niao He. Global convergence and variance reduction for a class of nonconvex-nonconcave minimax problems. In *Advances in Neural Information Processing Systems*, volume 33, 2020. URL https://papers.neurips.cc/paper_files/paper/2020/hash/0cc6928e741d75e7a92396317522069e-Abstract.html.

Junchi Yang, Antonio Orvieto, Aurélien Lucchi, and Niao He. Faster single-loop algorithms for minimax optimization without strong concavity. In *Proceedings of the 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pages 5485–5517, 2022. URL https://proceedings.mlr.press/v151/yang22b.html.

Siqi Zhang, Junchi Yang, Cristóbal Guzmán, Negar Kiyavash, and Niao He. The complexity of nonconvex-strongly-concave minimax optimization. In *Proceedings of the Thirty-Seventh Conference on Uncertainty in Artificial Intelligence*, volume 161 of *Proceedings of Machine Learning Research*, pages 482–492, 2021. URL https://proceedings.mlr.press/v161/zhang21c.html.

APPENDIX CONTENTS

A Proof of main result **15**

A.1 Proof of Theorem 4.1 15

A.2 Proof of Theorem 5.1 15

A.3 Proof of Theorem 6.3 16

A.4 Upper bound for SGDA_{Sim} 17

B Preliminary lemmas **19**

B.1 NC-PL games with a predetermined outer function 19

B.2 NC-PL geometry 19

C Proofs for the high-ratio lower bound **20**

C.1 Proof of the lower bound for f_{unstable} 20

C.2 Construction of g_{corridor} 24

C.3 Construction of g_σ 25

C.4 Proof of the lower bound for f_{stable} 29

D Proofs for the low-ratio lower bound **31**

D.1 Construction of g_{cycle} 31

D.2 Construction of f_{cycle} 33

E Comparison with Smoothed-AGDA **36**

E.1 Initialization budget of the hard instances 37

F Comparison conditions for Table 1 **38**

A Proof of main result

We first combine the two ratio regimes to prove Theorem 4.1. We then prove Theorems 5.1 and 6.3 using the construction lemmas from Sections 5 and 6, and establish the matching upper bound for $\mathrm{SGDA}_{\mathrm{Sim}}$.

A.1 Proof of Theorem 4.1

Theorem 4.1 (A lower bound for two-timescale SGDA in NC-PL games)**.** *There are universal constants $c > 0$ and $\kappa_0 \geq 2$, depending on c_{init}, such that, for every $\kappa \geq \kappa_0$, $\ell > 0$, $\Delta > 0$, $\sigma \geq 0$, and integer $T \geq 1$,*

$$\mathcal{R}^{\star}_{T,\mathrm{Sim}}, \mathcal{R}^{\star}_{T,\mathrm{Alt}} \geq c \min\left\{\ell\Delta, \max\left\{\frac{\kappa^2\ell\Delta}{T}, \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right\}\right\}. \tag{4.1}$$

Proof. It remains to cover every ratio. Choose κ_0 above both theorem thresholds. The construction constants can be chosen so that $c_{\mathrm{high}} \leq c_{\mathrm{low}}$, as verified in the calibration below. In the high-ratio bound, $\rho\ell\Delta/T \geq c_{\mathrm{high}}\kappa^2\ell\Delta/T$; the stochastic term already matches (4.1). The low-ratio bound supplies the cap $c\ell\Delta$. Taking the ratio infimum proves (4.1).

For $\varepsilon^2 \leq c\ell\Delta$ with sufficiently small universal $c > 0$, inversion gives the lower complexity bound in (4.2). The matching constant-step upper bounds cited in Section 4 complete the proof. □

Calibrating the overlapping ratio regimes. The fixed cycle and its extension in Lemma 6.2 determine a universal inner-PL constant μ. Let C_{low} be the smoothness-normalization constant chosen in Appendix A.3; the low-ratio construction then applies up to $c_{\mathrm{low}} = (\mu/C_{\mathrm{low}})^2$. These constants are fixed before we calibrate the noise component. Let $\alpha_\star$ be the refined constant from the proof of Lemma 5.4, let c_0, C_0 be the constants from Lemma C.3, and choose

$$0 < \gamma \leq \min\left\{\sqrt{c_{\mathrm{low}}\alpha_\star}, \frac{c_0}{2C_0}\right\}, \qquad \theta := \frac{\gamma}{c_0}.$$

These choices calibrate only the high-ratio block: they give (C.40) and $c_{\mathrm{high}} = \gamma^2/\alpha_\star \leq c_{\mathrm{low}}$. Thus its ratio range begins no later than the endpoint of the low-ratio range. The low-ratio constants do not depend on γ or θ.

A.2 Proof of Theorem 5.1

Theorem 5.1 (High-ratio lower bound)**.** *There are universal $c, c_{\mathrm{high}} > 0$ and $\kappa_{\mathrm{high}} \geq 2$ such that, for every $\ell, \Delta > 0$, $\sigma \geq 0$, integer $T \geq 1$, $\kappa \geq \kappa_{\mathrm{high}}$, and $\rho \geq c_{\mathrm{high}}\kappa^2$, both $\mathrm{SGDA}_{\mathrm{Sim}}$ and $\mathrm{SGDA}_{\mathrm{Alt}}$ satisfy*

$$\mathcal{R}_{T,\mathrm{Sim}}(\rho), \mathcal{R}_{T,\mathrm{Alt}}(\rho) \geq c \min\left\{\ell\Delta, \max\left\{\frac{\rho\ell\Delta}{T}, \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right\}\right\}. \tag{5.1}$$

Proof. For each of $\mathrm{SGDA}_{\mathrm{Sim}}$ and $\mathrm{SGDA}_{\mathrm{Alt}}$ with predetermined non-increasing step sizes, let m count the initial excessively large steps $\ell\eta_t \geq 8$ and set $S = \sum_{t=m}^{T-2}\eta_t$. If $m + 1 \geq T/2$, Lemma 5.2 gives an average squared outer gradient at least $c\ell\Delta$ on f_{unstable}, proving the capped bound. Otherwise, Lemma 5.6 applies to f_{stable} and also gives the cap when $S = 0$. For $S > 0$, the remaining steps satisfy $S \leq 8T/\ell$, while the maximum of two nonnegative numbers is at least their geometric mean. Hence

$$\max\left\{\frac{\rho\Delta}{S}, \frac{\kappa^4\ell\sigma^2 S}{\rho T}\right\} \geq \max\left\{\frac{\rho\ell\Delta}{8T}, \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right\}. \tag{A.1}$$

Substituting into (5.12) and adjusting the universal constant proves (5.1). In each case, the instance is chosen after fixing the schedule. Taking the instance supremum and then the schedule infimum proves the theorem. □

A.3 Proof of Theorem 6.3

Theorem 6.3 (Necessity of an order-κ^2 fixed ratio)**.** *There are universal $c, c_{\mathrm{low}} > 0$ and $\kappa_{\mathrm{low}} \geq 2$ such that, for every $\kappa \geq \kappa_{\mathrm{low}}$, $\ell, \Delta > 0$, $\sigma \geq 0$, integer $T \geq 1$, $1 \leq \rho \leq c_{\mathrm{low}}\kappa^2$, and predetermined nonnegative non-increasing step sizes, each of SGDA$_{Sim}$ and SGDA$_{Alt}$ admits a deterministic instance f_{cycle} in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with $d_x = 2$, $d_y = 5$, a strongly convex quadratic outer objective, and*

$$\min_{0 \leq t < T} \|\nabla \Phi(x_t)\|^2 \geq c\ell\Delta. \tag{6.5}$$

Proof. We rescale the fixed-dimensional game while retaining its strongly convex quadratic outer objective, then verify smoothness, inner PL, and the initial-gap budgets.

Let $C \geq 1$ bound the smoothness constant, initial inner gap, and $\|\nabla_u \tilde{g}_{\mathrm{cycle}}(e_1, y_0)\|$ in Lemma 6.2; Appendix E.1 verifies the last bound. Choose a universal $C_{\mathrm{low}} \geq 2C$, and fix a sufficiently small universal $0 < \zeta \leq \min\{1/2, c_{\mathrm{init}}/(2C)\}$. For a ratio ρ, set

$$A := \frac{2\zeta\Delta}{\sqrt{\rho}}, \qquad s_x^2 := \frac{2C_{\mathrm{low}}\zeta\Delta}{\ell\rho}, \qquad s_y^2 := \frac{2C_{\mathrm{low}}\zeta\Delta}{\ell}. \tag{A.2}$$

For a fixed horizon and predetermined non-increasing step sizes (η_t), put $h_t = \ell\eta_t/(C_{\mathrm{low}}\sqrt{\rho})$. Let m count the initial steps with $h_t > h_\star$. Apply Lemma 6.2 to those steps, writing $\tilde{g}_{\mathrm{cycle}}$ for its game. For both SGDA$_{\mathrm{Sim}}$ and SGDA$_{\mathrm{Alt}}$, the lemma keeps $u_t = e_1$ through time m and then identifies the continuation with the small-step dynamics of g_{cycle}. Hence Lemma 6.1 keeps all later states in $\mathcal{U} \times \{0\}$, so $\|u_t\| \geq 1/2$ at every iterate. This includes $m = 0$, $m = T - 1$, $T = 1$, and zero later steps.

Define the game $f_{\mathrm{cycle}} : \mathbb{R}^2 \times \mathbb{R}^5 \to \mathbb{R}$ by

$$f_{\mathrm{cycle}}(x, y) := A\, \tilde{g}_{\mathrm{cycle}}\left(\frac{x}{s_x}, \frac{y}{s_y}\right), \tag{A.3}$$

initialized at $x_0 = s_x e_1$ and $y_0^{\mathrm{cycle}} = s_y y_0$. Its outer objective and Hessian are

$$\Phi_{\mathrm{cycle}}(x) = \frac{A}{2s_x^2}\|x\|^2 = \frac{\ell\sqrt{\rho}}{2C_{\mathrm{low}}}\|x\|^2, \qquad \nabla^2\Phi_{\mathrm{cycle}} = \frac{\ell\sqrt{\rho}}{C_{\mathrm{low}}} I_2 \succ 0. \tag{A.4}$$

Thus the full outer objective, not merely one component, is strongly convex and quadratic. This does not assert strong convexity of $f_{\mathrm{cycle}}(\cdot, y)$, which is affine.

After substituting $u = x/s_x$ and $v = y/s_y$ into the updates, both gradient coefficients equal

$$\frac{\eta_t A}{\rho s_x^2} = \frac{\eta_t A}{s_y^2} = \frac{\ell\eta_t}{C_{\mathrm{low}}\sqrt{\rho}} = h_t. \tag{A.5}$$

The outer–outer Hessian of f_{cycle} vanishes. Its mixed and inner–inner blocks have operator norms bounded, respectively, by

$$C\frac{A}{s_x s_y} = C\frac{\ell}{C_{\mathrm{low}}}, \qquad C\frac{A}{s_y^2} = C\frac{\ell}{C_{\mathrm{low}}\sqrt{\rho}}.$$

The norm of the full Hessian is at most the sum of these two block bounds, which is at most $2C\ell/C_{\mathrm{low}} \leq \ell$ since $\rho \geq 1$. Thus the full gradient is ℓ-Lipschitz. If μ is the inner-PL constant of $\tilde{g}_{\mathrm{cycle}}$, the scaled constant is $\mu A/s_y^2 = \mu\ell/(C_{\mathrm{low}}\sqrt{\rho})$, which is at least ℓ/κ when

$$\rho \leq c_{\mathrm{low}}\kappa^2, \qquad c_{\mathrm{low}} := \left(\frac{\mu}{C_{\mathrm{low}}}\right)^2. \tag{A.6}$$

All maxima are attained, as they are for $\tilde{g}_{\mathrm{cycle}}$.

Both initial gaps satisfy Assumption 3.3:

$$\begin{aligned} \Phi_{\mathrm{cycle}}(x_0) - \inf_x \Phi_{\mathrm{cycle}}(x) &= \frac{A}{2} = \frac{\zeta\Delta}{\sqrt{\rho}} \leq \Delta, \\ \Phi_{\mathrm{cycle}}(x_0) - f_{\mathrm{cycle}}(x_0, y_0^{\mathrm{cycle}}) &= A\left[\frac{1}{2} - \tilde{g}_{\mathrm{cycle}}(e_1, y_0)\right] \leq CA \leq 2C\zeta\Delta \leq c_{\mathrm{init}}\Delta. \end{aligned} \tag{A.7}$$

The last inequalities use $\rho \geq 1$ and the chosen ζ. Finally, every iterate satisfies

$$\|\nabla \Phi_{\mathrm{cycle}}(x_t)\|^2 = \frac{A^2}{s_x^2}\|u_t\|^2 = \frac{2\zeta\ell\Delta}{C_{\mathrm{low}}}\|u_t\|^2 \geq \frac{\zeta}{2C_{\mathrm{low}}}\ell\Delta. \tag{A.8}$$

The instance therefore belongs to $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ for every $\sigma \geq 0$, using exact gradients. Its dimension is $(d_x, d_y) = (2, 5)$, independent of the horizon and schedule. Taking the schedule infimum proves the risk bound.

For any possibly randomized index $\tau \in \{0, \ldots, T-1\}$ selected after observing the trajectory, the same instance satisfies $\|\nabla \Phi_{\mathrm{cycle}}(x_\tau)\|^2 \geq c\ell\Delta$ surely, since the bound holds for every generated iterate. This assertion does not concern a new point formed by arbitrary postprocessing. □

A.4 Upper bound for SGDA$_{\mathrm{Sim}}$

We prove the matching upper bound for SGDA$_{\mathrm{Sim}}$ under the same assumptions and average squared-gradient criterion as the lower bound. The deterministic one-sided-PL counterpart is already established by Chen et al. (2022, Appendix B.2); the argument below includes unbiased stochastic gradients with bounded variance and our convention of averaging over $x_0, \ldots, x_{T-1}$. The main issue is that $\nabla_x f(x_t, y_t)$ differs from $\nabla\Phi(x_t)$. Inner PL controls this discrepancy through the inner gradient by Lemma B.2; the potential below uses a faster inner step to absorb it.

Theorem A.1 (Upper bound for SGDA$_{\mathrm{Sim}}$)**.** *There are constants $c_{\mathrm{step}}, C > 0$, depending at most on c_{init}, such that the following holds for every $\kappa \geq 2$, $\ell, \Delta > 0$, $\sigma \geq 0$, and integer $T \geq 1$. On any instance in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$, run SGDA$_{\mathrm{Sim}}$ with*

$$\rho = 16\kappa^2, \qquad \alpha := c_{\mathrm{step}} \min\left\{\frac{1}{\kappa^2\ell}, \sqrt{\frac{\Delta}{\kappa^4\ell\sigma^2 T}}\right\}, \qquad \eta_t \equiv 16\kappa^2\alpha, \tag{A.9}$$

where the second term is $+\infty$ when $\sigma = 0$. Then

$$\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\|\nabla\Phi(x_t)\|^2 \leq C\left(\frac{\kappa^2\ell\Delta}{T} + \sqrt{\frac{\kappa^4\ell\Delta\sigma^2}{T}}\right). \tag{A.10}$$

Consequently, whenever $\varepsilon^2 \leq c\ell\Delta$ for a sufficiently small universal $c > 0$,

$$N_{\varepsilon,\mathrm{Sim}}(16\kappa^2), N^\star_{\varepsilon,\mathrm{Sim}} \leq C\left(\frac{\kappa^2\ell\Delta}{\varepsilon^2} + \frac{\kappa^4\ell\Delta\sigma^2}{\varepsilon^4}\right). \tag{A.11}$$

Proof of Theorem A.1. We combine outer descent with the change in the inner suboptimality. The ratio $16\kappa^2$ makes their combination absorb the squared primal-gradient discrepancy, leaving a telescoping bound and a variance term. We handle the final iterate x_{T-1} after summation.

