ALTGDA ACHIEVES GLOBAL $O(1/T)$ ERGODIC CONVERGENCE IN MATRIX GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

Alternating gradient descent–ascent (AltGDA) is a simple and practically effective method for solving finite two-player zero-sum matrix games. However, the theory of AltGDA remains limited: existing results either apply only to unconstrained settings or require restrictive assumptions on the equilibrium in constrained settings. We show that AltGDA converges globally at an $O(1/T)$ ergodic rate in every finite two-player zero-sum matrix game. Unlike prior results, our guarantee holds for every initialization and every horizon T: the uniform averages of the AltGDA iterates satisfy an $O(1/T)$ duality-gap bound. Our proof is inspired by numerical results obtained using a novel performance estimation programming (PEP) framework for Lyapunov function search over compact convex sets. Additionally, we provide simple counterexamples showing that the last-iterate duality gap of AltGDA does not converge to zero. This justifies why averaging of iterates is indeed necessary to achieve an $O(1/T)$ rate. We have formalized and machine-checked our global ergodic convergence result in Lean 4.

1 INTRODUCTION

Game-theoretic learning algorithms underpin systems that attain superhuman play in poker (Bowling et al., 2015; Moravčík et al., 2017; Brown & Sandholm, 2018; 2019) as well as expert-level play in Stratego (Perolat et al., 2022), and Diplomacy (Meta Fundamental AI Research Diplomacy Team (FAIR) et al., 2022). A canonical model for learning in games is the two-player zero-sum matrix game, where two players select their strategies probabilistically to optimize an expected payoff, with one player minimizing and the other maximizing.

The simplest learning dynamics is simultaneous gradient descent-ascent (SimGDA), where both players run projected gradient steps on their own strategies. SimGDA is known to achieve an $O(1/\sqrt{T})$ ergodic convergence rate in the duality gap. In this paper, we study the *alternating* variant of gradient descent-ascent (AltGDA), where the players take turns updating their strategies, while always observing the newly-updated strategy of the opposing player. Empirically, AltGDA exhibits much better performance in solving finite two-player zero-sum matrix games compared to SimGDA, yet our theoretical understanding of AltGDA remains limited. Our question is whether the AltGDA dynamics, without any optimism or other tricks, admit a global $O(1/T)$ ergodic guarantee.

The question is delicate because plain gradient play in adversarial problems need not settle. In unconstrained bilinear games, simultaneous gradient dynamics diverges away from the equilibrium, and the alternating variant cycles (Bailey et al., 2020); on the simplex, adversarial regularized learning cycles as well (Mertikopoulos et al., 2018b). The classical remedies modify the dynamics through extragradient steps (Korpelevich, 1976) or optimistic corrections (Popov, 1980; Rakhlin & Sridharan, 2013a). The stronger last-iterate guarantees known today are for those modified dynamics (Wei et al., 2021).

Alternation itself has a long practical record. It entered large-scale game solving as a numerical device in Counterfactual Regret Minimization+ (CFR+) (Tammelin et al., 2015), the algorithm used for superhuman poker AI. Alternation provably does not hurt the theoretical guarantees (Burch et al., 2019; Farina et al., 2019). On random matrix games, numerical studies in (Kroer, 2020; Nan et al., 2026) suggest an ergodic $O(1/T)$ convergence rate for AltGDA. On the theory side, (Wibisono et al., 2022; Katona et al., 2026) show that alternation provably helps for alternating *mirror* descent

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.

with Legendre regularizers. Legendre regularizers guarantee that the dynamics never touch the boundaries of the simplices. In contrast AltGDA's Euclidean projections cause collisions with the simplex boundary, which is the main difficulty in a theoretical analysis. Nan et al. (2026) sidestep this boundary issue by assuming that the game contains an interior equilibrium, or that the dynamics are initialized close to an equilibrium, and show $O(1/T)$ convergence for these special cases.

Contributions. Our contributions in this paper are threefold:

- **A global $O(1/T)$ ergodic convergence result for AltGDA (Section 4, Appendix A).** We show that for every finite matrix game, every initialization, and every iteration budget T, the uniform averages $(\overline{x}^T, \overline{y}^T)$ of the AltGDA iterates satisfy $\mathrm{Gap}(\overline{x}^T, \overline{y}^T) \leq 15/(2\eta T)$ for every fixed stepsize $0 < \eta \leq \delta/(2\sqrt{2}\,L)$, where $L > 0$ bounds the spectral norm of the payoff matrix and $\delta \in (0, 1]$ is a positive constant associated with a Goldman–Tucker saddle point. Our result provides the first global guarantee without any restrictions, showing that alternation alone suffices for an $O(1/T)$ ergodic rate in general matrix games. We have formalized and machine-checked our convergence results in Lean 4 (de Moura & Ullrich, 2021).
- **A computer-assisted framework for Lyapunov-function discovery in constrained bilinear games (Section 5, Appendix B).** We present a computer-assisted framework based on performance estimation programming (PEP) that transforms the search for quadratic Lyapunov certificates for AltGDA into a linear matrix inequality problem (LMI). The framework combines computer-assisted automated Lyapunov analysis (Upadhyaya et al., 2025a;b) with operator interpolation (Bousselmi et al., 2024) to construct an LMI over general compact convex bilinear problems, such that exact feasibility of this LMI certifies a quadratic Lyapunov certificate leading to an $O(1/T)$ ergodic convergence rate. We demonstrate numerically that this LMI is feasible subject to the tolerance of the SDP solver and the patterns present in computed numerical solutions guided us in deriving the analytical convergence results in the first contribution.
- **The last-iterate duality gap of AltGDA does not converge to zero (Section 6, Appendix C).** Our first two contributions involve the ergodic duality gap, so a natural next question is if we can achieve an $O(1/T)$ convergence rate for AltGDA in last-iterate duality gap. We provide simple counterexamples showing that the last-iterate duality gap of AltGDA does not converge to zero both in matrix games and in bilinear problems over general compact convex sets. First, for compact convex sets, we show a simple 2-dimensional example over the unit ball, where the last-iterate duality gap is exactly two for every iterate and every stepsize choice. Second, for the special case of matrix games, we present a 2×2 matrix such that the last-iterate duality gap is bounded away from zero no matter how its stepsize is tuned. Our explicit counterexamples are motivated again by a PEP-based framework that numerically computes the worst-case last-iterate duality gap of AltGDA over compact convex sets.

Julia and Lean code. The Lean code for verifying our global ergodic convergence result and Julia code for reproducing the numerical PEP results are available in the attached zip file accompanying the paper. Unless otherwise specified, we conducted the experiments on a laptop running macOS Tahoe with an Apple M5 Max processor and 128 GB of memory. The Lean verification uses Lean 4.32.0 and mathlib 4.32.0.

2 Related Work

Alternation in unconstrained minimax optimization. In unconstrained bilinear games, Bailey et al. (2020) showed that AltGDA with fixed stepsizes achieves constant alternating regret against any fixed comparator, yielding $O(1/T)$ average regret. For smooth strongly convex-strongly concave objectives, Zhang et al. (2022) established near-optimal local linear convergence with better condition-number dependence than simultaneous updates. Lee et al. (2024) subsequently established a global advantage by proving an iteration-complexity upper bound for AltGDA below the corresponding lower bound for simultaneous GDA. Feng et al. (2025) showed that smaller heavy-ball momentum improves local stability in continuous-time models.