The case $T = 1$ is immediate from (B.7): smoothness and lower boundedness of Φ imply $\|\nabla\Phi(x_0)\|^2 \leq 2\ell_\Phi[\Phi(x_0) - \Phi^*] \leq C\kappa^2\ell\Delta$. Assume henceforth that $T \geq 2$.

Let $\mathcal{F}_t$ be the history before the two oracle calls at iteration t and $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_t]$. The primal step is α from (A.9); write $\eta = 16\kappa^2\alpha$ for the constant dual step. Choose the universal constant $c_{\mathrm{step}} \leq 1/128$. Then

$$\ell_\Phi\alpha \leq \frac{1}{4}, \qquad \ell\alpha \leq \frac{1}{8}, \qquad \ell\eta \leq \frac{1}{8}. \tag{A.12}$$

By ℓ_Φ-smoothness of Φ and the unbiased primal oracle,

$$\begin{aligned}\mathbb{E}_t[\Phi(x_t) - \Phi(x_{t+1})] &\geq \alpha\langle\nabla\Phi(x_t), \nabla_x f(x_t, y_t)\rangle - \frac{\ell_\Phi\alpha^2}{2}(\|\nabla_x f(x_t, y_t)\|^2 + \sigma^2) \\ &= \frac{\alpha}{2}\|\nabla\Phi(x_t)\|^2 + \frac{\alpha}{2}(1 - \ell_\Phi\alpha)\|\nabla_x f(x_t, y_t)\|^2 \\ &\quad - \frac{\alpha}{2}\|\nabla_x f(x_t, y_t) - \nabla\Phi(x_t)\|^2 - \frac{\ell_\Phi\alpha^2}{2}\sigma^2.\end{aligned} \tag{A.13}$$

Assumption 3.1 implies that the full gradient is 2ℓ-Lipschitz in the Euclidean product norm. Since the primal and dual gradients in $\mathrm{SGDA}_{\mathrm{Sim}}$ are evaluated at (x_t, y_t), this gives

$$\mathbb{E}_t[f(x_{t+1}, y_{t+1}) - f(x_t, y_t)] \geq -\left(\alpha + \ell\alpha^2\right) \|\nabla_x f(x_t, y_t)\|^2 + \left(\eta - \ell\eta^2\right) \|\nabla_y f(x_t, y_t)\|^2 - \ell(\alpha^2 + \eta^2)\sigma^2. \tag{A.14}$$

Define

$$\mathcal{V}_t := \Phi(x_t) + \frac{1}{8}[\Phi(x_t) - f(x_t, y_t)]. \tag{A.15}$$

Since $\Phi(x_t) - f(x_t, y_t) \geq 0$, one has $\mathcal{V}_t \geq \Phi^*$. Combining $\frac{9}{8}$ times (A.13) with $\frac{1}{8}$ times (A.14), using $\|\nabla_x f(x_t, y_t) - \nabla\Phi(x_t)\|^2 \leq \kappa^2 \|\nabla_y f(x_t, y_t)\|^2$ from (B.6), and then applying (A.12), yields

$$\mathbb{E}_t[\mathcal{V}_t - \mathcal{V}_{t+1}] \geq \frac{9\alpha}{16} \|\nabla\Phi(x_t)\|^2 + \frac{9\alpha}{32} \|\nabla_x f(x_t, y_t)\|^2 + \frac{19\alpha\kappa^2}{16} \|\nabla_y f(x_t, y_t)\|^2 - 33\kappa^4\ell\alpha^2\sigma^2. \tag{A.16}$$

Indeed, the coefficient of $\|\nabla_x f(x_t, y_t)\|^2$ is at least $\frac{9\alpha}{16}(1 - \ell_\Phi\alpha) - \frac{\alpha}{8} - \frac{\ell\alpha^2}{8} \geq 9\alpha/32$, while the coefficient of $\|\nabla_y f(x_t, y_t)\|^2$ is at least $\frac{\eta}{8}(1 - \ell\eta) - 9\alpha\kappa^2/16 \geq 19\alpha\kappa^2/16$. The noise coefficient satisfies

$$\frac{9\ell_\Phi\alpha^2}{16} + \frac{\ell}{8}(\alpha^2 + \eta^2) \leq 33\kappa^4\ell\alpha^2 \qquad (\kappa \geq 2).$$

The two gap bounds in Assumption 3.3 give

$$\Phi(x_0) - f(x_0, y_0) \leq c_{\mathrm{init}}\Delta, \qquad \mathcal{V}_0 - \Phi^* \leq \left(1 + \frac{c_{\mathrm{init}}}{8}\right)\Delta. \tag{A.17}$$

Telescoping (A.16) for $0 \leq t \leq T-2$ therefore gives

$$\alpha \sum_{t=0}^{T-2} \mathbb{E}\left[\|\nabla\Phi(x_t)\|^2 + \|\nabla_x f(x_t, y_t)\|^2\right] \leq C\left[\Delta + \kappa^4\ell\alpha^2\sigma^2 T\right]. \tag{A.18}$$

It remains only to include the final iterate x_{T-1} in the performance criterion. By (B.7),

$$\mathbb{E}\|\nabla\Phi(x_{T-1})\|^2 \leq 2\mathbb{E}\|\nabla\Phi(x_{T-2})\|^2 + 2\ell_\Phi^2\alpha^2\mathbb{E}\left[\|\nabla_x f(x_{T-2}, y_{T-2})\|^2 + \sigma^2\right]. \tag{A.19}$$

Consequently,

$$\sum_{t=0}^{T-1} \mathbb{E}\|\nabla\Phi(x_t)\|^2 \leq 3\sum_{t=0}^{T-2} \mathbb{E}\|\nabla\Phi(x_t)\|^2 + 2\ell_\Phi^2\alpha^2 \sum_{t=0}^{T-2} \mathbb{E}\|\nabla_x f(x_t, y_t)\|^2 + 2\ell_\Phi^2\alpha^2\sigma^2.$$

Using $\ell_\Phi\alpha \leq 1/4$ and (A.18), we obtain

$$\sum_{t=0}^{T-1} \mathbb{E}\|\nabla\Phi(x_t)\|^2 \leq C\left[\frac{\Delta}{\alpha} + \kappa^4\ell\alpha\sigma^2 T\right]. \tag{A.20}$$

Dividing by T gives

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\|\nabla\Phi(x_t)\|^2 \leq C\left[\frac{\Delta}{\alpha T} + \kappa^4\ell\alpha\sigma^2\right]. \tag{A.21}$$

Substitute (A.9). If the step-size bound $1/(\kappa^2\ell)$ is active, then $\sigma^2 T \leq \ell\Delta$ and both terms in (A.21) are bounded by $C\kappa^2\ell\Delta/T$. Otherwise the two terms are both bounded by $C\sqrt{\kappa^4\ell\Delta\sigma^2/T}$. This proves (A.10).

Finally, choosing

$$T \geq C\left(\frac{\kappa^2\ell\Delta}{\varepsilon^2} + \frac{\kappa^4\ell\Delta\sigma^2}{\varepsilon^4}\right)$$

makes the right-hand side of (A.10) at most ε^2. The constant schedule is nonincreasing, so the minimax definitions in (3.7) give (A.11). □

B PRELIMINARY LEMMAS

We first record the completion-of-squares construction used to preserve the outer function and inner PL. We then prove the consequences of inner PL used in the matching upper bound.

B.1 NC-PL GAMES WITH A PREDETERMINED OUTER FUNCTION

We use an onto inner map to determine the outer maximum exactly. A lower bound on its transposed derivative then converts the residual norm into an inner-gradient bound. This lemma proves inner PL, not global smoothness; the latter is checked for each constructed map.

Lemma B.1 (Outer function and inner PL for a completion-of-squares construction)**.** *Let $b : \mathbb{R}^{d_x} \to \mathbb{R}^m$ and $T : \mathbb{R}^{d_y} \to \mathbb{R}^m$ be continuously differentiable, with T onto (surjective). Define*

$$g(x, y) := \langle b(x), T(y) \rangle - \frac{1}{2} \|T(y)\|^2 . \tag{B.1}$$

Then

$$\max_y g(x, y) = \frac{1}{2} \|b(x)\|^2 , \qquad \max_y g(x, y) - g(x, y) = \frac{1}{2} \|b(x) - T(y)\|^2 . \tag{B.2}$$

If, for some $\nu > 0$,

$$\left\| DT(y)^\mathsf{T} z \right\| \geq \nu \|z\| \qquad \text{for all } y, z, \tag{B.3}$$

then

$$\frac{1}{2} \|\nabla_y g(x, y)\|^2 \geq \nu^2 \left[\max_v g(x, v) - g(x, y) \right] . \tag{B.4}$$

Proof. Since T is onto, maximizing (B.1) over y is equivalent to maximizing $\langle b(x), z \rangle - \|z\|^2 /2$ over $z \in \mathbb{R}^m$. Completing the square gives (B.2). Moreover,

$$\nabla_y g(x, y) = DT(y)^\mathsf{T} [b(x) - T(y)],$$

so (B.3) and (B.2) imply (B.4). □

B.2 NC-PL GEOMETRY

The following are standard consequences of smoothness and inner PL; see Nouiehed et al. (2019, Lemma A.5), Chen et al. (2022, Lemmas A.7–A.8), and Ha (2026, Lemmas A.2–A.3). We prove them under Assumptions 3.1 and 3.2, without concavity or uniqueness of the maximizer.

Lemma B.2 (Inner error bound, gradient bias, and outer smoothness)**.** *Under Assumptions 3.1 and 3.2, Φ is differentiable with $\nabla \Phi(x) = \nabla_x f(x, y^\star)$ for every $y^\star \in Y^\star(x)$. For every (x, y),*

$$\operatorname{dist}(y, Y^\star(x)) \leq \frac{\kappa}{\ell} \|\nabla_y f(x, y)\| , \tag{B.5}$$

$$\|\nabla_x f(x, y) - \nabla \Phi(x)\| \leq \kappa \|\nabla_y f(x, y)\| . \tag{B.6}$$

Moreover, Φ has ℓ_Φ-Lipschitz gradient with

$$\ell_\Phi \leq \ell(1 + \kappa). \tag{B.7}$$

Proof. We first bound the inner gap after changing x while keeping a maximizer fixed; the resulting quadratic remainder proves differentiability of Φ. An inner gradient-flow argument then bounds the distance to a maximizer, giving the gradient-bias and outer-smoothness estimates.

Fix x, $y^\star \in Y^\star(x)$, and $h \in \mathbb{R}^{d_x}$. Since the inner domain is unconstrained, $\nabla_y f(x, y^\star) = 0$. Inner PL at $(x + h, y^\star)$ and Assumption 3.1 give

$$0 \leq \Phi(x + h) - f(x + h, y^\star) \leq \frac{\|\nabla_y f(x + h, y^\star)\|^2}{2\mu} \leq \frac{\ell\kappa}{2} \|h\|^2 .$$

Combining this with the Taylor bound for $f(\cdot, y^\star)$ yields

$$|\Phi(x + h) - \Phi(x) - \langle \nabla_x f(x, y^\star), h \rangle| \leq \frac{\ell(1 + \kappa)}{2} \|h\|^2 .$$

Thus Φ is differentiable and $\nabla\Phi(x) = \nabla_x f(x, y^\star)$ for every maximizer, including when the maximizer is not unique. If $d_y = 0$, this identity and all the claimed bounds follow directly from Assumption 3.1; henceforth assume $d_y \geq 1$.

Fix x and define the inner objective gap $g_x(y) := \Phi(x) - f(x, y) \geq 0$. It is ℓ-smooth, attains its minimum zero, and satisfies $\|\nabla g_x(y)\|^2 \geq 2(\ell/\kappa) g_x(y)$. Since ∇g_x is globally Lipschitz, the gradient flow $\dot{y}(s) = -\nabla g_x(y(s))$, initialized at $y(0) = y$, exists for all $s \geq 0$. If it reaches $g_x = 0$ in finite time, the desired distance bound follows immediately by integrating up to that time. Otherwise $g_x(y(s)) > 0$ for every finite s, and

$$-\frac{d}{ds}\sqrt{g_x(y(s))} = \frac{\|\nabla g_x(y(s))\|^2}{2\sqrt{g_x(y(s))}} \geq \sqrt{\frac{\ell}{2\kappa}}\,\|\nabla g_x(y(s))\|\,.$$

Hence the total distance traveled by the gradient-flow solution is finite:

$$\int_0^\infty \|\dot{y}(s)\|\, ds \leq \sqrt{\frac{2\kappa g_x(y)}{\ell}}.$$

The finite path length makes $y(s)$ a Cauchy curve in $\mathbb{R}^{d_y}$, so it converges to a limit $y^\star$. The limit must satisfy $g_x(y^\star) = 0$: otherwise PL would keep $\|\nabla g_x\|$ bounded away from zero and the integral of its speed would be infinite. Thus $y^\star \in Y^\star(x)$, and therefore

$$\operatorname{dist}(y, Y^\star(x)) \leq \sqrt{\frac{2\kappa g_x(y)}{\ell}} \leq \frac{\kappa}{\ell}\,\|\nabla g_x(y)\|\,,$$

where the second inequality is the PL inequality. This proves (B.5).

Choose the maximizer supplied by (B.5). The gradient identity just proved and Assumption 3.1 give

$$\|\nabla_x f(x, y) - \nabla\Phi(x)\| \leq \ell\,\|y - y^\star\| \leq \kappa\,\|\nabla_y f(x, y)\|\,,$$

which is (B.6).

Finally, fix x, x' and $y^\star \in Y^\star(x)$. Since $\nabla_y f(x, y^\star) = 0$, Assumption 3.1 gives $\|\nabla_y f(x', y^\star)\| \leq \ell\,\|x' - x\|$. Applying (B.5) to the inner problem at x' yields a maximizer $\bar{y} \in Y^\star(x')$ such that $\|\bar{y} - y^\star\| \leq \kappa\,\|x' - x\|$. Thus

$$\begin{aligned}\|\nabla\Phi(x') - \nabla\Phi(x)\| &= \|\nabla_x f(x', \bar{y}) - \nabla_x f(x, y^\star)\| \\ &\leq \ell\big(\|x' - x\| + \|\bar{y} - y^\star\|\big) \leq \ell(1 + \kappa)\,\|x' - x\|\,.\end{aligned}$$

This proves (B.7). □

C Proofs for the high-ratio lower bound

We prove the construction lemmas used in Section 5. The quadratic estimate accounts for initially large dual steps. The corridor objective controls the remaining primal movement. For the stochastic term, we first bound the quadratic perturbation response, then realize it by a globally smooth inner-PL game. Finally, we combine the corridor and perturbation components into one admissible instance and prove its lower bound.

C.1 Proof of the lower bound for f_{unstable}

For the quadratic game (5.2), the squared outer gradient is $|\Phi'_{\text{unstable}}(x_t)|^2 = (\ell + \mu)^2 x_t^2/64$. We first choose an initial pair (x_0, y_0) of norm one for which $\sum_{t=0}^m x_t^2$ is large. The following matrices express this sum as a quadratic form in the initial pair.