Alternation on constrained strategy sets. For constrained bilinear games, Wibisono et al. (2022) established an $O(T^{-2/3})$ ergodic rate for alternating mirror descent, which Katona et al. (2026) improved to $O(T^{-4/5})$ through symplectic analysis. These guarantees rely on Legendre regularizers and smoothness assumptions that do not cover Euclidean projection onto simplices. For projected AltGDA in matrix games, Nan et al. (2026) proved an ergodic $O(1/T)$ rate for sufficiently small constant stepsizes, globally when an interior equilibrium exists and locally without this condition. Our theorem removes the interiority requirement, extending the global guarantee to every finite matrix game and every initialization. In online learning, Cevher et al. (2023) constructed algorithms attaining $\widetilde{O}(T^{1/3})$ alternating regret for linear losses on simplices. Hait et al. (2025) extended these guarantees to bounded convex losses using continuous Hedge, with $\widetilde{O}(d^{2/3}T^{1/3})$ alternating regret in dimension d. The benefits of alternation depend on the learning rule: Lazarsfeld et al. (2025) showed that, in fictitious play, optimism can achieve constant regret without regularization, whereas alternation alone need not.

Optimistic first-order methods in game solving. Usually, optimistic-type of modifications such as extrapolation and prediction are used in first-order methods for saddle-point optimization to achieve fast convergence rate. Classical examples include the extragradient method (Korpelevich, 1976), Mirror-Prox (Nemirovski, 2004), dual extrapolation (Nesterov, 2007), and primal–dual splitting (Chambolle & Pock, 2011). The latter three achieve $O(1/T)$ ergodic convergence for various convex-concave problem classes. In online learning, optimistic mirror descent and optimistic follow-the-regularized-leader exploit predictability in the observed losses to obtain faster convergence in games (Rakhlin & Sridharan, 2013b; Syrgkanis et al., 2015). Mokhtari et al. (2020) provide a unified analysis of extragradient and optimistic gradient descent–ascent (OGDA) as approximations of the proximal-point method. A complementary line of work studies the last-iterate behavior of optimistic methods (Daskalakis et al., 2017; Mertikopoulos et al., 2018a; Daskalakis & Panageas, 2018). In particular, Wei et al. (2020) establish linear last-iterate convergence of OGDA in bilinear games over polytopes without assuming equilibrium uniqueness, and of optimistic multiplicative weights in matrix games with a unique equilibrium. Finite-time last-iterate guarantees for extragradient and OGDA also extend to constrained monotone games (Cai et al., 2022). Our result shows that the $O(1/T)$ ergodic rate can also be attained through *alternation alone*: the uniform average of the unmodified AltGDA iterates converges at the $O(1/T)$ rate in every finite matrix game, without additional correction, extrapolation, or prediction terms.

Performance estimation programming and Lyapunov analysis. The PEP framework was introduced by Drori & Teboulle (2014) and developed through exact interpolation formulations by Taylor et al. (2017a;b). Subsequently, Bousselmi et al. (2024) derived interpolation conditions for linear operators, while Upadhyaya et al. (2025a;b) developed automated Lyapunov analysis and corresponding software AutoLyap. For AltGDA, Nan et al. (2026) used PEP to optimize stepsizes against worst-case ergodic bounds. In contrast, we combine operator interpolation and automated Lyapunov analysis to search for certificates, with the numerical results interpreted through approximate feasibility in Section 5. We also study finite-horizon last-iterate performance over compact convex sets in Section 6.

3 PRELIMINARIES

Bilinear saddle-point problems. In this paper we consider the following bilinear saddle point problems (SPP) of the form:

$$\min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} y^\top A x, \tag{SPP}$$

where $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ are nonempty compact convex sets and $A \in \mathbb{R}^{m\times n}$. Here $\mathbb{R}^d$ denotes the d-dimensional Euclidean space. We say that a feasible pair $(x^\star, y^\star) \in \mathcal{X} \times \mathcal{Y}$ is a *saddle point* if

$$y^\top A x^\star \le (y^\star)^\top A x^\star \le (y^\star)^\top A x \text{ for all } x \in \mathcal{X},\ y \in \mathcal{Y}.$$

Compactness, convexity, and bilinearity ensure that a saddle point exists and that the two optimization orders have the common value

$$\nu = \min_{x\in\mathcal{X}} \max_{y\in\mathcal{Y}} y^\top A x = \max_{y\in\mathcal{Y}} \min_{x\in\mathcal{X}} y^\top A x.$$

For a feasible pair $(\widetilde{x}, \widetilde{y}) \in \mathcal{X} \times \mathcal{Y}$, define its *duality gap* by

$$\mathrm{Gap}(\widetilde{x}, \widetilde{y}) := \max_{y \in \mathcal{Y}} y^\top A\widetilde{x} - \min_{x \in \mathcal{X}} \widetilde{y}^\top Ax = \max_{x \in \mathcal{X},\ y \in \mathcal{Y}} \left(y^\top A\widetilde{x} - \widetilde{y}^\top Ax\right). \tag{1}$$

This gap is always nonnegative on $\mathcal{X} \times \mathcal{Y}$ and vanishes exactly at saddle points.

Matrix games. Finite two-player zero-sum matrix games are the specialization of (SPP) with $\mathcal{X} = \Delta_n$, $\mathcal{Y} = \Delta_m$, where $\Delta_d = \{z \in \mathbb{R}^d : z \geq 0,\ \mathbf{1}^\top z = 1\}$ is the probability simplex. The x-player mixes over n actions to minimize the expected payoff, and the y-player mixes over m actions to maximize it. A saddle point is then a *Nash equilibrium* (NE); every finite matrix game has at least one (von Neumann, 1928).

Alternating GDA. Let Π_C denote Euclidean projection onto a nonempty closed convex set C. From an initialization $(x^0, y^0) \in \mathcal{X} \times \mathcal{Y}$, alternating gradient descent–ascent (AltGDA) with a fixed stepsize $\eta > 0$ generates the iterates

$$\begin{aligned} x^{t+1} &= \Pi_{\mathcal{X}}\left(x^t - \eta A^\top y^t\right), \\ y^{t+1} &= \Pi_{\mathcal{Y}}\left(y^t + \eta A x^{t+1}\right), \end{aligned} \tag{AltGDA}$$

for $t = 0, 1, \ldots, T-1$, where $T \geq 1$ is the iteration budget. The x-update is a projected descent step; the y-update is a projected ascent step using the newly computed x^{t+1}. At the time horizon T, we distinguish the last iterate (x^T, y^T) from the uniform averages $\overline{x}^T = \frac{1}{T}\sum_{t=1}^T x^t$ and $\overline{y}^T = \frac{1}{T}\sum_{t=1}^T y^t$.

For the remainder of this section and the global ergodic convergence analysis, we specialize to $\mathcal{X} = \Delta_n$ and $\mathcal{Y} = \Delta_m$ associated with *AltGDA for matrix games*.

Goldman-Tucker strictly complementary saddle point. A saddle point is *strictly complementary* if, at each coordinate, either the equilibrium strategy has a strictly positive probability or the corresponding payoff slack is strictly positive. For a saddle point $(x^\star, y^\star)$, we define its payoff slack vectors as

$$s_x = A^\top y^\star - \nu\mathbf{1} \geq 0, \quad s_y = \nu\mathbf{1} - Ax^\star \geq 0. \tag{2}$$

Then, $(x^\star, y^\star)$ is called a strictly complementary saddle point if

$$x_i^\star s_{x,i} = 0, \quad x_i^\star + s_{x,i} > 0, \quad y_j^\star s_{y,j} = 0, \quad y_j^\star + s_{y,j} > 0, \tag{3}$$

for $i = 1, 2, \ldots, n$ and $j = 1, 2, \ldots, m$. By the Goldman-Tucker theorem (Goldman & Tucker, 1956; Williams, 1970), every matrix game admits a strictly complementary saddle point. We fix one such pair and call it the *Goldman-Tucker strictly complementary saddle point*.

Let $L > 0$ be an upper bound of the spectral norm of A: $\|A\|_2 \leq L$. Let $(x^\star, y^\star)$ be the Goldman-Tucker strictly complementary saddle point we choose, and its payoff slack vectors defined as in Eq. (2). Define the supports: $I := \mathrm{supp}(x^\star)$, $J := \mathrm{supp}(y^\star)$, and the smallest on-support probabilities: $\underline{x} := \min_{i \in I} x_i^\star > 0$ and $\underline{y} := \min_{j \in J} y_j^\star > 0$. When the complements of the supports are nonempty, we denote the smallest off-support slacks by $\sigma_x := \min_{i \in I^c} s_{x,i}$ and $\sigma_y := \min_{j \in J^c} s_{y,j}$. The Goldman-Tucker theorem guarantees that $s_{x,i} > 0$ for all $i \in I^c$ and $s_{y,j} > 0$ for all $j \in J^c$. We then define the following scalar parameter:

$$\delta := \min\left\{\underline{x},\ \underline{y},\ \frac{\sigma_x}{L} \text{ if } I^c \neq \varnothing,\ \frac{\sigma_y}{L} \text{ if } J^c \neq \varnothing\right\}, \tag{4}$$

where terms corresponding to empty complements are omitted. It follows from the above that $0 < \delta \leq 1$. We call δ the *Goldman-Tucker separation constant* associated with the selected saddle point and the chosen bound L. Intuitively, this constant records how decisively the Goldman-Tucker saddle point separates its supports from their complements, so that the larger the constant, the more decisive the separation (since all on-support coordinates are far from zero and all off-support coordinates have large payoff slacks).

4 Main Theorem and its Proof

4.1 Main theorem and discussion

We establish the $O(1/T)$ ergodic convergence rate in the following theorem.

Theorem 1. *For any matrix game with the payoff matrix $A \in \mathbb{R}^{m\times n}$, let $L > 0$ satisfy $\|A\|_2 \le L$, and define δ from a selected Goldman-Tucker strictly complementary saddle point as in Eq.* (4)*. For every initialization $(x^0, y^0) \in \Delta_n \times \Delta_m$, every $T \ge 1$, and every fixed stepsize satisfying $0 < \eta \le \frac{\delta}{2\sqrt{2}\,L}$, the averages of the iterates generated by* (AltGDA) *for matrix games with $\mathcal{X} = \Delta_n$ and $\mathcal{Y} = \Delta_m$ satisfy*

$$\mathrm{Gap}\left(\overline{x}^T, \overline{y}^T\right) \le \frac{15}{2\eta T}. \tag{5}$$

We present a proof sketch of Theorem 1 in the next subsection, and defer the complete proof to Appendix A.

Theorem 1 establishes the first global $O(1/T)$ ergodic convergence guarantee for AltGDA in *every* finite matrix game. Therefore, it closes the gap left by the previous theory of Nan et al. (2026): *alternation alone, without optimism or any additional correction step, is sufficient to achieve an $O(1/T)$ ergodic rate in general matrix games*. More specifically, our result strengthens the previous guarantees in the following respects:

- *No structural assumption on the game.* Theorem 1 applies to every finite matrix game, and does not require the existence of an interior Nash equilibrium. Note that every matrix game admits a Goldman–Tucker strictly complementary saddle point due to the Goldman-Tucker theorem (Goldman & Tucker, 1956; Williams, 1970). In contrast, the previous global $O(1/T)$ guarantee required the game to admit an interior Nash equilibrium, which is a restrictive assumption as many matrix games have equilibria supported only on proper faces of the simplices and admit no fully mixed equilibrium.
- *No locality assumption on the initialization.* Theorem 1 holds for every pair of initial strategies. For games without an interior equilibrium, the previous $O(1/T)$ theorem required the initial strategies to lie in a sufficiently small neighborhood of a maximum-support equilibrium, which is a restrictive condition.
- *A uniform theorem for every horizon.* For general matrix games, the previous evidence for a global $O(1/T)$ rate in terms of the ergodic duality gap came from numerical finite-horizon PEP computations for $T = 5, 6, \ldots, 50$ in Nan et al. (2026), which conjectured an $O(1/T)$ ergodic-rate. Our Theorem 1 shows that the finite-horizon numerical conjecture in Nan et al. (2026) is indeed correct with an analytical proof that holds for every $T \ge 1$.

4.2 Proof Sketch of Theorem 1

The standard path to an $O(1/T)$ ergodic rate often follow a Lyapunov function based argument: we show a nonnegative Lyapunov or energy function whose per-step decrease dominates the quantity to be summed. However, the presence of alternation and projection in AltGDA requires some careful modifications to the Lyapunov based argument. We briefly describe the proof sketch of our main convergence result in five steps below with the details of each step shown in full in Appendix A.

Establish an exact energy balance. By the KKT conditions (AltGDA) for matrix games AltGDA admits the following equivalent representation without requiring explicit projection steps: there are scalars γ_t, λ_t and vectors $\mu_t \in \mathbb{R}^n_+, \rho_t \in \mathbb{R}^m_+$ such that (Lemma A.1)

$$\Delta x_t := x^{t+1} - x^t = \eta(v_t - \gamma_t \mathbf{1} + \mu_t), \quad \mu_t \odot x^{t+1} = 0, \tag{6}$$

$$\Delta y_t := y^{t+1} - y^t = \eta(u_t - \lambda_t \mathbf{1} + \rho_t), \quad \rho_t \odot y^{t+1} = 0, \tag{7}$$

where $v_t := -A^\top y^t$ and $u_t := Ax^{t+1}$, and $\odot$ denotes the coordinatewise product. We call γ_t, λ_t *threshold scalars*, and μ_t, ρ_t are *multiplier vectors*.

Let us fix a Goldman-Tucker saddle point $(x^\star, y^\star)$ as the selected NE. We can then define several important quantities that are useful throughout the proof.