For $s_t = (\ell + \mu)\eta_t/2$, define

$$M_t^{\text{Alt}} := \begin{pmatrix} 1 & -s_t/(2\rho) \\ s_t/2 & 1 - s_t - s_t^2/(4\rho) \end{pmatrix}, \qquad M_t^{\text{Sim}} := \begin{pmatrix} 1 & -s_t/(2\rho) \\ s_t/2 & 1 - s_t \end{pmatrix}. \tag{C.1}$$

For SGDA_{Sim} and SGDA_{Alt}, using their respective matrices, let $P_0 = I$, $P_t = M_{t-1} \cdots M_0$, and

$$W_m := \sum_{t=0}^{m} P_t^\mathsf{T} e_1 e_1^\mathsf{T} P_t. \tag{C.2}$$

Thus an initial vector $v = (x_0, y_0)^\mathsf{T}$ generates $x_t = e_1^\mathsf{T} P_t v$ and $\sum_{t=0}^{m} x_t^2 = v^\mathsf{T} W_m v$. A lower bound on $\lambda_{\max}(W_m)$ supplies a suitable initial direction. Here $M_t = M_t^{\mathrm{Sim}}$ for $\mathrm{SGDA}_{\mathrm{Sim}}$ and $M_t = M_t^{\mathrm{Alt}}$ for $\mathrm{SGDA}_{\mathrm{Alt}}$.

C.1.1 $\mathrm{SGDA}_{\mathrm{ALT}}$

The required bound follows by combining persistence over shorter runs with the growth of determinants over longer runs.

Lemma C.1 (Finite-horizon estimate for $\mathrm{SGDA}_{\mathrm{Alt}}$)**.** *There is a universal $c > 0$ such that, for every $\rho \geq 1$ and every non-increasing sequence $s_0 \geq \cdots \geq s_{m-1} \geq 4$,*

$$\lambda_{\max}(W_m) \geq c(m+1). \tag{C.3}$$

Proof. We first prove the intermediate bound

$$\lambda_{\max}(W_m) \geq c \min\{m+1, \rho\} \quad \text{for non-increasing } s_t \geq 2. \tag{C.4}$$

We compare the rows mapping the initial state to consecutive outer coordinates. If their determinants are large, their squared norms accumulate. Otherwise, a backward construction keeps the outer coordinate away from zero. The stronger threshold $s_t \geq 4$ then extends the bound to every run length.

Write $u_t = e_1^\mathsf{T} P_t$. Since $\det M_j^{\mathrm{Alt}} = 1 - s_j$, for $0 \leq t < m$,

$$D_t := \left| \det \begin{pmatrix} u_t \\ u_{t+1} \end{pmatrix} \right| = \frac{s_t}{2\rho} \prod_{j=0}^{t-1} (s_j - 1). \tag{C.5}$$

For $0 \leq t < m-1$,

$$\frac{D_{t+1}}{D_t} = s_{t+1} \left(1 - \frac{1}{s_t}\right) \geq 1,$$

so (D_t) is nondecreasing.

If $\min\{m+1, \rho\} \leq 12$, then $W_m \succeq e_1 e_1^\mathsf{T}$ proves the claim after decreasing c. Assume $\min\{m+1, \rho\} > 12$, set $N = \min\{m, \lfloor \rho \rfloor\}$ and $k = \lfloor N/2 \rfloor$. Then $N \geq 6$, $k \geq N/3$, and $N \geq c \min\{m+1, \rho\}$. Fix a sufficiently small universal $\delta_0 > 0$.

If $D_k \geq \delta_0$, then $D_t \geq \delta_0$ for $k \leq t < N$, and

$$\|u_t\|^2 + \|u_{t+1}\|^2 \geq 2 \left| \det \begin{pmatrix} u_t \\ u_{t+1} \end{pmatrix} \right| \geq 2\delta_0.$$

Summing and accounting for overlap gives $\operatorname{tr} W_m \geq cN$, hence $\lambda_{\max}(W_m) \geq \frac{1}{2} \operatorname{tr} W_m \geq c \min\{m+1, \rho\}$.

Suppose $D_k < \delta_0$. Since $s_k \geq 2$, (C.5) gives

$$A := \prod_{j=0}^{k-1} (s_j - 1) < \delta_0 \rho. \tag{C.6}$$

All factors are at least one, so $A \geq 1 + \sum_{j<k} (s_j - 2)$ and $\sum_{j<k} s_j \leq \delta_0 \rho + 2k$. Also $A \geq (s_k - 1)^k$, so $s_j \leq s_k \leq 1 + A^{1/k}$ for $j \geq k$. Since $\rho \geq N$ and $k \geq N/3$,

$$N\rho^{1/k} \leq C\rho, \qquad N\rho^{1/k-1} \leq N N^{1/k-1} = N^{1/k} \leq e^3.$$

Consequently,

$$\sum_{j=0}^{N-1} s_j \leq C\rho. \tag{C.7}$$

The bound on the step sum allows a trajectory with bounded total relative decrease. To construct it, eliminate y_t and work backward from the desired ratio between consecutive outer coordinates. The two-dimensional iteration gives

$$x_{t+2} = (1 - B_t - C_t) x_{t+1} + B_t x_t, \quad B_t := \frac{s_{t+1}}{s_t}(s_t - 1) \geq 1, \quad C_t := \frac{s_t s_{t+1}}{4\rho}, \tag{C.8}$$

for $0 \le t \le N-2$. Set $\varrho_{N-1} = 1$ and define backward

$$\varrho_t := \frac{B_t}{\varrho_{t+1} + B_t + C_t - 1}, \qquad t = N-2, \dots, 0. \tag{C.9}$$

With $d_t = 1 - \varrho_t$,

$$d_t = \frac{C_t - d_{t+1}}{B_t + C_t - d_{t+1}}. \tag{C.10}$$

The sequence C_t is nonincreasing. Backward induction from $d_{N-1} = 0$ gives

$$0 \le d_t \le \frac{C_t}{B_t} = \frac{s_t^2}{4\rho(s_t - 1)} \le \frac{s_t + 2}{4\rho}. \tag{C.11}$$

At $t = N-2$, use $d_{N-1} = 0$ directly. At earlier indices, if $0 \le d_{t+1} \le C_{t+1}/B_{t+1}$, then $d_{t+1} \le C_{t+1} \le C_t$, while the denominator in (C.10) is at least B_t. Since $s_0 - 1 \le A < \delta_0 \rho$, choosing δ_0 sufficiently small and using $\rho > 12$ gives $d_t \le 1/4$. Equations (C.7)–(C.11) imply

$$\sum_{t=0}^{N-2} d_t \le C. \tag{C.12}$$

Set $x_0 = 1$ and $x_{t+1} = \varrho_t x_t$. This sequence satisfies (C.8), and

$$x_t = \prod_{j=0}^{t-1} (1 - d_j) \ge \exp\left(-2 \sum_{j=0}^{t-1} d_j \right) \ge c, \qquad 0 \le t \le N. \tag{C.13}$$

The corresponding initial dual coordinate is

$$y_0 = \frac{2\rho(1 - \varrho_0)}{s_0} = \frac{2\rho d_0}{s_0} \le \frac{s_0}{2(s_0 - 1)} \le 1. \tag{C.14}$$

Thus $v = (1, y_0)^\mathsf{T}$ satisfies $\|v\| \le \sqrt{2}$, while $x_t = u_t v$, and

$$\lambda_{\max}(W_m) \ge \frac{v^\mathsf{T} W_m v}{\|v\|^2} = \frac{\sum_{t=0}^{m} x_t^2}{\|v\|^2} \ge cN \ge c \min\{m+1, \rho\}.$$

This proves (C.4).

For the required bound (C.3), if $m \le 2\rho$, then $\min\{m+1, \rho\} \ge (m+1)/3$, so (C.4) suffices. If $m > 2\rho$, then $D_0 = s_0/(2\rho) \ge 2/\rho$, and $D_{t+1}/D_t \ge 3$. Hence $D_t \ge 1$ after at most $\lceil \log_3(\rho/2) \rceil_+ \le \rho$ updates. At least $m/2$ subsequent adjacent pairs remain, and the determinant estimate gives $\operatorname{tr} W_m \ge cm$. □

C.1.2 SGDA$_{\text{SIM}}$

Here the update matrices have common eigenvectors. We use their two scalar products to show that suppressing one contribution cannot suppress their sum, which controls the norm of the outer-coordinate row.

Lemma C.2 (Finite-horizon estimate for SGDA$_{\text{Sim}}$)**.** *There is a universal $c > 0$ such that, for every $\rho \ge 1$ and every sequence $s_0, \dots, s_{m-1} \ge 4$,*

$$\lambda_{\max}(W_m) \ge c(m+1). \tag{C.15}$$

Proof. The two common eigenvectors show why large steps cannot suppress all outer responses. Put $\nu = \sqrt{1 - \rho^{-1}}$ and $\mu_\pm = (1 \pm \nu)/2$. The matrices M_t^{Sim} have eigenvectors

$$v_- = \begin{pmatrix} 1 \\ (1+\nu)^{-1} \end{pmatrix}, \qquad v_+ = \begin{pmatrix} 1 \\ \rho(1+\nu) \end{pmatrix}, \qquad M_t^{\text{Sim}} v_\pm = (1 - \mu_\pm s_t) v_\pm. \tag{C.16}$$

For $1 \le \rho \le 2$, $|1 - \mu_+ s_t| \ge 1$ and $\|v_+\| \le \sqrt{17}$. Thus $u_t = e_1^\mathsf{T} P_t$ satisfies $\|u_t\| \ge 1/\sqrt{17}$ for every t, including $\rho = 1$ where the eigenvalues coincide.

For $\rho \geq 2$, define the two response magnitudes and their sum by

$$X_t = \prod_{j<t} |1 - \mu_- s_j|, \qquad Y_t = \frac{1}{\rho} \prod_{j<t} (\mu_+ s_j - 1), \qquad H_t = X_t + Y_t. \tag{C.17}$$

Since $|u_t v_-| = X_t$, $|u_t v_+| = \rho Y_t$, $\|v_-\| \leq \sqrt{2}$, and $\|v_+\| \leq \sqrt{5}\rho$, we have $\|u_t\| \geq H_t/(2\sqrt{5})$. For $s \geq 4$, write $a(s) = |1 - \mu_- s|$ and $b(s) = \mu_+ s - 1$. The inequality $b(s) \geq a(s)$ gives $Y_t/X_t \geq 1/\rho$ whenever $X_t > 0$.

For every $\varrho \geq 1/\rho$, the function $a(s) + \varrho b(s)$ is increasing on $[4, \infty)$. Indeed, its slope before $s = 1/\mu_-$ is $\varrho\mu_+ - \mu_- \geq \mu_-(4\mu_+^2 - 1) \geq 0$, and the slope afterward is positive. Since $\nu \geq 1/\sqrt{2}$, evaluation at $s = 4$ gives

$$a(s) + \varrho b(s) - (1 + \varrho) \geq -4\mu_- + \frac{2\nu}{\rho} = 2(1 - \nu)(\nu + \nu^2 - 1) \geq 0. \tag{C.18}$$

Hence $H_{t+1} \geq H_t$ when $X_t > 0$. If $X_t = 0$, it remains zero and $b(s_t) \geq 1$ gives the same conclusion. Thus $H_t \geq H_0 \geq 1$, so the row norm is uniformly bounded below. In both ratio ranges,

$$\lambda_{\max}(W_m) \geq \tfrac{1}{2} \operatorname{tr} W_m = \tfrac{1}{2} \sum_{t=0}^{m} \|u_t\|^2 \geq c(m+1).$$

□

C.1.3 Conversion to the outer-gradient lower bound

Lemma 5.2 (Quadratic lower bound for large dual steps)**.** *There is a universal $c > 0$ such that, for $\ell, \Delta > 0$, $\kappa \geq 2$, $\sigma \geq 0$, $\rho \geq 1$, integers $T \geq 1$ and $0 \leq m < T$, and predetermined nonnegative non-increasing step sizes with $\eta_t \geq 8/\ell$ for $t < m$, each of SGDA$_{\text{Sim}}$ and SGDA$_{\text{Alt}}$ admits an initialization (x_0, y_0) making f_{unstable} with $\mu = \ell/\kappa$ an exact-gradient instance in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with*

$$\frac{1}{T} \sum_{t=0}^{T-1} |\Phi'_{\text{unstable}}(x_t)|^2 \geq c \frac{m+1}{T} \ell \Delta. \tag{5.3}$$

Proof. For the chosen algorithm, SGDA$_{\text{Sim}}$ or SGDA$_{\text{Alt}}$, let $w = (w_x, w_y)^\top$ be a unit eigenvector of W_m for its largest eigenvalue. Among all initial pairs of norm one, w maximizes $\sum_{t=0}^{m} x_t^2 = w^\top W_m w = \lambda_{\max}(W_m)$. We set $x_0 = \sqrt{\Delta/(\ell+\mu)}\, w_x$ and $y_0 = \sqrt{\Delta/(\ell+\mu)}\, w_y$. By linearity of the updates, this multiplies every iterate by $\sqrt{\Delta/(\ell+\mu)}$ and the accumulated squared outer gradient by $\Delta/(\ell+\mu)$. We now check that this scale satisfies the initial-gap budgets.

Since $\mu = \ell/\kappa \leq \ell/2$, the Hessian is $(\ell+\mu)\left(\begin{smallmatrix} 0 & 1/4 \\ 1/4 & -1/2 \end{smallmatrix}\right)$ and the inner maximum is attained at $y = x/2$. Hence

$$\begin{aligned} \left\|\nabla^2 f_{\text{unstable}}\right\|_{\text{op}} &= \frac{1+\sqrt{2}}{4}(\ell+\mu) \leq \frac{3(1+\sqrt{2})}{8}\ell < \ell, \\ \Phi_{\text{unstable}}(x) - f_{\text{unstable}}(x,y) &= \frac{\ell+\mu}{4}(y - x/2)^2, \\ \frac{1}{2}|\nabla_y f_{\text{unstable}}(x,y)|^2 &= \frac{\ell+\mu}{2}\big(\Phi_{\text{unstable}}(x) - f_{\text{unstable}}(x,y)\big). \end{aligned}$$

Since $(\ell+\mu)/2 \geq \mu$, the game is NC-μ-PL. Our initialization satisfies $x_0^2 + y_0^2 = \Delta/(\ell+\mu)$, so Cauchy–Schwarz gives $(y_0 - x_0/2)^2 \leq 5\Delta/[4(\ell+\mu)]$ and

$$\Phi_{\text{unstable}}(x_0) \leq \frac{\Delta}{16}, \qquad \Phi_{\text{unstable}}(x_0) - f_{\text{unstable}}(x_0, y_0) \leq \frac{5\Delta}{16}.$$

These inequalities verify both initial-gap budgets, since $c_{\text{init}} \geq 1$.

Finally, the definition of W_m and Lemmas C.1 and C.2 yield

$$\frac{1}{T} \sum_{t<T} |\Phi'_{\text{unstable}}(x_t)|^2 \geq \frac{(\ell+\mu)^2}{64T} \sum_{t=0}^{m} x_t^2 = \frac{(\ell+\mu)\Delta}{64T} \lambda_{\max}(W_m) \geq c \frac{m+1}{T} \ell \Delta. \tag{C.19}$$

Here the cutoff gives $s_t = (\ell+\mu)\eta_t/2 \geq \ell\eta_t/2 \geq 4$, and $\ell + \mu \geq \ell$ gives the final inequality. Exact gradients are admissible for every $\sigma \geq 0$, so this proves the claim. □

C.2 Construction of g_{corridor}

Lemma 5.3 (Corridor lower bound)**.** *For $\ell, \Delta > 0$, $\rho \geq 1$, integers $T \geq 1$, $0 \leq m < T$, and nonnegative steps with $S = \sum_{t=m}^{T-2} \eta_t$, there is an ℓ-smooth corridor $g_{\mathrm{corridor}} : \mathbb{R} \to \mathbb{R}$ with gap Δ at zero and slope $G > 0$, where $G^2 \asymp \min\{\ell\Delta, \rho\Delta/S\}$. Gradient descent from zero with steps η_t/ρ satisfies $g'_{\mathrm{corridor}}(u_t) < 0$ for all $t < T$ and $g'_{\mathrm{corridor}}(u_t) = -G$ for $m \leq t < T$. Consequently,*

$$\frac{1}{T}\sum_{t<T} |g'_{\mathrm{corridor}}(u_t)|^2 \gtrsim \frac{T-m}{T} \min\left\{\ell\Delta, \frac{\rho\Delta}{S}\right\}. \tag{5.5}$$

Proof. We prescribe the derivative before choosing its magnitude. An initial linear segment makes the first m updates reach a constant-derivative interval. Its length is the distance traveled by the remaining steps; a final linear segment joins it to a minimum. Integrating the derivative then calibrates the exact initial gap.