We leverage an *energy function* V_t that consists of the squared Euclidean distance between the current iterates and the selected NE, plus a bilinear cross term to account for the alternation:

$$V_t := \|x^t - x^\star\|_2^2 + \|y^t - y^\star\|_2^2 - \eta(y^t - y^\star)^\top A(x^t - x^\star). \tag{8}$$

We then define the *dissipation* term $\mathscr{D}_t$ that scales with the magnitude of the mismatches between the selected NE and the current iterates on both on-support and off-support coordinates.

$$\mathscr{D}_t := \eta\left(s_x^\top(x^t + x^{t+1}) + s_y^\top(y^t + y^{t+1})\right) + 2\eta\left(\mu_t^\top x^\star + \rho_t^\top y^\star\right) \tag{9}$$

where the first term measures the probability masses of the current iterates on the off-support coordinates, and the second term measures the current iterates colliding with the simplex boundary on the on-support coordinates. We notice that both terms vanish when the supports of the selected NE are correctly identified.

Finally, we define the *residual* term $r_t := \eta v_t^\top \Delta x_t + \eta u_t^\top \Delta y_t - \|\Delta x_t\|_2^2 - \|\Delta y_t\|_2^2$. By using the KKT representations, we can show that the residual term is exactly (Lemma A.2)

$$r_t = \eta(\mu_t^\top x^t + \rho_t^\top y^t) \geq 0. \tag{10}$$

This term is positive exactly when there exists at least one coordinate that has both positive probability mass in the previous iteration and a positive projection multiplier, and thus zero probability mass in the new iteration, by Eqs. (6) and (7). Therefore, this term is positive only when the trajectories of (AltGDA) collide with the simplex boundary.

As the first step of the proof, we establish an exact identity that balances V_t, $\mathscr{D}_t$, and r_t (Lemma A.3)

$$V_{t+1} - V_t + \mathscr{D}_t = r_t. \tag{11}$$

This equation explicitly describes the change of the energy function per step, and plays a central role in our proof.

Control the residual by the dissipation. The next central step we need is showing that the residual term is small relative to the dissipation term. In particular, we pursue an upper bound for the residual term r_t which is smaller than $\mathscr{D}_t$ when we choose a small stepsize.

To achieve this, we need to handle the on-support and off-support coordinates separately. The difficult part is the off-support block of coordinates. It is turns out that we can prove the following useful decomposition of the off-support block. For the x-block, letting $(v)_{I^c} \in \mathbb{R}^{|I^c|}$ denote the restriction of a vector v to coordinates in I^c, we have (Lemma A.5)

$$\begin{aligned}\eta(\mu_t)_{I^c}^\top (x^t)_{I^c} = &-\langle G_x(\Delta x_t)_{I^c}, (\Delta x_t)_{I^c}\rangle \\ &- \eta(d_t^x)^\top(\Delta x_t)_{I^c} - \eta\left(\tfrac{1}{|I|}\sum\nolimits_{i\in I}\mu_{t,i}\right)\left(\sum\nolimits_{i\in I^c}(\Delta x_t)_i\right),\end{aligned} \tag{12}$$

where G_x is a positive semidefinite matrix and $d_t^x \in \mathbb{R}^{|I^c|}$ encodes the off-support payoff compared to the average of the on-support payoff. Here, we present this decomposition using the same notations as before, and we write a complete statement and its proof with a set of more comprehensive notations in Appendix (A.5.1).

For the on-support part, we have the following decomposition:

$$\eta(\mu_t)_I^\top x_I^t = -\eta(\mu_t)_I^\top\left((\Delta x_t)_I - (\tfrac{1}{|I|}\sum\nolimits_{i\in I}(\Delta x_t)_i)\mathbf{1}_{|I|}\right) + \eta\left(\tfrac{1}{|I|}\sum\nolimits_{i\in I}\mu_{t,i}\right)\left(\sum\nolimits_{i\in I^c}(\Delta x_t)_i\right).$$

Combining both on-support and off-support components, the following upper bound holds for the x-block:

$$\eta\mu_t^\top x^t \leq B_{x,t} - B_{x,t+1} + \sqrt{1+\tfrac{1}{|I|}}\,\tfrac{\eta^2L^2}{\sigma_x}(s_x^\top(x^t + x^{t+1})) + \sqrt{1-\tfrac{1}{|I|}}\,\tfrac{\eta^2 L}{\underline{x}}(\mu_t^\top x^\star), \tag{13}$$

where $B_{x,t} \coloneqq \eta(d_{t-1}^x)^\top(x^t)_{I^c}$, with $d_{-1}^x = d_0^x$, as in Eq. (A.34). If $I^c = \varnothing$, set $B_{x,t} = 0$ and omit the σ_x-term in Eq. (13). Eq. (13) bounds the x-block contribution $\eta\mu_t^\top x^t$ to r_t by the consecutive difference of a bounded storage term plus $O(\eta)\mathscr{D}_t$, where the implied constant depends only on the game and the selected saddle point. Combining this bound with the corresponding y-block bound (Eq. (A.48)), and setting $B_t := B_{x,t} + B_{y,t}$ with $B_{y,t}$ as in Eq. (A.44), the stepsize condition in Theorem 1 yields

$$r_t \leq B_t - B_{t+1} + \tfrac{1}{2}\mathscr{D}_t. \tag{14}$$

Prove the residual is summable. The summability of the residual term follows directly from Eqs. (11) and (14):

$$\sum_{t=0}^{T-1} r_t \le B_0 - B_T + \frac{1}{2}\sum_{t=0}^{T-1} \mathscr{D}_t = B_0 - B_T + \frac{1}{2}\sum_{t=0}^{T-1}(r_t - V_{t+1} + V_t),$$

which leads to a constant upper bound for $\sum_{t=0}^{T-1} r_t$ for any $T \ge 1$ because both V_t and B_t are bounded. The rest of the proof follows the same path as that in Nan et al. (2026, Lemma 1), which essentially proves the following descent inequality:

$$\eta(y^\top A x^t - (y^t)^\top A x) + \eta(y^\top A x^{t+1} - (y^{t+1})^\top A x) \le \Phi_t - \Phi_{t+1} + r_t, \quad \forall\, t \ge 1, \tag{15}$$

where Φ_t is a bounded Lyapunov function. Since the left-hand side of Eq. (15) is connected to the duality gap of the averaged iterates by averaging and maximizing, the summability of $\sum_{t=0}^{T-1} r_t$ is sufficient to prove an $O(1/T)$ convergence rate.

5 A PEP Framework for Lyapunov Certificates

Our Lyapunov-style analysis behind Theorem 1 was motivated by a PEP-based framework that numerically searches for quadratic Lyapunov certificates for AltGDA over general compact convex sets. Our framework is inspired by the recent automated PEP-based Lyapunov analysis of Upadhyaya et al. (2025a;b) and the operator interpolation results of Bousselmi et al. (2024).