Fix $m \in \{0, \dots, T-1\}$ and let $S = \sum_{t=m}^{T-2} \eta_t$. When $S = 0$, the minimum in the statement is understood as $\ell\Delta$. Define

$$P_0 := 1, \qquad P_t := \prod_{i=0}^{t-1} \left(1 + \frac{\ell\eta_i}{\rho}\right) \quad (1 \leq t \leq m), \tag{C.20}$$

$$G^2 := \frac{\Delta}{S/\rho + \ell^{-1}(1 - 1/(2P_m^2))}, \qquad G_0 := \frac{G}{P_m}, \qquad b := \frac{G - G_0}{\ell} + \frac{GS}{\rho}. \tag{C.21}$$

For $G > 0$, write $[z]_{[0,G]} := \min\{G, \max\{0, z\}\}$. Add a constant to g_{corridor} so that $\inf_u g_{\mathrm{corridor}}(u) = 0$, and define

$$g'_{\mathrm{corridor}}(u) := -[G_0 + \ell u]_{[0,G]} + [\ell(u - b)]_{[0,G]}. \tag{C.22}$$

The two intervals on which the clipping terms vary are disjoint, so g'_{corridor} is continuous and ℓ-Lipschitz.

Starting from $u_0 = 0$, induction gives

$$u_t = \frac{P_t - 1}{\ell} G_0, \qquad g'_{\mathrm{corridor}}(u_t) = -P_t G_0, \qquad 0 \leq t \leq m, \tag{C.23}$$

$$u_t = \frac{G - G_0}{\ell} + G\sum_{j=m}^{t-1} \frac{\eta_j}{\rho}, \qquad g'_{\mathrm{corridor}}(u_t) = -G, \qquad m \leq t \leq T-1. \tag{C.24}$$

Direct integration gives

$$\begin{aligned} g_{\mathrm{corridor}}(0) &= \frac{G^2 - G_0^2}{2\ell} + \frac{G^2 S}{\rho} + \frac{G^2}{2\ell} \\ &= G^2\left(\frac{S}{\rho} + \frac{1}{\ell} - \frac{1}{2\ell P_m^2}\right) = \Delta, \end{aligned} \tag{C.25}$$

and

$$G^2 \geq \frac{\Delta}{S/\rho + 1/\ell} \geq \frac{1}{2}\min\left\{\ell\Delta, \frac{\rho\Delta}{S}\right\}. \tag{C.26}$$

Since $P_m \geq 1$, the denominator in (C.21) is at least $S/\rho + 1/(2\ell)$, so also $G^2 \leq 2\min\{\ell\Delta, \rho\Delta/S\}$. This proves the two-sided comparison in the statement.

The game is the single-player objective $f(u) = g_{\mathrm{corridor}}(u)$, equivalently an instance with no inner variable. Hence the inner-PL condition is vacuous. Equations (C.24) and (C.26), summed over the last $T - m$ iterates, prove (5.5). Since $G_0, G > 0$ and $P_t > 0$, the derivative is negative at every iterate, so each positive update strictly increases the scalar iterate. The initial gap equals Δ by (C.25). Exact gradients are admissible for every σ. □

C.3 Construction of g_σ

The target is the average squared outer-gradient bound for the noise components in Lemma C.4. We first bound the primal displacement caused by one inner-gradient error, construct a smooth inner-PL game with these dynamics, and count the iterates with a large outer gradient. Summing the perturbations with enough subsequent cumulative step gives the bound. The separable construction afterward supplies one admissible oracle for all perturbations.

C.3.1 Primal displacement under a fixed shift

Lemma 5.4 (Persistent displacement under a fixed shift)**.** *There are universal constants $c, C, C_0 > 0$ such that, for $\ell, \delta > 0$, $\kappa \geq 1$, $\rho \geq C_0\kappa^2$, and any finite step-size sequence with $0 \leq \ell\eta_t \leq c\kappa$, both SGDA$_{Sim}$ and SGDA$_{Alt}$ with exact gradients on the linear quadratic game g_δ given by (5.6), initialized at $(0,0,0)$, satisfy $|z_t| \leq C\kappa\delta/\rho$ and $|a_t| + |b_t| \leq C\delta$ whenever $\sum_{i<t}\eta_i \leq c\rho/(\ell\kappa)$, and*

$$z_t \geq c\frac{\kappa\delta}{\rho} \quad \text{if} \quad C\frac{\kappa}{\ell} \leq \sum_{i<t}\eta_i \leq c\frac{\rho}{\ell\kappa}. \tag{5.7}$$

Proof. We first normalize the quadratic dynamics. Inner contraction bounds the response in terms of the largest outer displacement; an exact summed identity then bounds that displacement from above and below. Set

$$X_n = \frac{\kappa z_n}{\delta}, \qquad Y_n = \frac{a_n}{\delta}, \qquad Z_n = 1 + \frac{b_n}{\delta}, \qquad h_n = \frac{\ell\eta_n}{\kappa}, \qquad \alpha = \frac{\kappa^2}{\rho}.$$

The quadratic updates become

$$\begin{aligned} X_{n+1} &= X_n - \alpha h_n Y_n, \\ Y_{n+1} &= Y_n + h_n(X'_n - 2Y_n - Z_n), \qquad (X_0, Y_0, Z_0) = (0,0,1), \\ Z_{n+1} &= Z_n - h_n(Y_n + Z_n), \end{aligned} \tag{C.27}$$

where $X'_n = X_n$ for SGDA$_{\text{Sim}}$ and $X'_n = X_{n+1}$ for SGDA$_{\text{Alt}}$. We choose universal positive constants $\alpha_\star, \bar{h}, \tau_-, \tau_+$ below. For $\alpha \leq \alpha_\star$ and $h_n \leq \bar{h}$, we will bound $|X_n|$ by a universal multiple of α up to cumulative step τ_+/α, and $X_n \geq \alpha/4$ after cumulative step τ_-.

For the normalized inner dynamics, let $B = \left(\begin{smallmatrix} -2 & -1 \\ -1 & -1 \end{smallmatrix}\right)$ and let $q_n = (Y_n, Z_n)^\mathsf{T}$ be the inner-coordinate vector. The eigenvalues of $-B$ are $(3 \pm \sqrt{5})/2$. Write $\lambda_0 = (3-\sqrt{5})/2$ and choose $\bar{h} \leq 2/(3+\sqrt{5})$. Then $I + hB$ is positive semidefinite and its norm is $1 - \lambda_0 h$ for $0 \leq h \leq \bar{h}$.

Fix an endpoint N, and write $M_N = \max_{0\leq i\leq N}|X_i|$ and $\tau_n = \sum_{i<n} h_i$ only within this proof. Since $|X'_n| \leq M_N$ for $n < N$, the inner update with the additional term $h_n X'_n$ gives

$$\|q_{n+1}\| \leq (1 - \lambda_0 h_n)\,\|q_n\| + h_n M_N, \tag{C.28}$$

$$\|q_n\| \leq \prod_{i<n}(1-\lambda_0 h_i) + \frac{M_N}{\lambda_0}\left[1 - \prod_{i<n}(1-\lambda_0 h_i)\right] \leq e^{-\lambda_0\tau_n} + \frac{M_N}{\lambda_0}. \tag{C.29}$$

The second line is a scalar induction from $\|q_0\| = 1$, followed by $1 - a \leq e^{-a}$.

Subtracting the two inner updates yields $Y_{n+1} - Z_{n+1} - (Y_n - Z_n) = h_n(X'_n - Y_n)$. Sum this equality, use $Y_0 - Z_0 = -1$, and then sum the outer update to obtain

$$X_n = \alpha\left[1 + Y_n - Z_n - \sum_{i<n} h_i X'_i\right]. \tag{C.30}$$

For every $n \leq N$, $|Y_n - Z_n| \leq \sqrt{2}\,\|q_n\|$ and $\sum_{i<n} h_i|X'_i| \leq \tau_N M_N$. Hence

$$M_N \leq \alpha(1+\sqrt{2}) + \alpha\left(\frac{\sqrt{2}}{\lambda_0} + \tau_N\right) M_N. \tag{C.31}$$

Put $C_X = 2(1+\sqrt{2})$. If $\alpha_\star\sqrt{2}/\lambda_0 \leq 1/4$ and $\tau_+ \leq 1/4$, then $\tau_N \leq \tau_+/\alpha$ makes the coefficient of M_N on the right at most $1/2$. Thus $M_N \leq C_X\alpha$, and (C.29) gives the uniform inner-coordinate bound.

Using the sign in (C.30) now gives

$$\frac{X_N}{\alpha} \geq 1 - \sqrt{2}e^{-\lambda_0 \tau_N} - \frac{\sqrt{2}}{\lambda_0} M_N - \tau_N M_N \geq 1 - \sqrt{2}e^{-\lambda_0 \tau_N} - \frac{\sqrt{2}C_X}{\lambda_0}\alpha - C_X \alpha \tau_N. \quad \text{(C.32)}$$

Choose $\tau_- = \lambda_0^{-1} \log(4\sqrt{2})$, take $\tau_+ \leq \min\{1/4, 1/(4C_X)\}$, and reduce $\alpha_\star$ so that

$$\alpha_\star \leq \min\left\{ \frac{\lambda_0}{4\sqrt{2}C_X},\ 1,\ \frac{\tau_+}{2(\tau_- + 2\bar{h})} \right\}.$$

On the lower-bound interval, each of the three errors in (C.32) is at most $1/4$, so $X_N \geq \alpha/4$. These choices also imply the earlier absorption condition and ensure that the interval has positive length. Since N was arbitrary, both conclusions hold at every stated iterate. Zero steps cause no change, and the proof covers both choices of X_n'.

To recover the stated constants, choose $c > 0$ no larger than $\bar{h}$, τ_+, and $1/4$, and take $C \geq \tau_-$ large enough to bound $|X_n|/\alpha$ and $|Y_n| + |Z_n - 1|$. Reduce c so that $cC \leq 1/4$, and then reduce $\alpha_\star$ so that $c/\alpha_\star \geq 2(C + 2c)$. Set $C_0 = \alpha_\star^{-1}$. These choices are fixed before the cutoff and subsequent normalization.

If $\rho \geq C_0 \kappa^2$ and $\ell \eta_n \leq c\kappa$, then $\alpha \leq \alpha_\star$ and $h_n \leq \bar{h}$. Substitution into the normalized bounds gives

$$|z_n| \leq C\frac{\kappa\delta}{\rho}, \qquad |a_n| + |b_n| \leq C\delta \quad \text{when} \quad \sum_{i<n} \eta_i \leq c\frac{\rho}{\ell\kappa}.$$

Within the stated cumulative-step range, if also $\sum_{i<n} \eta_i \geq C\kappa/\ell$, then $\sum_{i<n} h_i \geq \tau_-$ and $z_n \geq c\kappa\delta/\rho$. The inequality $c/\alpha_\star \geq 2(C + 2c)$ makes the cumulative-step interval nonempty and leaves room for the first-crossing correction in the single-noise proof. No monotonicity of the steps was used. □

C.3.2 A SMOOTH INNER-PL GAME WITH THE SAME PRIMAL DISPLACEMENT

We retain the quadratic updates on the region containing the relevant inner iterates, and set the nonlinear shift to zero outside a larger region. Fix R_0 larger than the inner-coordinate bound in Lemma 5.4, and choose a smooth cutoff $\chi : \mathbb{R} \to [0, 1]$ equal to one on $[-R_0, R_0]$, zero outside $[-2R_0, 2R_0]$, with $|\chi'| \leq 1/4$ and bounded first two derivatives. The constant R_0 is fixed before all smoothness and PL constants below.

For a provisional condition parameter $K \geq 1$, define the residual $r_K : \mathbb{R}^3 \to \mathbb{R}$ and game $g_K : \mathbb{R} \times \mathbb{R}^3 \to \mathbb{R}$ by

$$r_K(a, b, s) = a + b + \sqrt{K}\,\chi\left(\frac{a+b}{\sqrt{K}}\right)\sin^2\left(\frac{\pi s}{2}\right), \qquad g_K(u; a, b, s) = ua - \frac{a^2 + r_K(a, b, s)^2}{2K}. \quad \text{(C.33)}$$

The profile vanishes at $s = 0$ and has zero derivative there. At $s = \pm 1$, it supplies the same shift and again has zero derivative. Thus a gradient error can create a nonzero shift from an initialization with zero inner gap.

Lemma C.3 (Smoothness and inner PL of the single-perturbation game)**.** *There are universal $c, C > 0$ such that, for every $K \geq 1$, g_K has a C-Lipschitz gradient,*

$$\max_{a,b,s} g_K(u; a, b, s) = \frac{K}{2}u^2, \qquad \frac{1}{2}\left\|\nabla_{a,b,s} g_K\right\|^2 \geq \frac{c}{K}\left[\frac{K}{2}u^2 - g_K(u; a, b, s)\right]. \quad \text{(C.34)}$$

The origin is stationary.

Proof. Completing the square gives

$$g_K(u; a, b, s) = \frac{K}{2}u^2 - \frac{(a - Ku)^2 + r_K(a, b, s)^2}{2K}. \quad \text{(C.35)}$$

Both residuals vanish at $(a, b, s) = (Ku, -Ku, 0)$, proving the outer-objective formula and attainment. For the PL inequality, write $r = r_K(a, b, s)$, $e = a - Ku$, and

$$p = \partial_a r_K = \partial_b r_K = 1 + \chi'\left(\frac{a+b}{\sqrt{K}}\right)\sin^2\left(\frac{\pi s}{2}\right) \in [3/4, 5/4].$$

The two gradient coordinates are $\partial_a g_K = -(e+pr)/K$ and $\partial_b g_K = -pr/K$. Since $e = (e+pr) - pr$ and $r^2 \le 16(pr)^2/9$,

$$e^2 + r^2 \le 2(e+pr)^2 + \left(2 + \frac{16}{9}\right)(pr)^2 \le 4\big[(e+pr)^2 + (pr)^2\big], \tag{C.36}$$

$$\frac{1}{2}\left\|\nabla_{a,b,s} g_K\right\|^2 \ge \frac{e^2+r^2}{8K^2} = \frac{1}{4K}\left[\frac{K}{2}u^2 - g_K(u;a,b,s)\right]. \tag{C.37}$$

This proves the inner-PL inequality.

For smoothness, $|\partial_a r_K|, |\partial_b r_K| \le C$ and $|\partial_s r_K| \le C\sqrt{K}$. Wherever a second derivative is nonzero, the cutoff implies $|a+b| \le 2R_0\sqrt{K}$ and $|r_K| \le C\sqrt{K}$. Direct differentiation gives second-derivative bounds of order $K^{-1/2}$ in (a,b), order one for mixed derivatives with s, and order $K^{1/2}$ for $\partial_{ss} r_K$. Hence, for $y = (a,b,s)$,

$$\left\|\nabla_y^2 (r_K^2/(2K))\right\| \le \frac{\left\|\nabla_y r_K\right\|^2 + |r_K|\left\|\nabla_y^2 r_K\right\|}{K} \le C. \tag{C.38}$$

Outside that region the residual is affine, so the same bound holds globally. The remaining terms $ua - a^2/(2K)$ also have bounded Hessian for $K \ge 1$. All first derivatives vanish at the origin. □

At $s = \pm 1$ and $|a+b| \le R_0\sqrt{K}$, the residual satisfies

$$r_K(a,b,s) = a + b + \sqrt{K}, \qquad \partial_s r_K(a,b,s) = 0. \tag{C.39}$$

Thus s stays fixed and the other coordinates follow the shifted quadratic of Lemma 5.4. The resulting inner gap is $1/2$, whereas perturbing b by the same amount at $s = 0$ produces only $1/(2K)$. The inner quadratic in (a,b) retains curvature of order $1/K$.

C.3.3 Changing variables and the contribution of one perturbation

Let c_0, C_0 be the constants in Lemma C.3, and let $\alpha_\star$ be the refined constant fixed in the proof of Lemma 5.4. Choose universal $\gamma, \theta > 0$, fixed after the low-ratio construction, such that

$$\theta C_0 \le 1, \qquad \frac{\theta c_0}{\gamma} \ge 1. \tag{C.40}$$

For condition number κ, set

$$K := \gamma\kappa, \qquad \lambda := \theta\ell. \tag{C.41}$$

In (5.11), set $g_{\sigma,j} \equiv 0$ whenever $\eta_j\sigma = 0$; these components have zero gradients and gaps. Increase the universal lower bound on κ so that $K \ge 1$. For $\sigma > 0$ and a block j with $\eta_j > 0$, write $y[j] = (a[j], b[j], s[j])$. The game in Lemma 5.5 is

$$g_{\sigma,j}(v[j]; a[j], b[j], s[j]) := \lambda(\eta_j\sigma)^2 g_K\left(\frac{v[j]}{\eta_j\sigma}; \frac{a[j]}{\eta_j\sigma}, \frac{b[j]}{\eta_j\sigma}, \frac{s[j]}{\eta_j\sigma}\right). \tag{C.42}$$

The squared amplitude compensates exactly for the spatial rescaling. Writing $w \in \mathbb{R}^4$ for the component coordinates, we have

$$\nabla_w^2\left[(\eta_j\sigma)^2 g_K\left(\frac{w}{\eta_j\sigma}\right)\right] = \nabla^2 g_K\left(\frac{w}{\eta_j\sigma}\right). \tag{C.43}$$

Thus arbitrarily small positive noise displacements do not increase the Hessian bound. Define the shift as a function of the perturbed inner coordinate

$$\delta_j(s) = \sqrt{K}\,\eta_j\sigma \sin^2\left(\frac{\pi s}{2\eta_j\sigma}\right). \tag{C.44}$$

The profile and its derivative vanish at zero, while $\delta_j(\pm\eta_j\sigma) = \sqrt{K}\,\eta_j\sigma$ and $\delta_j'(\pm\eta_j\sigma) = 0$. Thus either sign of the gradient error creates the same fixed shift. With ℓ and κ in (5.6) replaced by λ and K, respectively, the game in (C.42) is that shifted quadratic with displacement $\chi((a[j] + b[j])/(\sqrt{K}\eta_j\sigma))\delta_j(s[j])$. The cutoff bounds second derivatives outside the region containing the relevant iterates; the game equals g_δ in (5.6), under the same parameter substitution and with $\delta = \delta_j(s[j])$ whenever $|a[j] + b[j]| \le R_0\sqrt{K}\eta_j\sigma$.

Its full-gradient Lipschitz constant is at most $\lambda C_0 \le \ell$, and its inner-PL parameter is at least $\lambda c_0/K \ge \ell/\kappa$. Also $\Phi_{\sigma,j}(v[j]) = \lambda K(v[j])^2/2$, and the origin is both an outer minimizer and an inner maximizer.

Lemma 5.5 (Contribution of one inner-gradient perturbation)**.** *There are universal $c, C, c_{\mathrm{step}} > 0$ such that, under Theorem 5.1's parameter conditions with $\sigma > 0$, $T \geq 2$, any predetermined steps $\eta_0 \geq \cdots \geq \eta_{T-2} \geq 0$ and index $0 \leq j \leq T-2$ with $0 < \ell\eta_j \leq c_{\mathrm{step}}\kappa$ admit an ℓ-smooth NC-(ℓ/κ)-PL component $g_{\sigma,j} : \mathbb{R} \times \mathbb{R}^3 \to \mathbb{R}$. Its quadratic outer objective $\Phi_{\sigma,j}$ has curvature $\Phi''_{\sigma,j} \asymp \ell\kappa$, and both gaps vanish at the origin. Started there with one equiprobable $\pm\sigma$ error in the third inner gradient at update j and exact gradients otherwise, both SGDA$_{\mathrm{Sim}}$ and SGDA$_{\mathrm{Alt}}$ generate outer iterates $v_t[j]$ satisfying*

$$\sum_{t<T} \mathbb{E}|\Phi'_{\sigma,j}(v_t[j])|^2 \geq c\frac{\ell\kappa^4\sigma^2\eta_j}{\rho} \quad \text{if} \quad \sum_{t=j+1}^{T-2} \eta_t \geq C\frac{\rho}{\ell\kappa}. \tag{5.10}$$

Proof. The cutoff game and its exact normalization were defined above. The inner-gradient error creates the fixed shift in Lemma 5.4. Non-increasing step sizes then give a lower bound on the number of iterates with a large primal displacement.

Let ξ take values ± 1 with equal probability. All coordinates remain zero until update j, which sends $s[j]$ to $\eta_j\sigma\xi$ and leaves the other coordinates zero. In the variables of g_K, this is $s = \xi \in \{-1, 1\}$. Since the sine-squared profile equals one and has derivative zero at $s = \pm 1$, both signs produce the same active trajectory and s stays fixed.

Write c_r, C_r for c, C in Lemma 5.4, and retain the refined $\alpha_\star$ from its proof. Choose

$$c_{\mathrm{high}} = \frac{\gamma^2}{\alpha_\star}, \qquad c_{\mathrm{step}} = \frac{c_r\gamma}{\theta},$$

and take the constant in the subsequent-sum condition at least $c_r/(\theta\gamma)$. Index the state after the perturbation by $n = 0$ (original time $j+1$), put $N = T-j-2$, and define

$$\alpha = \frac{K^2}{\rho}, \qquad \tau_n = \frac{\lambda}{K}\sum_{t=j+1}^{j+n} \eta_t \quad (0 \leq n \leq N).$$

An empty sum is zero. The hypotheses give $\alpha \leq \alpha_\star$, increments at most $\lambda\eta_j/K \leq c_r$, and $\tau_N \geq c_r/\alpha$. Dividing all component coordinates by $\eta_j\sigma$ gives the updates of λg_K. The gradient error produces the fixed shift $\sqrt{K}$ in these coordinates, so Lemma 5.4 applies with parameters $(\lambda, K, \sqrt{K})$. The corresponding outer coordinate u_n satisfies

$$u_n \geq c_r\frac{K^{3/2}}{\rho} \quad \text{when} \quad C_r \leq \tau_n \leq \frac{c_r}{\alpha}, \qquad |a_n| + |b_n| \leq C_r\sqrt{K} \quad \text{when} \quad \tau_n \leq \frac{c_r}{\alpha}.$$

Let i and k be the first indices with $\tau_i \geq C_r$ and $\tau_k \geq c_r/\alpha$. By the choice of $\alpha_\star$ in the proof of Lemma 5.4, $c_r/\alpha \geq 2(C_r + 2c_r)$, so $i < k \leq N$. We first identify this quadratic trajectory with the smooth game. For $0 \leq n < k$, the inner-coordinate bound gives $|a_n + b_n| \leq C_r\sqrt{K}$. The cutoff radius $R_0 > C_r$ was fixed before the smoothness constants and numerical rescalings. Induction therefore keeps the cutoff identically one at every gradient query needed to generate iterate $k-1$, including the intermediate primal query of SGDA$_{\mathrm{Alt}}$, because it depends only on $a+b$. Thus the displacement estimate holds for the actual smooth game before we count the iterates satisfying the lower bound.

Monotonicity and the first-crossing overshoot give

$$(k-i)\frac{\lambda\eta_j}{K} \geq \tau_k - \tau_i \geq \frac{c_r}{\alpha} - C_r - c_r \geq \frac{c_r}{2\alpha}. \tag{C.45}$$

For $i \leq n < k$, the displacement bound applies, and the corresponding original iteration indices $j+1+n$ are smaller than T. The raw outer objective $\Phi_K(u) = Ku^2/2$ therefore has squared gradient at least $c_r^2K^5/\rho^2 = c_r^2K\alpha^2$. Returning to the component multiplies this gradient by $\lambda\eta_j\sigma$, and hence

$$\sum_{t<T} \mathbb{E}|\Phi'_{\sigma,j}(v_t[j])|^2 \geq c\lambda^2(\eta_j\sigma)^2K\alpha^2\frac{K}{\alpha\lambda\eta_j} = c\frac{\lambda K^4\sigma^2\eta_j}{\rho} \geq c'\frac{\ell\kappa^4\sigma^2\eta_j}{\rho}. \tag{C.46}$$

All constants are independent of the horizon and step-size sequence, and the same displacement estimate holds for either noise sign. Appendix C.4 supplies a fixed oracle for the prescribed single-gradient perturbations. □

The preceding estimate applies only to perturbations with enough subsequent cumulative step. The excluded final perturbations carry a bounded total step size, so their removal leaves a constant fraction of S when S is sufficiently large.

Lemma C.4 (Lower bound from the noise components)**.** *There are universal $c, C > 0$ with the following property. Fix $\ell, \sigma > 0$, $\kappa \geq \kappa_{\mathrm{high}}$, $\rho \geq \max\{1, c_{\mathrm{high}}\kappa^2\}$, an integer $T \geq 2$, a predetermined schedule $\eta_0 \geq \cdots \geq \eta_{T-2} \geq 0$, and $m \in \{0, \ldots, T-1\}$ with $\ell\eta_t < 8$ for $m \leq t \leq T-2$. Set $S = \sum_{t=m}^{T-2} \eta_t$. The components in (5.11) can be chosen ℓ-smooth and NC-(ℓ/κ)-PL, with $\Phi_{\sigma,j}(0) = \min \Phi_{\sigma,j} = g_{\sigma,j}(0,0) = 0$. Both SGDA$_{\mathrm{Sim}}$ and SGDA$_{\mathrm{Alt}}$ on the noise component g_σ, initialized at the origin with exact gradients except an equiprobable $\pm\sigma$ error in the gradient for component j's third inner coordinate at update j, independently across components, satisfy*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla\Phi_\sigma(v_t)\right\|^2 \geq c\frac{\kappa^4\ell\sigma^2 S}{\rho T} \quad \text{if } S \geq \frac{C\rho}{\kappa\ell}. \tag{C.47}$$

Proof. Set $g_{\sigma,j} \equiv 0$ for $\eta_j = 0$, and choose each positive-step block of g_σ from Lemma 5.5. Since the coordinates are disjoint, each nonzero block follows its single-perturbation trajectory and

$$\left\|\nabla\Phi_\sigma(v_t)\right\|^2 = \sum_{j\in J} \left\|\nabla\Phi_{\sigma,j}(v_t[j])\right\|^2.$$

Let C_{one} be the remaining-step constant in Lemma 5.5, and enlarge κ_{high} so that $8 \leq c_{\mathrm{step}}\kappa$. Every positive step with index $t \geq m$ then satisfies its step-size condition. Let $J_{\mathrm{good}} := \{j \in J : \sum_{t=j+1}^{T-2} \eta_t \geq C_{\mathrm{one}}\rho/(\kappa\ell)\}$. Non-increasingness ensures $\eta_j > 0$ for every $j \in J_{\mathrm{good}}$. If $J \setminus J_{\mathrm{good}}$ is nonempty, its first index j_0 satisfies

$$\sum_{j\in J\setminus J_{\mathrm{good}}} \eta_j \leq \sum_{t=j_0}^{T-2} \eta_t < \frac{8}{\ell} + \frac{C_{\mathrm{one}}\rho}{\kappa\ell} \leq \frac{C_1\rho}{\kappa\ell}, \qquad C_1 := C_{\mathrm{one}} + \frac{8}{c_{\mathrm{high}}},$$

where the last inequality uses $\rho \geq c_{\mathrm{high}}\kappa^2$ and $\kappa \geq 1$. If the excluded set is empty, its total step size is zero. Taking $C := 2C_1$, the condition $S \geq C\rho/(\kappa\ell)$ implies $\sum_{j\in J_{\mathrm{good}}} \eta_j \geq S/2$. Summing (5.10) over these indices gives

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla\Phi_\sigma(v_t)\right\|^2 \geq c\frac{\kappa^4\ell\sigma^2}{\rho T} \sum_{j\in J_{\mathrm{good}}} \eta_j \geq \frac{c}{2}\frac{\kappa^4\ell\sigma^2 S}{\rho T}.$$

Reducing the universal constant c proves the claim. □

C.4 Proof of the lower bound for f_{stable}

The component bounds must hold along one run with a fixed oracle whose total variance is at most σ^2. The deterministic corridor coordinate identifies the update at which each component receives its perturbation. We then combine the aggregate noise bound with the corridor bound.

Lemma 5.6 (Stable-step lower bound and admissibility)**.** *Under Theorem 5.1's conditions, let m count the steps $\eta_t \geq 8/\ell$ in a predetermined nonnegative non-increasing sequence, and assume $m + 1 < T/2$. For each of SGDA$_{\mathrm{Sim}}$ and SGDA$_{\mathrm{Alt}}$, the game f_{stable} given by (5.4) and a fixed oracle can be chosen as an instance in $\mathcal{C}(\ell, \kappa, \Delta, \sigma)$ with initial outer gap Δ and inner gap zero at the origin. Writing $S = \sum_{t=m}^{T-2} \eta_t$ and $\rho\Delta/0 := +\infty$, its iterates satisfy, for a universal $c > 0$,*

$$\frac{1}{T}\sum_{t=0}^{T-1} \mathbb{E}\left\|\nabla\Phi_{\mathrm{stable}}(u_t, v_t)\right\|^2 \geq c\min\left\{\ell\Delta, \max\left\{\frac{\rho\Delta}{S}, \frac{\kappa^4\ell\sigma^2 S}{\rho T}\right\}\right\}. \tag{5.12}$$

Proof. For the chosen algorithm, SGDA$_{\mathrm{Sim}}$ or SGDA$_{\mathrm{Alt}}$, use the predetermined non-increasing step sizes in the statement. Let $J = \{j : m \leq j \leq T-2\}$, and choose the corridor from Lemma 5.3. For $\sigma > 0$, choose the aggregate noise game g_σ from Lemma C.4; for $\sigma = 0$, set $g_{\sigma,j} \equiv 0$ for every $j \in J$ and use exact gradients. Use these games in (5.4) and initialize every coordinate at zero. In either case, the dimensions are $d_x = 1 + |J| \leq T - m$ and $d_y = 3|J| \leq 3(T - m - 1)$.