The core idea behind our PEP-based Lyapunov framework is modeling and solving a linear matrix inequality (LMI) over the space of appropriate quadratic Lyapunov functions, and if this LMI is feasible, then there exists a suitable Lyapunov function that can lead to the desired ergodic convergence rate for AltGDA. In our Lyapunov function search three constraints (denoted by C_1, C_2, C_3) are modeled in the LMI: constraint C_1 requires one-step Lyapunov descent up to an auxiliary quadratic term, C_2 requires a nonnegative potential, and C_3 compares the auxiliary term with the one-step comparator gap, with the decision variables being two symmetric Lyapunov matrices, an auxiliary quadratic pair, and condition-specific positive semidefinite (PSD) blocks with interpolation and normal-cone multipliers. Also rather than considering only one iteration of AltGDA, we consider a variable number of iterations modeled by a history parameter h, as proving the desirable convergence rate may require more than one iteration. Here $h = 0$ corresponds to one iteration of AltGDA, $h = 1$ means two iterations and so on.

An *exactly* feasible solution to the Lyapunov LMI satisfies C_1, C_2, C_3 on every window, and telescoping gives the resultant $O(1/T)$ ergodic convergence rate for AltGDA in Theorem 2 below. However, the standard SDP solvers can only compute solutions with a numerical tolerance (e.g., Mosek (ApS, 2026) has a default feasibility tolerance of 10^{-8}). Such a numerically computed feasible solution is only *approximate* and *not exact*. However, the structure present in the approximately feasible solution to this PEP-based LMI often reveals specific patterns (Goujaud et al., 2023), which we can then exploit to arrive at an analytical Lyapunov function for AltGDA leading to the desired $O(1/T)$ ergodic convergence rate. In Appendix B, we describe the formulation for the LMI in detail and also discuss how the feasibility of this LMI leads to an $O(1/T)$ convergence rate for AltGDA. For brevity we just present the core result below.

Theorem 2 (Conditional Lyapunov–LMI rate)**.** *Let $A \in \mathbb{R}^{m\times n}$ satisfy $\|A\|_2 \le 1$, and let $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ be compact convex sets with $\|x\|_2 \le D$ for $x \in \mathcal{X}$ and $\|y\|_2 \le D$ for $y \in \mathcal{Y}$. Let $\{(x^t, y^t)\}_{t\ge 0}$ be generated by the AltGDA iteration* (AltGDA) *with initialization $(x^0, y^0) \in \mathcal{X}\times\mathcal{Y}$ and stepsize $\eta > 0$, and let* Gap *be defined by* (1)*. Fix an integer $h \ge 0$. Suppose the Lyapunov LMI of Appendix B is exactly feasible for (h, η), and fix any feasible certificate. Let Q^x and Q^y be the Lyapunov matrices in that certificate, and define $\varsigma_x^+ = \max\{\lambda_{\max}(Q^x), 0\}$ and $\varsigma_y^+ = \max\{\lambda_{\max}(Q^y), 0\}$. Fix an integer $T \ge 1$. Then*

$$\mathrm{Gap}\left(\frac{1}{T}\sum_{k=0}^{T-1} x^k,\ \frac{1}{T}\sum_{k=0}^{T-1} y^k\right) \le \frac{D^2\left((h+4)\varsigma_x^+ + (h+5)\varsigma_y^+ + (h+1)\eta^2\left(\varsigma_x^+ + \varsigma_y^+\right)\right)}{T}.$$

Numerical results. In our numerical experiments, for a numerical solution to the Lyapunov LMI to be considered approximately feasible, we require:

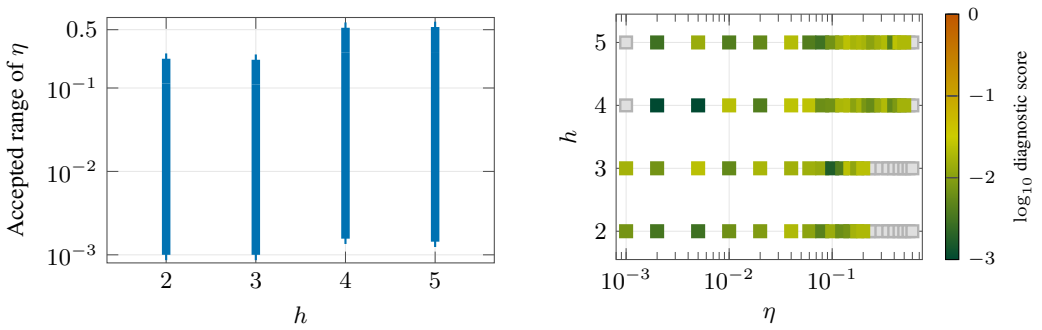

(a) History parameter h

(b) Diagnostic score over η

Figure 1: Numerical Lyapunov search for AltGDA. (a) Range of accepted stepsizes yielding approximately feasible solutions for the LMI for each history parameter $h \geq 2$. Note that for $h = 0, 1$, Mosek could not find any feasible solution. (b) Quality of the numerical solutions for (h, η) with darker green indicating higher precision solutions.

- MOSEK reports `OPTIMAL` and `FEASIBLE_POINT`.
- The largest absolute equality residual is less than or equal to 10^{-6}.
- The smallest eigenvalue among the required PSD blocks is -10^{-7} or larger.
- The smallest nonnegativity-constrained dual multiplier is -10^{-8} or larger.

Figure 1 summarizes the numerical Lyapunov search over the stepsize η and history parameter h. We find that for $h = 0, 1$, the LMI in consideration does not yield an approximately feasible solution based on our acceptance criteria. Starting from $h \geq 2$, we find that MOSEK can find an approximately feasible solution to the LMI for a range of stepsizes, as shown in Figure 1. Figure 1(a) shows the range of accepted stepsizes. On the other hand, Figure 1(b) shows how closely the computed candidates satisfy the LMI constraints based on a diagnostic score that we define as follows. Let r be the maximum absolute affine-identity residual, e the minimum eigenvalue over the required PSD blocks, and ℓ the minimum nonnegativity-constrained multiplier. Then the diagnostic score for a numerical solution is given by:

$$s = \max\left\{\frac{r}{10^{-6}}, \frac{(-10^{-7} - e)_+}{10^{-7}}, \frac{(-10^{-8} - \ell)_+}{10^{-8}}\right\}, \text{ where } (a)_+ = \max\{a, 0\},$$

where Figure 1(b) displays $\log_{10}(\max\{s, 10^{-16}\})$. All grids with dark green, light green, or yellow-green color satisfy the aforementioned numerical acceptance tests. The gray grids show the pairs that failed the solver-status requirement.

6 Last-Iterate Behavior

In the previous sections, we established the global $O(1/T)$ convergence rate of AltGDA inspired by the PEP-based framework that seeks a quadratic Lyapunov certificate. A natural question is: *can we achieve a last-iterate convergence rate for AltGDA by tuning its stepsize $\eta > 0$?* Unfortunately, the answer is no: we give simple counterexamples showing that the last-iterate duality gap of AltGDA does not converge to zero both in matrix games and in bilinear problems over general compact convex sets. For AltGDA over compact convex sets, we provide a simple 2-dimensional counterexample over the unit ball such that no matter what $\eta > 0$ we use, the last-iterate duality gap remains a constant across all iterations. Second, for matrix games with $\mathcal{X} = \Delta_n$ and $\mathcal{Y} = \Delta_m$, we present one 2×2 matrix and initialization such that AltGDA has its last-iterate duality gap bounded away from zero for every $\eta > 0$. Our explicit counterexamples are motivated again by a PEP-based framework (Drori & Teboulle, 2014; Taylor et al., 2017b) that numerically computes the worst-case last-iterate duality gap of AltGDA over compact convex sets.