We first verify the game assumptions. Each component has full-gradient Lipschitz constant at most ℓ. Since the components use disjoint coordinates, summing the squared component inequalities gives

$$\|\nabla f_{\text{stable}}(u,v;y) - \nabla f_{\text{stable}}(\widetilde{u},\widetilde{v};\widetilde{y})\|^2 \le \ell^2\big(|u-\widetilde{u}|^2 + \|v-\widetilde{v}\|^2 + \|y-\widetilde{y}\|^2\big).$$

Thus the smoothness assumption holds without a dimension factor. Each inner maximum is finite and attained, so their finite product attains the maximum of f_{stable}. Its outer objective is

$$\Phi_{\text{stable}}(u,v) = g_{\text{corridor}}(u) + \Phi_\sigma(v) = g_{\text{corridor}}(u) + \frac{\lambda K}{2} \sum_{\substack{j\in J\\ \eta_j\sigma>0}} (v[j])^2.$$

It is differentiable and lower bounded. Summing the component inner-PL inequalities gives

$$\begin{aligned}
\frac{1}{2}\|\nabla_y f_{\text{stable}}(u,v;y)\|^2 &= \frac{1}{2}\sum_{j\in J}\big\|\nabla_{y[j]} g_{\sigma,j}(v[j];y[j])\big\|^2 \\
&\ge \frac{\ell}{\kappa}\sum_{j\in J}\big(\Phi_{\sigma,j}(v[j]) - g_{\sigma,j}(v[j];y[j])\big) \\
&= \frac{\ell}{\kappa}\big(\Phi_{\text{stable}}(u,v) - f_{\text{stable}}(u,v;y)\big).
\end{aligned}$$

The corridor supplies outer gap Δ, each perturbation component has zero outer gap at the origin, and every component is initially maximized in its inner variables. Hence the initial outer and inner gaps are Δ and zero. These conclusions also hold for the corridor alone, with a vacuous inner-PL inequality.

For $\sigma > 0$, the corridor coordinate follows deterministic gradient descent and strictly increases at every positive update by Lemma 5.3. Its inner-query values $c_j = u_{j+1}$ for SGDA_{Alt} and $c_j = u_j$ for SGDA_{Sim} therefore distinguish all positive updates. Define $e : \mathbb{R} \to \mathbb{R}^{d_y}$ to be the unit vector in the $s[j]$-coordinate at c_j for each $j \in J$ with $\eta_j\sigma > 0$, and zero elsewhere. This finite-support map is measurable and has norm at most one. For $x = (u,v)$ and a uniform sign $\xi \in \{-1,1\}$, set

$$G_x(x,y;\xi) = \nabla_x f_{\text{stable}}(u,v;y), \qquad G_y(x,y;\xi) = \nabla_y f_{\text{stable}}(u,v;y) + \sigma\xi e(u).$$

With fresh independent signs at the oracle calls, this fixed oracle is pointwise unbiased and its total inner-gradient variance is $\sigma^2\|e(u)\|^2 \le \sigma^2$. At positive update j for $j \in J$, only the $s[j]$-coordinate receives the prescribed error; all other components use exact gradients. Repeated selected values at zero-step calls leave the iterates unchanged. Consequently, each nonzero component follows the trajectory in Lemma 5.5, while the zero components remain at the origin. The combined origin-initialized instance belongs to $\mathcal{C}(\ell,\kappa,\Delta,\sigma)$.

It remains to prove the lower bound. Since the squared outer gradients add and $(T-m)/T > 1/2$, the corridor gives the required conclusion when $\sigma = 0$, and gives $c\ell\Delta$ when $S = 0$. Suppose $\sigma, S > 0$, and let C be the threshold constant in Lemma C.4. Increase κ_{high} so that $\kappa_{\text{high}} \ge C$. If $S < C\rho/(\kappa\ell)$, then $\rho\Delta/S \ge \ell\Delta$, so the corridor again gives $c\ell\Delta$ and proves the capped bound. Otherwise Lemma C.4 applies, and the two component bounds give

$$\begin{aligned}
\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\|\nabla\Phi_{\text{stable}}(u_t,v_t)\|^2 &\ge c\frac{T-m}{T}\min\left\{\ell\Delta, \frac{\rho\Delta}{S}\right\}, \\
\frac{1}{T}\sum_{t=0}^{T-1}\mathbb{E}\|\nabla\Phi_{\text{stable}}(u_t,v_t)\|^2 &\ge c\frac{\kappa^4\ell\sigma^2 S}{\rho T}.
\end{aligned} \tag{C.48}$$

Their maximum is at least a constant multiple of the right-hand side of (5.12). Reducing its universal constant proves the lemma. □

D PROOFS FOR THE LOW-RATIO LOWER BOUND

We prove the two properties used by Theorem 6.3: small steps preserve a region of nonzero outer gradient, and the initially larger steps can be made to end inside that region. The first requires a smooth game with a stable periodic solution; the second modifies that specific game away from the solution. The final rescaling and initialization checks appear in Appendix A.3.

D.1 CONSTRUCTION OF g_{cycle}

Lemma 6.1 (An attracting periodic orbit and its discrete dynamics)**.** *There are universal $C, \mu, h_\star > 0$ and a smooth residual R such that the base game $g_{\mathrm{cycle}} : \mathbb{R}^2 \times \mathbb{R}^4 \to \mathbb{R}$ in (6.2) is jointly C-smooth and affine in u, has outer objective $\Phi_{g_{\mathrm{cycle}}}(u) = \frac{1}{2}\|u\|^2$, and satisfies the inner μ-PL inequality. The path (6.3), with $y_\star = (a_\star, b_\star)$, traces a locally stable 2π-periodic orbit Γ of GDA flow. There is a compact neighborhood $\mathcal{U}$ of Γ such that both SGDA$_{\mathit{Sim}}$ and SGDA$_{\mathit{Alt}}$, started in $\mathcal{U}$ with exact gradients and equal primal and dual steps $h_t \in [0, h_\star]$, satisfy $(u_t, y_t) \in \mathcal{U}$ and $\left\|\nabla \Phi_{g_{\mathrm{cycle}}}(u_t)\right\| = \|u_t\| \geq \frac{1}{2}$ for all $t \geq 0$, where $y_t = (a_t, b_t)$.*

Proof. We verify the global geometry before studying the periodic trajectory. Uniform invertibility of the inner coordinate map gives the PL inequality; the support of its nonlinear part controls smoothness. We then substitute the prescribed curve and check the eigenvalues governing departures from it. A local Lyapunov function transfers this attraction to an invariant neighborhood for both discrete updates.

We construct $g_{\mathrm{cycle}}(u, a, b) = \sqrt{2}\, u^\mathsf{T} a - \|a\|^2 - \|R(a,b)\|^2 / 2$ on $\mathbb{R}^2 \times \mathbb{R}^2 \times \mathbb{R}^2$. Its affine dependence on u will be essential to the ratio scaling. The residual has the form $R(a,b) = b + \chi_b(b)\psi(a)$: the values and first derivatives of ψ produce the periodic solution, its second derivatives control departures from that solution, and the cutoffs enforce global smoothness.

Let $J = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ be rotation by a right angle. The target path is $u_\star(t) = (\cos t, \sin t)^\mathsf{T}$, with $a_\star = -Ju_\star/\sqrt{2}$ and $b_\star = u_\star/\sqrt{2}$. Since $\dot{u} = -\sqrt{2}a$ and, where the cutoffs equal one, $\dot{b} = -(b + \psi(a))$, the first and third equations require $\psi(a_\star) = a_\star - Ja_\star$. The remaining equation requires $D\psi(a_\star)^\mathsf{T} a_\star = -2a_\star + Ja_\star$. We satisfy these value and derivative conditions first, then check contraction through the remaining derivatives.

Choose a smooth cutoff $\chi_a : [0, \infty) \to [0, 1]$ that equals one on a neighborhood of $\frac{1}{2}$ and vanishes outside a slightly larger compact interval. Define the polynomials

$$c(q) := 1 - 3(q - \tfrac{1}{2}) + \frac{7}{2}(q - \tfrac{1}{2})^2, \qquad d(q) := -1 + (q - \tfrac{1}{2}), \tag{D.1}$$

and set

$$\psi(a) := \chi_a(\|a\|^2)\left[c(\|a\|^2)a + d(\|a\|^2)Ja\right]. \tag{D.2}$$

The orbit conditions require $c(1/2) = 1$, $d(1/2) = -1$, and $c'(1/2) = -3$. The remaining choices $c''(1/2) = 7$ and $d'(1/2) = 1$ make the transverse linearization stable, as verified below. The map ψ is smooth and compactly supported. Choose a smooth radial cutoff $\chi_b : \mathbb{R}^2 \to [0, 1]$ that equals one on a neighborhood of the circle $\|b\| = 1/\sqrt{2}$, vanishes outside a sufficiently large ball, and satisfies

$$\|\psi\|_\infty \|\nabla \chi_b\|_\infty \leq \frac{1}{2}. \tag{D.3}$$

Define

$$R(a,b) := b + \chi_b(b)\psi(a), \tag{D.4}$$

and

$$g_{\mathrm{cycle}}(u, a, b) := \sqrt{2}\, u^\mathsf{T} a - \|a\|^2 - \frac{1}{2}\|R(a,b)\|^2. \tag{D.5}$$

For every fixed a and target $w \in \mathbb{R}^2$, the map $b \mapsto w - \chi_b(b)\psi(a)$ is a contraction with factor at most $1/2$ by (D.3). Its unique fixed point solves $R(a,b) = w$. Thus $b \mapsto R(a,b)$ is bijective. Moreover,

$$D_b R = I + \psi(a)\nabla\chi_b(b)^\mathsf{T}, \qquad \|D_b R - I\| \leq \tfrac{1}{2}, \qquad \left\|(D_b R)^{-1}\right\| \leq 2.$$

The inverse function theorem gives a smooth inverse. Boundedness of $D_a R$ and $(D_b R)^{-1}$ then bounds the inverse derivative of the triangular map $(a, b) \mapsto (a, R(a, b))$ uniformly.

All derivatives of ψ are bounded. Whenever a second derivative of R is nonzero, both a and b lie in fixed compact sets; elsewhere R is affine. Hence DR and RD^2R are globally bounded, and so is $\nabla^2 \|R\|^2/2$. The full gradient of g_{cycle} is therefore universally Lipschitz.

For each a, maximizing over b cancels the residual. Maximizing over a then gives $\Phi_{g_{\mathrm{cycle}}}(u) = \|u\|^2/2$ and

$$\Phi_{g_{\mathrm{cycle}}}(u) - g_{\mathrm{cycle}}(u, a, b) = \frac{1}{2}\left\|u - \sqrt{2}a\right\|^2 + \frac{1}{2}\|R(a, b)\|^2 . \tag{D.6}$$

Applying Lemma B.1 to the outer vector $(u, 0)$ and inner map $(a, b) \mapsto (\sqrt{2}a, R(a, b))$ gives the inner μ-PL inequality for a universal $\mu > 0$.

Along (6.3), $b_\star = Ja_\star$, $\|a_\star\|^2 = \frac{1}{2}$, and both cutoffs are constant. Hence

$$\psi(a_\star) = a_\star - Ja_\star, \qquad R(a_\star, b_\star) = a_\star.$$

On the neighborhood where both cutoff functions equal one, writing $q = \|a\|^2$,

$$D\psi(a)^\mathsf{T} a = [c(q) + 2qc'(q)]a - d(q)Ja. \tag{D.7}$$

The values $c(\frac{1}{2}) = 1$, $c'(\frac{1}{2}) = -3$, and $d(\frac{1}{2}) = -1$ give

$$\dot{u}_\star = -\sqrt{2}a_\star = Ju_\star, \quad \dot{a}_\star = \sqrt{2}u_\star - 2a_\star - D\psi(a_\star)^\mathsf{T} a_\star = Ja_\star, \quad \dot{b}_\star = -R(a_\star, b_\star) = Jb_\star.$$

Thus (6.3) solves the GDA differential equation and is 2π-periodic.

Along this solution, $-\nabla_u g_{\mathrm{cycle}} = Ju_\star$, $\nabla\Phi_{g_{\mathrm{cycle}}} = u_\star$, and $\|\nabla_{a,b} g_{\mathrm{cycle}}\|^2 = \|\dot{a}_\star\|^2 + \left\|\dot{b}_\star\right\|^2 = 1$. The derivative of the inner gap is therefore zero; this identity explains the persistent tracking error but is not an additional hypothesis.

The definitions imply $\psi(e^{tJ}a) = e^{tJ}\psi(a)$ and $R(e^{tJ}a, e^{tJ}b) = e^{tJ}R(a, b)$. Thus the derivative calculation at any point of the periodic solution is the rotation of the calculation at $t = 0$. Direct differentiation at $t = 0$, using $c''(\frac{1}{2}) = 7$ and $d'(\frac{1}{2}) = 1$, gives

$$A_\star := D\psi(a_\star(0)) = \begin{pmatrix} 1 & 0 \\ -1 & -2 \end{pmatrix}, \qquad S_\star := D_a[D\psi(a)^\mathsf{T} R(a, b)]_{(a_\star(0), b_\star(0))} = \begin{pmatrix} -1 & 3 \\ 3 & 2 \end{pmatrix}. \tag{D.8}$$

Write $z = (u, a, b)$ and $z_\star(t) = (u_\star(t), a_\star(t), b_\star(t))$. Apply the time-dependent change of variables $\widetilde{z}(t) = \mathcal{Q}_t^\mathsf{T} z(t)$ with $\mathcal{Q}_t = \mathrm{diag}(e^{tJ}, e^{tJ}, e^{tJ})$. Along the periodic solution, $\widetilde{z}(t) = z_\star(0)$ is constant, and the Jacobian of the transformed differential equation at this fixed point is

$$M_\star = \begin{pmatrix} -J & -\sqrt{2}I_2 & 0 \\ \sqrt{2}I_2 & -2I_2 - S_\star - J & -A_\star^\mathsf{T} \\ 0 & -A_\star & -I_2 - J \end{pmatrix}. \tag{D.9}$$

A direct determinant calculation gives

$$\det(\lambda I - M_\star) = \lambda(\lambda^5 + 7\lambda^4 + 7\lambda^3 + 21\lambda^2 + 10\lambda + 4). \tag{D.10}$$

The Routh–Hurwitz criterion determines the number of roots in the open right half-plane from the sign changes in the first column of the Routh array. For the quintic factor, this column is $1, 7, 4, 9/2, 370/63, 4$. It has no sign change, so the quintic has no root in the open right half-plane. Moreover, the Routh array has neither a zero first-column entry nor an all-zero row, which rules out roots on the imaginary axis. Hence all five roots have strictly negative real parts. The remaining eigenvalue is zero and has eigenvector $\dot{z}_\star(0)$. Since $\dot{z}_\star(0)$ is the derivative of the periodic solution with respect to time, a small displacement in this direction changes the time label from $z_\star(t)$ to $z_\star(t + s)$ to first order and remains on the same periodic solution.

To control departures from the orbit, choose a basis $S \in \mathbb{R}^{6\times 5}$ of the stable invariant subspace of $M_\star$, and let A_- be its restriction in this basis. For sufficiently small $\varepsilon > 0$, the map $\mathcal{X}(\vartheta, \xi) = \mathcal{Q}_\vartheta(z_\star(0) + S\xi)$, with $\vartheta \in \mathbb{R}/(2\pi\mathbb{Z})$ and $\xi \in \mathbb{R}^5$, $\|\xi\| < \varepsilon$, gives smooth coordinates near Γ: the

columns of S complement its tangent direction, and the orbit is compact and embedded. Rotational equivariance gives $\dot{\xi} = A_-\xi + O(\|\xi\|^2)$ uniformly in ϑ. Since A_- is Hurwitz, choose $P \succ 0$ satisfying $A_-^\mathsf{T}P + PA_- = -I$, and define $V(\mathcal{X}(\vartheta,\xi)) = \xi^\mathsf{T}P\xi$. On a smaller neighborhood,

$$V(z) \asymp \operatorname{dist}(z,\Gamma)^2, \qquad \dot{V} = -\|\xi\|^2 + O(\|\xi\|^3) \le -cV$$

for a universal $c > 0$. This proves local exponential orbital stability.

Choose a fixed small $\delta > 0$ so that $\mathcal{U} := \{V \le \delta\}$ is compactly contained in this neighborhood and in $\{\|u\| \ge 1/2\}$. Let $F = (-\nabla_u g_{\text{cycle}}, \nabla_{a,b} g_{\text{cycle}})$. Simultaneous GDA gives $z^+_{\text{Sim}} = z + hF(z)$, while primal-first alternating GDA gives $z^+_{\text{Alt}} - z^+_{\text{Sim}} = (0, -2h^2a, 0)$. For $z \in \mathcal{U}$ and sufficiently small h, both updates, the line segments joining them to z, and the intermediate primal queries remain in a fixed compact neighborhood where V is smooth. Taylor expansion therefore gives, uniformly for either update,

$$V(z^+) \le V(z) + h\nabla V(z)^\mathsf{T}F(z) + Ch^2 \le (1-ch)V(z) + Ch^2 \le (1-ch)\delta + Ch^2 \le \delta,$$

where we choose $h_\star > 0$ small enough that $ch_\star \le 1$ and $Ch_\star \le c\delta$. Thus $\mathcal{U}$ is invariant for every $h \in [0, h_\star]$, including $h = 0$. Induction and $\nabla\Phi_{g_{\text{cycle}}}(u) = u$ give the discrete conclusion for every such step sequence. All constants depend only on the fixed game.

For the next lemma, fix a bounded open neighborhood containing $\mathcal{U}$ and all intermediate primal queries $(u - h\nabla_u g_{\text{cycle}}(u,y), y)$ for $(u,y) \in \mathcal{U}$ and $h \in [0, h_\star]$. We protect this entire neighborhood from the local modifications, so both update rules remain unchanged there. □

D.2 Construction of f_{cycle}

To construct f_{cycle}, we first extend the base game to accommodate finitely many initial steps $h_t > h_\star$. At the initial query states y_t to be chosen, for $0 \le t < m$ and $e_1 = (1,0)^\mathsf{T}$, the required gradients are

$$\nabla_u \tilde{g}_{\text{cycle}}(e_1, y_t) = 0, \qquad \nabla_y \tilde{g}_{\text{cycle}}(e_1, y_t) = \frac{y_{t+1} - y_t}{h_t}. \tag{D.11}$$

They keep the outer variable fixed and send each inner state to the next, ending on the cycle. A local change of inner coordinates prescribes these gradients through its derivative. We preserve the first two coordinates, a, because they determine the primal gradient $\sqrt{2}a$. The following lemma realizes one derivative without changing the map elsewhere; disjoint modification regions will keep the bounds independent of the number of steps. Appendix A.3 then rescales the game to the desired parameters.

Lemma D.1 (Local derivative modification that leaves the first two coordinates unchanged)**.** *Fix $M < \infty$ and $\delta > 0$. There is $C = C(M,\delta)$ such that, for every $z \in \mathbb{R}^5$ and every*

$$N = \begin{pmatrix} I_2 & 0 \\ P & Q \end{pmatrix}, \qquad \det N > 0, \qquad \|N\| + \left\|N^{-1}\right\| \le M, \tag{D.12}$$

there is a smooth bijection $G : \mathbb{R}^5 \to \mathbb{R}^5$ with a smooth inverse. For every w, the first two coordinates of $G(w)$ equal those of w. Moreover, $G(z) = z$, $DG(z) = N$, G equals the identity outside $B(z,\delta)$, G maps this ball onto itself, and

$$\|DG\|_\infty + \left\|DG^{-1}\right\|_\infty + \left\|D^2G\right\|_\infty \le C. \tag{D.13}$$

Proof. We join the identity matrix to N through uniformly invertible matrices, then realize this matrix path as the derivative of a compactly supported flow. Since $\det N > 0$ and the first diagonal block is I_2, the matrix Q is invertible and has positive determinant. Write $Q = US$, where $U^\mathsf{T}U = I$, $\det U = 1$, and S is symmetric positive definite. Choose a skew-symmetric matrix A with $e^A = U$ and $\|A\| \le \pi$, and define

$$Q_\tau = e^{\tau A}((1-\tau)I + \tau S), \qquad N_\tau = \begin{pmatrix} I_2 & 0 \\ \tau P & Q_\tau \end{pmatrix}.$$

The matrices N_τ, their inverses, and their derivatives are uniformly bounded in terms of M. Put $A_\tau = \dot{N}_\tau N_\tau^{-1}$. Choose a smooth radial cutoff χ that equals one on $B(z,\delta/3)$ and vanishes outside $B(z,2\delta/3)$, and set

$$X_\tau(w) = \chi(w)A_\tau(w - z).$$

This time-dependent differential equation has zero first two components and uniformly bounded first two derivatives. Let $G_\tau(w)$ denote its solution at time τ starting from w at time zero. Then G_τ fixes z, preserves the first two coordinates, and satisfies

$$\frac{d}{d\tau} DG_\tau(z) = A_\tau DG_\tau(z), \qquad DG_0(z) = I,$$

so $DG_\tau(z) = N_\tau$. Differentiating this differential equation once and twice with respect to the initial point, and then applying Gronwall's inequality, gives (D.13) for $G = G_1$ and its inverse. Since the vector field vanishes in a neighborhood of the boundary, the ball is mapped onto itself. □

Lemma 6.2 (Large initial steps reach the cycle)**.** *There are universal $C, \mu > 0$ with the following property. For any finite sequence of steps $h_0, \dots, h_{m-1} > h_\star$, including $m = 0$, we can modify the game g_{cycle} in Lemma 6.1 to obtain $\tilde{g}_{\mathrm{cycle}} : \mathbb{R}^2 \times \mathbb{R}^5 \to \mathbb{R}$. This game is jointly C-smooth and affine in u, satisfies the inner μ-PL inequality, and has outer objective $\Phi_{\tilde{g}_{\mathrm{cycle}}}(u) = \|u\|^2/2$. With $y = (a, b, r)$ and $e_1 = (1, 0)^\mathsf{T}$, there is an initialization (e_1, y_0) with inner gap at most C. From this initialization, both SGDA$_{\mathit{Sim}}$ and SGDA$_{\mathit{Alt}}$, using exact gradients and equal primal and dual steps, keep $u_t = e_1$ for $0 \le t \le m$ and reach $y_m = (y_\star(0), 0)$. For every continuation with $h_t \in [0, h_\star]$, each algorithm follows its g_{cycle} dynamics in (u, a, b), while r remains zero.*

Proof. We choose separated inner states whose required update vectors are uniformly bounded. At each state, a local coordinate modification realizes the desired gradient while preserving the first two inner coordinates. The residual representation then verifies inner PL, and the disjoint modification regions give uniform smoothness.

The case $m = 0$ uses $\tilde{g}_{\mathrm{cycle}}(u, v, r) = g_{\mathrm{cycle}}(u, v)$ and $y_0 = (y_\star(0), 0)$. All initialization quantities are bounded on the fixed periodic solution. Suppose $m \ge 1$ and write $y = (a, b, r) \in \mathbb{R}^2 \times \mathbb{R}^2 \times \mathbb{R}$. The extra coordinate r separates the query states even when their first four coordinates agree. Its value is not itself penalized, so this separation need not increase the initial gap. Define

$$\begin{aligned} r_t &= -\sum_{j=t}^{m-1} h_j \quad (0 \le t < m), \quad r_m = 0, \\ y_t &= (0, -Be_1, r_t) \quad (t < m), \quad y_m = (a_\star(0), b_\star(0), 0). \end{aligned} \tag{D.14}$$

Here B is a fixed universal constant large enough that a neighborhood of $(a, b) = (0, -Be_1)$ is outside the support of the nonlinear part of R and separated from the projection onto inner variables of the protected neighborhood chosen above. Let $\bar{H}(a, b, r) = ((a, R(a, b)), r)$. At all y_t with $t < m$, it is the identity in a fixed neighborhood.

Put $v_t = (y_{t+1} - y_t)/h_t = (v_t^a, v_t^b, 1)$. Its first four coordinates are zero except at $t = m - 1$; their magnitude at the last step is bounded by a universal constant since $h_{m-1} > h_\star$ and the first four endpoint coordinates are fixed. Prescribe $\widetilde{H}(y_t) = y_t$ and

$$D\widetilde{H}(y_t) = \begin{pmatrix} I_2 & 0 & 0 \\ (v_t^a - \sqrt{2}e_1)^\mathsf{T}/B & (v_t^b)^\mathsf{T}/B & 1/B \\ 0 & e_2^\mathsf{T} & 0 \\ 0 & -Be_1^\mathsf{T} & 0 \end{pmatrix}. \tag{D.15}$$

The determinant after eliminating the first two identity rows is

$$\det \begin{pmatrix} v_{t,b_1}/B & v_{t,b_2}/B & 1/B \\ 0 & 1 & 0 \\ -B & 0 & 0 \end{pmatrix} = 1.$$

Both this matrix and its inverse are uniformly bounded. Distinct y_t with $t < m$ are separated by at least $h_\star$, because their scalar coordinates differ by sums of steps exceeding $h_\star$. Choose a universal ball radius smaller than one quarter of this separation and small enough that all balls avoid the protected neighborhood.

Apply Lemma D.1 in these disjoint balls. Each local map fixes its center, has the prescribed derivative, preserves the first two coordinates, and maps its ball onto itself. Let G be their composition and set

$\widetilde{H} = G \circ \bar{H}$. Since $\bar{H}$ is the identity near every center, (D.15) holds. At most one factor of G is nonidentity along the image of any point, so the derivative, inverse-derivative, and second-derivative bounds do not grow with m. Since $\bar{H}$ is bijective and equals the identity on each modification ball, no point outside those balls is mapped into them. Thus $\widetilde{H}$ agrees with $\bar{H}$ throughout the protected neighborhood, including a fixed interval of scalar coordinates around zero.

Let π_b project onto the two coordinates corresponding to b. Define the game and its residual map by

$$\tilde{g}_{\mathrm{cycle}}(u, y) = \left\langle (u, 0), \widetilde{T}(y) \right\rangle - \frac{1}{2} \left\| \widetilde{T}(y) \right\|^2, \qquad \widetilde{T}(a, b, r) = (\sqrt{2}a, \pi_b \widetilde{H}(a, b, r)). \tag{D.16}$$

Since $\widetilde{H}$ is onto and preserves its first two coordinates, $\widetilde{T} : \mathbb{R}^5 \to \mathbb{R}^4$ is onto. For $q = (q_a, q_b) \in \mathbb{R}^4$,

$$D\widetilde{T}(y)^{\mathsf{T}} q = D\widetilde{H}(y)^{\mathsf{T}}(\sqrt{2}q_a, q_b, 0), \qquad \left\| D\widetilde{T}(y)^{\mathsf{T}} q \right\| \geq c \left\| q \right\|,$$

where the last inequality uses the uniform inverse-derivative bound for $\widetilde{H}$. Lemma B.1 now gives $\max_y \tilde{g}_{\mathrm{cycle}}(u, y) = \|u\|^2 / 2$ and global inner PL with a universal constant.

The first two residual coordinates are affine. On the support of the original residual's nonlinear part, its value and the relevant derivatives are uniformly bounded. On each modification ball, the last two residual coordinates remain in a fixed bounded neighborhood of $-Be_1$, regardless of the potentially large scalar center r_t. Consequently $D\widetilde{T}$ and $\nabla^2 \left\| \widetilde{T} \right\|^2 / 2$ are uniformly bounded. Elsewhere those residual coordinates are affine. Expanding (D.16) gives $\tilde{g}_{\mathrm{cycle}}(u, a, b, r) = \sqrt{2}u^{\mathsf{T}}a - \|a\|^2 - \left\| \pi_b \widetilde{H}(a, b, r) \right\|^2 / 2$, proving affinity in u and universal full-gradient smoothness. Agreement with g_{cycle} and independence of r hold throughout the protected neighborhood.

Finally, at y_t with $t < m$ one has $\widetilde{T}(y_t) = (0, -Be_1)$. The prescribed derivative therefore gives exactly the target gradients in (D.11):

$$\nabla_u \tilde{g}_{\mathrm{cycle}}(e_1, y_t) = 0, \qquad \nabla_y \tilde{g}_{\mathrm{cycle}}(e_1, y_t) = D\widetilde{T}(y_t)^{\mathsf{T}}(e_1, Be_1) = v_t,$$

so both $\mathrm{SGDA}_{\mathrm{Sim}}$ and $\mathrm{SGDA}_{\mathrm{Alt}}$ follow the prescribed states. The initial inner gap is $(1 + B^2)/2$ and the initial primal gradient is zero. After time m, agreement with g_{cycle} throughout the protected neighborhood covers $\mathrm{SGDA}_{\mathrm{Sim}}$ and $\mathrm{SGDA}_{\mathrm{Alt}}$, including every intermediate primal query in $\mathrm{SGDA}_{\mathrm{Alt}}$. Independence of r there keeps $r_t = 0$, while Lemma 6.1 keeps $(u_t, a_t, b_t) \in \mathcal{U}$ for every small-step continuation. All constants depend only on the fixed small-step construction, not on the number or values of the prescribed steps. □

E Comparison with Smoothed-AGDA

We compare the expected squared outer gradient at a uniformly selected SGDA iterate with that at a point computed from a Smoothed-AGDA iterate pair. The mean-square estimates of Yang et al. (2022, Theorem 4.1, Proposition 2.1(b), and Corollary 4.1) give the following guarantee, including its additional initialization condition and conversion cost.

Proposition E.1 (Outer stationarity for Smoothed-AGDA)**.** *Suppose Assumptions 3.1–3.5 hold with $\kappa \geq 2$, and assume in addition that*

$$\Phi(x_0) - \inf_x \{f(x, y_0) + \ell \|x - x_0\|^2\} \leq c_{\mathrm{init}} \Delta. \tag{E.1}$$

For every $0 < \varepsilon^2 \leq \ell\Delta$, Smoothed-AGDA followed by the conversion below returns a point $\widehat{x}$ with $\mathbb{E}\|\nabla\Phi(\widehat{x})\|^2 \leq \varepsilon^2$, using at most

$$C\left[\frac{\kappa\ell\Delta}{\varepsilon^2} + \frac{\kappa^2\ell\Delta\sigma^2}{\varepsilon^4} + \left(\kappa + \frac{\kappa^5\sigma^2}{\varepsilon^2}\right)\log(2\kappa)\right] \tag{E.2}$$

stochastic partial-gradient calls, where C depends at most on c_{init}. In particular, if $\varepsilon^2 \leq c\ell\Delta/(\kappa^3\log(2\kappa))$ for a sufficiently small universal $c > 0$, this cost is

$$O\left(\frac{\kappa\ell\Delta}{\varepsilon^2} + \frac{\kappa^2\ell\Delta\sigma^2}{\varepsilon^4}\right). \tag{E.3}$$

Proof. We first obtain a pair with small squared game gradients. Its exact proximal point has a small outer gradient, and stochastic iterations approximate that point in mean square. Their step sizes and count depend on the accuracy and problem parameters, not on the realized pair.