Appendix C.8 proves the results associated with both counterexamples below.

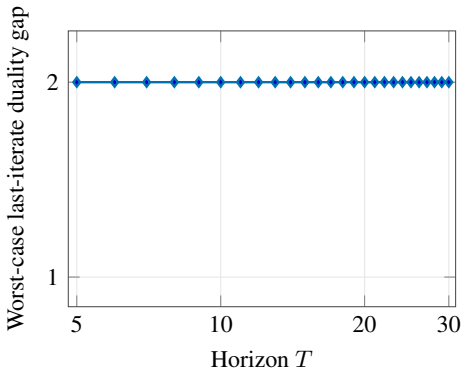

Figure 2: Numerical confirmation of the exact worst-case last-iterate gap 2 over normalized compact convex sets (Example 1), for $T = 5, \ldots, 30$.

Example 1 (Last-iterate duality gap of AltGDA for bilinear problems over compact convex sets)**.** Let $\mathcal{P}_T(\eta)$ denote the worst-case last-iterate duality gap of projected AltGDA over all dimensions, all compact convex sets $\mathcal{X}, \mathcal{Y}$ contained in Euclidean unit balls, all A with $\|A\|_2 \leq 1$, and all feasible initializations. For every constant stepsize $\eta > 0$ and time horizon $T \geq 1$, we have

$$\mathcal{P}_T(\eta) = 2, \text{ and } \inf_{\eta>0} \mathcal{P}_T(\eta) = 2.$$

The upper bound is attained on two-dimensional unit balls with $A = I_2$. The construction has gap 2 at every iterate, including initialization.

Example 2 (Last-iterate duality gap of AltGDA for matrix games)**.** On $\Delta_2 \times \Delta_2$, consider

$$A = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \text{ with } x^0 = (3/4, 1/4), \text{ and } y^0 = (1/2, 1/2).$$

The unique Nash equilibrium for this game is $x^\star = y^\star = (1/2, 1/2)$. For this matrix game, for every fixed stepsize $\eta > 0$, AltGDA has its last-iterate duality gap bounded away from zero for every $\eta > 0$, i.e., it satisfies $\inf_{t\geq 0} \mathrm{Gap}(x^t, y^t) > 0$.

By our convergence results, for every fixed $0 < \eta \leq 1/2$, while in last-iterate duality gap we have $\mathrm{Gap}(x^t, y^t) \geq 1/2$, AltGDA converges in ergodic duality gap at a rate $\mathrm{Gap}(\overline{x}^T, \overline{y}^T) \leq 5/(4\eta T)$. Thus averaging yields convergence on a trajectory whose last-iterate gap never vanishes. We provide the proof to this claim in Remark C.1.

Figure 2 validates Example 1 numerically using the PEP framework. It shows the computed worst-case last-iterate duality gaps for $T = 5, \ldots, 30$ are essentially constant at 2 for all practical purposes over a stepsize grid containing 25 points in $[1/64, 2]$. In Appendix C.7 we describe the PEP framework in detail.

7 Conclusion

In this paper we established a global $O(1/T)$ ergodic convergence rate for AltGDA in every finite two-player zero-sum matrix game, for every initialization and every horizon T. Thus, alternation alone suffices to obtain this rate in general matrix games. We have formalized and machine-checked this global ergodic convergence result in Lean 4. Our proof was guided by numerical observations from a PEP framework for quadratic Lyapunov function search over compact convex sets, which illustrates how systematic numerical searches of Lyapunov functions can inform analytical convergence proofs. Additionally, we provided simple counterexamples showing that the last-iterate duality gap of AltGDA does not converge to zero both in matrix games and in bilinear problems over general compact convex sets. This last-iterate behavior justifies why averaging of iterates is indeed necessary to achieve an $O(1/T)$ rate.

AI USE STATEMENT

For developing our numerical PEP framework and deriving analytical convergence results in this paper, we designed a *domain-specific harness* using our domain expertise and used it with GPT 5.6 Sol. A domain-specific harness is a user-written framework built around a general purpose LLM to enhance the LLM's domain specific capabilities. Our harness for the design and analysis of first-order methods for constrained bilinear games is inspired by the AutoOPT harness of Kim et al. (2026) that is designed for an unconstrained single function setup. Our harness is composed of three agent *skill*s; an agent skill is a collection of simple markdown files and scripts for extending agentic capabilities of an LLM, for more information see `https://agentskills.io/home`. We designed the first *skill* to assist us in modeling and solving the PEP problems arising in constrained bilinear games to assist us with the formulation, derivation, implementation, and validation of the Lyapunov LMI PEP problem and the finite-horizon last-iterate duality-gap PEP problem in Section 5 and Section 6. The second skill of our harness helped us in developing analytical convergence arguments described in Section 4 by observing and discovering analytical patterns obtained by solving the Lyapunov LMI PEP problem numerically. After manually verifying the correctness of the analytical convergence proofs, in the final step we used a third component of our harness to help us in formalizing and machine checking all these convergence results in Lean 4.32.0 with mathlib 4.32.0. For our Lean code, we ran two rounds of checks. In the first check, we used Lean to verify the ergodic convergence rate results with every step being explicit and the underlying logic verified, whereas in the second step we used `Comparator` independently (`https://github.com/leanprover/comparator`) to ensure that the Lean code proves exactly the statements it claims, does not use any forbidden axioms, and the submitted proofs are accepted by Lean kernel.

We thoroughly reviewed any AI-assisted mathematical arguments, code, computations, and prose and we accept responsibility for the correctness of the results in the manuscript.

ETHICS STATEMENT

Our work is theoretical and computational, and it uses no human subjects, personal or sensitive data, user studies, or deployed decision systems. Our paper studies convergence in bilinear saddle point problems and we make no deployment recommendation either.

REPRODUCIBILITY STATEMENT

The zip file accompanying with this paper provides both the Lean verification code and Julia code for reproducing the analytical and numerical results in this paper, respectively. All the convergence results in Sections 3 and 4.2 and in Appendix A of the paper are formalized and machine-checked using Lean 4.32.0 with mathlib 4.32.0. The numerical PEP frameworks for the Lyapunov-function search in Section 5 and Appendix B and for the finite-horizon PEP study for investigating last-iterate behavior of AltGDA can be reproduced using Julia.

REFERENCES

MOSEK ApS. *The MOSEK Optimizer API for Julia. Version 11.2.*, 2026. URL `https://docs.mosek.com/latest/juliaapi/index.html`.

James P. Bailey, Gauthier Gidel, and Georgios Piliouras. Finite regret and cycles with fixed step-size via alternating gradient descent-ascent. In *Conference on Learning Theory*, pp. 391–407, 2020.

Heinz H. Bauschke and Patrick L. Combettes. *Convex Analysis and Monotone Operator Theory in Hilbert Spaces*. Springer, second edition, 2017. doi: 10.1007/978-3-319-48311-5.