The iterate pair from Smoothed-AGDA. Initialize the auxiliary center of Yang et al. (2022, Algorithm 2) at $z_0 = x_0$. Their initial quantity b_0 is twice the left side of (E.1), so $b_0 \leq 2c_{\mathrm{init}}\Delta$. Before tuning the primal step α, the proof of their Theorem 4.1, Eqs. (40)–(41), bounds the average squared game gradients by

$$\frac{1}{N}\sum_{t=0}^{N-1}\mathbb{E}\Big[\|\nabla_x f(x_t, y_t)\|^2 + \kappa\|\nabla_y f(x_t, y_t)\|^2\Big] \leq C\frac{\kappa\Delta}{N\alpha} + C\kappa\ell\alpha\sigma^2, \qquad 0 < \alpha \leq \frac{1}{3\ell}. \tag{E.4}$$

Here the dual step is $\alpha/48$, the regularization coefficient is 2ℓ, and the center relaxation is $\alpha\ell/(76800\kappa)$, as in their proof. Their descent estimates use only the conditional means and variances of the two gradient calls; no smoothness of individual sampled gradients is required. They therefore apply to the oracle formulation in Assumptions 3.4 and 3.5.

Choose $\alpha = \min\{1/(3\ell), \sqrt{\Delta/(N\ell\sigma^2)}/2\}$, omitting the second term when $\sigma = 0$. With $N \geq C(\kappa\ell\Delta/\varepsilon^2 + \kappa^2\ell\Delta\sigma^2/\varepsilon^4)$, a uniformly selected pair $(\widetilde{x}, \widetilde{y})$ satisfies $\mathbb{E}e \leq \varepsilon^2/16$, where $e := \|\nabla_x f(\widetilde{x}, \widetilde{y})\|^2 + \kappa\|\nabla_y f(\widetilde{x}, \widetilde{y})\|^2$. Using the upper budget Δ in this choice, rather than the actual initial outer gap, also covers initializations with zero outer gap.

The exact proximal point. Condition on the selected pair, and define the auxiliary game $\widehat{f}(x, y) := f(x, y) + \ell\|x - \widetilde{x}\|^2$. It is ℓ-strongly convex in x, retains inner PL with parameter ℓ/κ, and is 3ℓ-smooth in the partial-gradient convention of Assumption 3.1. Since $f(\cdot, y) + \ell\|\cdot\|^2/2$ is convex for every y, its supremum has the same property. Thus $\Phi(x) + \ell\|x - \widetilde{x}\|^2$ is ℓ-strongly convex and has a unique minimizer $p := \arg\min_x\{\Phi(x) + \ell\|x - \widetilde{x}\|^2\}$. Inner PL and strong convexity in x give

$$\begin{aligned}\Phi(\widetilde{x}) - \min_x \widehat{f}(x, \widetilde{y}) &= \Phi(\widetilde{x}) - f(\widetilde{x}, \widetilde{y}) + \widehat{f}(\widetilde{x}, \widetilde{y}) - \min_x \widehat{f}(x, \widetilde{y}) \leq \frac{e}{2\ell}, \\ \frac{\ell}{2}\|p - \widetilde{x}\|^2 &\leq \Phi(\widetilde{x}) - \Phi(p) - \ell\|p - \widetilde{x}\|^2 \leq \frac{e}{2\ell}.\end{aligned} \tag{E.5}$$

The second line uses strong convexity of the auxiliary outer objective, then $\Phi(p) + \ell\|p - \widetilde{x}\|^2 \geq \min_x \widehat{f}(x, \widetilde{y})$. Optimality of p now yields the pointwise bound

$$\nabla\Phi(p) = 2\ell(\widetilde{x} - p), \qquad \|\nabla\Phi(p)\|^2 \leq 4e. \tag{E.6}$$

This is the squared-gradient estimate underlying Yang et al. (2022, Proposition 2.1(b), Eq. (7)); no conversion from a first-moment norm bound is used.

A finite stochastic conversion. Starting at $(x^0, y^0) = (\widetilde{x}, \widetilde{y})$, apply the dual-first alternating method used in that proposition to $\widehat{f}$, with dual step τ and primal step 162τ. Explicitly, use fresh independent samples to update

$$\begin{aligned} y^{k+1} &= y^k + \tau G_y(x^k, y^k; \xi_k^y), \\ x^{k+1} &= x^k - 162\tau\big[G_x(x^k, y^{k+1}; \xi_k^x) + 2\ell(x^k - \widetilde{x})\big], \qquad 0 < \tau \le \frac{1}{486\ell}. \end{aligned} \tag{E.7}$$

These are an additional postprocessing procedure, not the primal-first recursion defining $\mathrm{SGDA}_{\mathrm{Alt}}$ in Section 3.

We use the potential contraction of Yang et al. (2022, Theorem A.1 and the proof of Proposition 2.1(b)), then obtain the required distance bound directly from strong convexity. Write $\Psi(y) := \min_x \widehat{f}(x, y)$ and $\Psi^* := \max_y \Psi(y)$. The auxiliary game has a saddle with outer coordinate p: any inner maximizer at p also minimizes $\widehat{f}(\cdot, y)$ there, by Lemma B.2 and strong convexity. Let

$$P_k := \mathbb{E}_{\mathrm{conv}}\left[\Psi^* - \Psi(y^k) + \frac{1}{10}\big(\widehat{f}(x^k, y^k) - \Psi(y^k)\big)\right],$$

where $\mathbb{E}_{\mathrm{conv}}$ averages the fresh conversion samples conditional on $(\widetilde{x}, \widetilde{y})$. The initialization estimate and their contraction give

$$P_0 \le \frac{11}{10}\left[\Phi(\widetilde{x}) - \min_x \widehat{f}(x, \widetilde{y})\right] \le \frac{11e}{20\ell}, \qquad P_k \le \left(1 - \frac{\ell\tau}{2\kappa}\right)^k P_0 + C\kappa\tau\sigma^2. \tag{E.8}$$

To bound distance to p, let $x_y := \arg\min_x \widehat{f}(x, y)$. Strong convexity at this minimizer and the saddle inequality $\widehat{f}(p, y) \le \Psi^*$ give

$$\begin{aligned} \|x - p\|^2 &\le 2\|x - x_y\|^2 + 2\|x_y - p\|^2 \\ &\le \frac{4}{\ell}\big[\widehat{f}(x, y) - \Psi(y) + \Psi^* - \Psi(y)\big]. \end{aligned} \tag{E.9}$$

Hence $\mathbb{E}_{\mathrm{conv}}\left\|x^k - p\right\|^2 \le 40P_k/\ell$. The step choices satisfy their conditions with smoothness 3ℓ and PL parameters ℓ and ℓ/κ. In particular, their variance term is $O(\kappa\tau\sigma^2)$. Substitution and expectation over the selected pair give

$$\mathbb{E}\left\|x^k - p\right\|^2 \le \frac{C}{\ell^2}\left(1 - \frac{\ell\tau}{2\kappa}\right)^k \mathbb{E}e + \frac{C\kappa\tau\sigma^2}{\ell}. \tag{E.10}$$

For $\sigma > 0$, choose $\tau = \min\{1/(486\ell), c\varepsilon^2/(\kappa^4\ell\sigma^2)\}$ with a sufficiently small universal $c > 0$; use $\tau = 1/(486\ell)$ when $\sigma = 0$. After

$$k = \left\lceil \frac{C\kappa}{\ell\tau}\log(2\kappa)\right\rceil = O\left(\left[\kappa + \frac{\kappa^5\sigma^2}{\varepsilon^2}\right]\log(2\kappa)\right)$$

updates, (E.10) is at most $\varepsilon^2/(16\kappa^2\ell^2)$. This choice works in expectation over the random starting pair, without requiring its realized game gradients to satisfy a deterministic accuracy bound. With $\widehat{x} = x^k$, Lemma B.2 and (E.6) conclude

$$\mathbb{E}\left\|\nabla\Phi(\widehat{x})\right\|^2 \le 2\mathbb{E}\left\|\nabla\Phi(p)\right\|^2 + 8\kappa^2\ell^2\mathbb{E}\left\|\widehat{x} - p\right\|^2 \le 8\mathbb{E}e + \frac{\varepsilon^2}{2} \le \varepsilon^2. \tag{E.11}$$

There are two partial-gradient calls per update in both stages. Adding their costs proves (E.2). Under the stated small-error condition, the term $\kappa\log(2\kappa)$ is absorbed by $\kappa\ell\Delta/\varepsilon^2$, and $\kappa^5\sigma^2\varepsilon^{-2}\log(2\kappa)$ is absorbed by $\kappa^2\ell\Delta\sigma^2\varepsilon^{-4}$, giving (E.3). □

E.1 Initialization budget of the hard instances

We verify the additional budget (E.1) for each hard instance used in Theorem 4.1. The ordinary outer and inner gap checks remain in their construction proofs; here we control the regularized initial gap using the same declared budget Δ and the same initializations.

The unstable quadratic. For f_{unstable} in (5.2), Appendix C.1 chooses $x_0^2 + y_0^2 = \Delta/(\ell+\mu)$, where $\mu = \ell/\kappa \le \ell/2$. Minimizing the scalar quadratic gives

$$\begin{aligned}\Phi_{\mathrm{unstable}}(x_0) - \inf_x\{f_{\mathrm{unstable}}(x,y_0) + \ell(x-x_0)^2\} &= \frac{\ell+\mu}{4}(y_0 - x_0/2)^2 + \frac{(\ell+\mu)^2 y_0^2}{64\ell} \\ &\le \frac{43\Delta}{128} \le c_{\mathrm{init}}\Delta.\end{aligned}$$

The bound uses $(y_0 - x_0/2)^2 \le 5\Delta/[4(\ell+\mu)]$, $y_0^2 \le \Delta/(\ell+\mu)$, and $c_{\mathrm{init}} \ge 1$. This verification includes the one-iterate case $T = 1$, for which the high-ratio proof uses this quadratic.

The corridor and noise components. The game f_{stable} in (5.4) is initialized at $(u_0, v_0; y_0) = (0,0;0)$. Equation (C.42) and the zero-component convention give $g_{\sigma,j}(v[j];0) = 0$ for every component, so $f_{\mathrm{stable}}(u,v;0) = g_{\mathrm{corridor}}(u)$. Also $\Phi_{\mathrm{stable}}(0,0) = g_{\mathrm{corridor}}(0)$, and (C.25) gives $g_{\mathrm{corridor}}(0) - \inf_u g_{\mathrm{corridor}}(u) = \Delta$. Therefore,

$$\begin{aligned}&\Phi_{\mathrm{stable}}(0,0) - \inf_{u,v}\{f_{\mathrm{stable}}(u,v;0) + \ell(u^2 + \|v\|^2)\} \\ &\quad = g_{\mathrm{corridor}}(0) - \inf_u\{g_{\mathrm{corridor}}(u) + \ell u^2\} \le \Delta \le c_{\mathrm{init}}\Delta.\end{aligned}$$

This also verifies the corridor alone, including $\sigma = 0$ and schedules with no positive noise-component steps, when all noise components vanish. Each noise component alone has regularized initial gap zero.

The low-ratio game. Retain the constants $C, C_{\mathrm{low}}, \zeta$ and scales A, s_x, s_y from Appendix A.3, with f_{cycle} defined by (A.3) and initialized at $x_0 = s_x e_1$, $y_0^{\mathrm{cycle}} = s_y y_0$. We first justify the bound on the unscaled initial outer gradient used in choosing C. Affinity in u and $\tilde{g}_{\mathrm{cycle}}(u, y_0) \le \|u\|^2/2$ imply, by completing the square,

$$\frac{1}{2}\|\nabla_u \tilde{g}_{\mathrm{cycle}}(e_1, y_0) - e_1\|^2 \le \frac{1}{2} - \tilde{g}_{\mathrm{cycle}}(e_1, y_0) \le C.$$

Thus $\|\nabla_u \tilde{g}_{\mathrm{cycle}}(e_1, y_0)\| \le 1 + \sqrt{2C}$ is universally bounded. The constant C in Appendix A.3 is chosen large enough to bound this gradient as well, before C_{low} is fixed. Affinity in x, the gap bound (A.7), and the scales in (A.2) now give

$$\begin{aligned}&\Phi_{\mathrm{cycle}}(x_0) - \inf_x\{f_{\mathrm{cycle}}(x, y_0^{\mathrm{cycle}}) + \ell\|x - x_0\|^2\} \\ &\quad = \Phi_{\mathrm{cycle}}(x_0) - f_{\mathrm{cycle}}(x_0, y_0^{\mathrm{cycle}}) + \frac{\left\|\nabla_x f_{\mathrm{cycle}}(x_0, y_0^{\mathrm{cycle}})\right\|^2}{4\ell} \\ &\quad \le \left(2C + \frac{C^2}{2C_{\mathrm{low}}}\right)\zeta\Delta.\end{aligned}$$

Reducing the universal ζ makes this at most $c_{\mathrm{init}}\Delta$ without affecting c_{low}. The check is uniform over the horizon and schedule, including $T = 1$. Both this game and the unstable quadratic use exact gradients, so their checks apply for every $\sigma \ge 0$.

Scope of the separation. The verification above establishes (E.1) for all hard instances, so the comparison holds on the same class and uses the declared gap budget Δ. SGDA returns a uniformly selected iterate, whereas $\widehat{x}$ is computed by additional iterations. This does not rule out improving SGDA through arbitrary postprocessing.

F Comparison conditions for Table 1

Table 1 fixes initial-gap budgets and displays the dependence on ℓ, κ, σ, and ε. The rows use different feedback models and stationarity criteria. Expected-norm guarantees use tolerance ε, whereas expected-squared-norm guarantees use ε^2; an expected-norm upper bound does not by itself imply an expected-squared-norm guarantee.

SGDA and Smoothed-AGDA. The matching SGDA bounds use the average expected squared outer-gradient criterion in (3.7) and optimize over fixed timescale ratios and predetermined nonincreasing step sizes. Proposition E.1 gives the Smoothed-AGDA comparison under the additional initialization condition (E.1), which our hard instances satisfy. Its full complexity (E.2) includes the conversion to an outer-stationary point. The table displays the leading terms in (E.3), valid for $\varepsilon^2 \le c\ell\Delta/(\kappa^3 \log(2\kappa))$. SGDA returns a uniformly selected iterate; Smoothed-AGDA uses the point computed by this conversion.

Sampled-gradient momentum methods. The MSGDA and AdaMSGDA rows show only the leading stochastic contributions for $0 < \varepsilon \le \sigma$ under Lipschitz sampled gradients (Huang et al., 2025). These methods evaluate the same sampled function at different iterates, using more information than fresh-sample SGDA. AdaMSGDA controls a gradient mapping, with its additional spectral constants held fixed in the table. The displayed noise terms omit deterministic and lower-order contributions.

Hessian-corrected momentum. HCMM-2 controls the average expected outer-gradient norm and uses stochastic Hessian-vector products (Cai et al., 2026). Its assumptions include Lipschitz population Hessians. The table displays only its leading noise contribution, with Hessian variance at most σ^2; O_ℓ leaves ℓ-dependence implicit. Its full guarantee also contains deterministic, initialization, and Hessian-regularity terms.

Finite-sum upper bounds. SPIDER-GDA and SGDA-RR assume smooth components, with one component-gradient evaluation as the unit cost. SPIDER-GDA controls the expected outer-gradient norm (Chen et al., 2022). The SGDA-RR row specializes the component-variance condition to a uniform bound σ^2 and controls the average expected squared outer gradient at epoch starts (Cho and Yun, 2023). Component reuse and sampling without replacement differ from the fresh-sample oracle used by SGDA.

Lower bounds and the nonconvergence entry. The finite-sum linear-span lower bound assumes $\kappa \ge 2n$ and averaged component smoothness, which is weaker than smoothness of every component (Zhang et al., 2021). The stochastic NC–SC lower bound uses a projected outer-gradient mapping under its domain and initialization assumptions (Li et al., 2021). These oracle lower bounds concern broader algorithm classes than the SGDA recursion studied here. The $\Omega(1)$ entry records a nonvanishing stationarity error, rather than an iteration complexity: Theorem 6.3 gives an error floor proportional to $\ell\Delta$ at every prescribed horizon.