Nizar Bousselmi, Julien M. Hendrickx, and François Glineur. Interpolation conditions for linear operators and applications to performance estimation problems. *SIAM Journal on Optimization*, 34(3):3033–3063, 2024. doi: 10.1137/23M1575391.

Michael Bowling, Neil Burch, Michael Johanson, and Oskari Tammelin. Heads-up limit hold'em poker is solved. *Science*, 347(6218):145–149, 2015.

Noam Brown and Tuomas Sandholm. Superhuman AI for heads-up no-limit poker: Libratus beats top professionals. *Science*, 359(6374):418–424, 2018.

Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890, 2019.

Neil Burch, Matej Moravcik, and Martin Schmid. Revisiting CFR+ and alternating updates. *Journal of Artificial Intelligence Research*, 64:429–443, 2019. doi: 10.1613/jair.1.11370.

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

Volkan Cevher, Ashok Cutkosky, Ali Kavis, Georgios Piliouras, Stratis Skoulakis, and Luca Viano. Alternation makes the adversary weaker in two-player games. In *Advances in Neural Information Processing Systems*, volume 36, pp. 18263–18290, 2023.

Antonin Chambolle and Thomas Pock. A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision*, 40(1):120–145, 2011. doi: 10.1007/s10851-010-0251-1.

Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. *arXiv preprint arXiv:1807.04252*, 2018.

Constantinos Daskalakis, Andrew Ilyas, Vasilis Syrgkanis, and Haoyang Zeng. Training gans with optimism. *arXiv preprint arXiv:1711.00141*, 2017.

Leonardo de Moura and Sebastian Ullrich. The lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28: 28th International Conference on Automated Deduction, Virtual Event, July 12–15, 2021, Proceedings*, pp. 625–635, Berlin, Heidelberg, 2021. Springer-Verlag. ISBN 978-3-030-79875-8. doi: 10.1007/978-3-030-79876-5_37. URL https://doi.org/10.1007/978-3-030-79876-5_37.

Yoel Drori and Marc Teboulle. Performance of first-order methods for smooth convex minimization: A novel approach. *Mathematical Programming*, 145(1–2):451–482, 2014. doi: 10.1007/s10107-013-0653-0.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Online convex optimization for sequential decision processes and extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1917–1925, 2019.

Yi Feng, Kaito Fujii, Stratis Skoulakis, Xiao Wang, and Volkan Cevher. Continuous-time analysis of heavy ball momentum in min-max games. In *International Conference on Machine Learning*, pp. 16670–16710, 2025.

A. J. Goldman and A. W. Tucker. Theory of linear programming. In H. W. Kuhn and A. W. Tucker (eds.), *Linear Inequalities and Related Systems*, volume 38 of *Annals of Mathematics Studies*, pp. 53–97. Princeton University Press, Princeton, NJ, 1956.

Baptiste Goujaud, Aymeric Dieuleveut, and Adrien B. Taylor. On fundamental proof structures in first-order optimization. In *IEEE Conference on Decision and Control*, 2023. URL https://arxiv.org/abs/2310.02015.

Soumita Hait, Ping Li, Haipeng Luo, and Mengxiao Zhang. Alternating regret for online convex optimization. In *Conference on Learning Theory*, pp. 2632–2633, 2025.

Jonas Katona, Xiuyuan Wang, and Andre Wibisono. A symplectic analysis of alternating mirror descent. *Journal of Machine Learning Research*, 27(44):1–61, 2026.

Heechang Kim, Ernest K. Ryu, and Shuvomoy Das Gupta. A domain-specific harness for end-to-end automation of optimization research. *arXiv preprint arXiv:2608.07407*, 2026. doi: 10.48550/arXiv.2608.07407.

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

Christian Kroer. Lecture note 5: Computing Nash equilibrium via regret minimization. Lecture notes for IEOR8100: Economics, AI, and Optimization, Columbia University, February 2020. URL https://web.archive.org/web/20221002062403/http://www.columbia.edu/~ck2945/files/s20_8100/lecture_note_5_nash_from_rm.pdf.

John Lazarsfeld, Georgios Piliouras, Ryann Sim, and Stratis Skoulakis. Optimism without regularization: Constant regret in zero-sum games. In *Advances in Neural Information Processing Systems*, 2025.

Jaewook Lee, Hanseul Cho, and Chulhee Yun. Fundamental benefit of alternating updates in minimax optimization. In *International Conference on Machine Learning*, pp. 26439–26514, 2024.

Miles Lubin, Oscar Dowson, Joaquim Dias Garcia, Joey Huchette, Benoît Legat, and Juan Pablo Vielma. JuMP 1.0: Recent improvements to a modeling language for mathematical optimization. *Mathematical Programming Computation*, 2023. doi: 10.1007/s12532-023-00239-3.

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. *arXiv preprint arXiv:1807.02629*, 2018a.

Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 2703–2717, 2018b.

Meta Fundamental AI Research Diplomacy Team (FAIR), Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of Diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *International Conference on Artificial Intelligence and Statistics*, pp. 1497–1507. PMLR, 2020.

Matej Moravčík, Martin Schmid, Neil Burch, Viliam Lisỳ, Dustin Morrill, Nolan Bard, Trevor Davis, Kevin Waugh, Michael Johanson, and Michael Bowling. DeepStack: Expert-level artificial intelligence in heads-up no-limit poker. *Science*, 356(6337):508–513, 2017.

Tianlong Nan, Shuvomoy Das Gupta, Garud Iyengar, and Christian Kroer. On the $O(1/T)$ convergence of alternating gradient descent–ascent in bilinear games. In *International Conference on Learning Representations*, 2026.

Arkadi Nemirovski. Prox-method with rate of convergence $O(1/t)$ for variational inequalities with Lipschitz continuous monotone operators and smooth convex-concave saddle point problems. *SIAM Journal on Optimization*, 15(1):229–251, 2004. doi: 10.1137/S1052623403425629.

Yurii Nesterov. Dual extrapolation and its applications to solving variational inequalities and related problems. *Mathematical Programming*, 109(2–3):319–344, 2007. doi: 10.1007/s10107-006-0034-z.

Julien Perolat, Bart De Vylder, Daniel Hennes, Eugene Tarassov, Florian Strub, Vincent de Boer, Paul Muller, Jerome T. Connor, Neil Burch, Thomas Anthony, et al. Mastering the game of Stratego with model-free multiagent reinforcement learning. *Science*, 378(6623):990–996, 2022.

Leonid D. Popov. A modification of the Arrow-Hurwicz method for search of saddle points. *Mathematical Notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980. doi: 10.1007/BF01141092.

Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Conference on Learning Theory*, pp. 993–1019, 2013a.

Alexander Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, pp. 3066–3074, 2013b.

Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, pp. 2989–2997, 2015.

Oskari Tammelin, Neil Burch, Michael Johanson, and Michael Bowling. Solving heads-up limit Texas Hold'em. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.

Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Exact worst-case performance of first-order methods for composite convex optimization. *SIAM Journal on Optimization*, 27(3): 1283–1313, 2017a. doi: 10.1137/16M108104X.

Adrien B. Taylor, Julien M. Hendrickx, and François Glineur. Smooth strongly convex interpolation and exact worst-case performance of first-order methods. *Mathematical Programming*, 161(1–2): 307–345, 2017b. doi: 10.1007/s10107-016-1009-3.

Manu Upadhyaya, Sebastian Banert, Adrien B. Taylor, and Pontus Giselsson. Automated tight Lyapunov analysis for first-order methods. *Mathematical Programming*, 209:133–170, 2025a. doi: 10.1007/s10107-024-02061-8.

Manu Upadhyaya, Shuvomoy Das Gupta, Adrien B. Taylor, Sebastian Banert, and Pontus Giselsson. The AutoLyap software suite for computer-assisted Lyapunov analyses of first-order methods. *arXiv preprint arXiv:2506.24076*, 2025b.

J. von Neumann. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen*, 100(1):295–320, 1928. doi: 10.1007/BF01448847.

Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. *arXiv preprint arXiv:2006.09517*, 2020.

Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Last-iterate convergence of decentralized optimistic gradient descent/ascent in infinite-horizon competitive Markov games. In *Conference on Learning Theory*, 2021.

Andre Wibisono, Molei Tao, and Georgios Piliouras. Alternating mirror descent for constrained min-max games. In *Advances in Neural Information Processing Systems*, volume 35, pp. 35201–35212, 2022.

A. C. Williams. Complementarity theorems for linear programming. *SIAM Review*, 12(1):135–137, 1970. doi: 10.1137/1012015.

Guodong Zhang, Yuanhao Wang, Laurent Lessard, and Roger B. Grosse. Near-optimal local convergence of alternating gradient descent-ascent for minimax optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 7659–7679, 2022.

A Full Proof of the Main Theorem

A.1 Overview, setting, and notation

This appendix gives a self-contained proof of Theorem 1. It first restates the setting of Section 3 and then constructs the analysis witness (we call the fixed Goldman-Tucker strictly complementary saddle point a *witness* in the proof) needed to state the result again as Theorem A.1. The roadmap at the end of this subsection aligns the detailed proof with the proof sketch in Section 4.2.

Consider a finite two-player zero-sum game with payoff matrix $A \in \mathbb{R}^{m\times n}$: one player mixes over n actions to minimize the expected payoff and the other mixes over m actions to maximize it. The simplest learning dynamics for such a game is letting both players run projected gradient steps simultaneously. We study the *alternating* variant, in which the maximizer moves second, against the minimizer's freshly updated strategy. The averages of the players' strategies then approach

an equilibrium at the rate $O(1/T)$ for any $T \geq 1$, from every initialization, with any appropriate stepsize.

Let

$$\Delta_d = \left\{ z \in \mathbb{R}^d : \; z \geq 0, \; \mathbf{1}^\top z = 1 \right\}$$

denote the probability simplex in $\mathbb{R}^d$, where $\mathbb{R}^d$ is the d dimensional Euclidean space. The matrix game can be represented by

$$\min_{x \in \Delta_n} \max_{y \in \Delta_m} y^\top A x, \tag{A.1}$$

where the x-player minimizes and the y-player maximizes the payoff. By the minimax theorem (von Neumann, 1928), the game has a value ν and a saddle point. The duality gap measures how near a candidate pair $(\widetilde{x}, \widetilde{y}) \in \Delta_n \times \Delta_m$ is to being a saddle point:

$$\begin{aligned} \mathrm{Gap}(\widetilde{x}, \widetilde{y}) &= \max_{y \in \Delta_m} y^\top A \widetilde{x} - \min_{x \in \Delta_n} \widetilde{y}^\top A x \\ &= \max_{x \in \Delta_n, y \in \Delta_m} \left(y^\top A \widetilde{x} - \widetilde{y}^\top A x \right), \end{aligned} \tag{A.2}$$

which is nonnegative everywhere and zero exactly at saddle points.

We study alternating projected gradient descent–ascent (AltGDA). With a fixed stepsize $\eta > 0$ and Π_C denoting Euclidean projection onto a closed convex set C, AltGDA generates

$$x^{t+1} = \Pi_{\Delta_n}\left(x^t - \eta A^\top y^t\right), \qquad y^{t+1} = \Pi_{\Delta_m}\left(y^t + \eta A x^{t+1}\right), \qquad \forall\, t \geq 0. \tag{A.3}$$

The alternation trick is applied as we use x^{t+1} to update y^{t+1} instead of using x^t. We track the uniform averages of the players' strategies:

$$\overline{x}^T = \frac{1}{T} \sum\nolimits_{t=1}^{T} x^t, \qquad \overline{y}^T = \frac{1}{T} \sum\nolimits_{t=1}^{T} y^t. \tag{A.4}$$

The main result. It is known that every matrix game possesses a strictly complementary saddle point, by the Goldman–Tucker theorem. From one such saddle point, we can extract a scalar parameter $\delta \in (0, 1]$, called here the *Goldman–Tucker separation constant* (see A.2): roughly speaking, this parameter captures the smallest equilibrium probability on the supports (of the selected saddle point) and the smallest normalized off-support payoff slack.

Then, Theorem A.1 states that, for any $L > 0$ with $\|A\|_2 \leq L$, AltGDA with every fixed stepsize $0 < \eta \leq \delta/(2\sqrt{2}\, L)$ guarantees that

$$\mathrm{Gap}\left(\overline{x}^T, \overline{y}^T\right) \leq \frac{15}{2\eta T} \quad \text{for every } T \geq 1 \text{ and every initialization } (x^0, y^0).$$

The main difficulty. The standard route to an $O(1/T)$ ergodic rate is an energy function argument: by exhibiting a nonnegative energy whose per-step decrease dominates the residual quantity to be summed, we can show that the residual terms in the regret analysis are bounded by a constant that is independent of T no matter how large the horizon T is, which leads to an $O(1/T)$ ergodic rate.

This is how the $O(1/T)$ convergence rate is proven under the assumption of the existence of an interior NE (Nan et al., 2026). However, in the general matrix games without an interior NE, the difference of the usual energy function no longer yields an upper bound for the residual terms.

In our proof, we identify that each AltGDA step obeys an *exact* energy balance identity (Lemma A.3),

$$V_{t+1} - V_t + \mathscr{D}_t = r_t,$$

where $\mathscr{D}_t \geq 0$ is the "dissipation" term, and $r_t \geq 0$ is the residual term that needs to be controlled. The residual term corresponds to a collision term supported on deactivated coordinates. It is positive exactly when at least one such coordinate has both positive previous mass and a positive projection multiplier (Lemma A.2). The proof then shows that the residual terms, summed over an arbitrary horizon T, cost only a constant.

The mechanism. Section A.3 first gives the KKT representation of the two projections and identifies the residual r_t exactly; Section A.4 then proves the energy balance for an alternation-corrected energy V_t: squared distance to the selected saddle point plus a bilinear cross term scaled to account for the alternation.

Section A.5 develops the following next steps. For each player, eliminating the shared projection threshold from the KKT system *reflects* the update into the off-support coordinates, where a bounded contrast field drives the dynamics. Telescoping the associated storage term leaves increments of size $\eta^2 L^2$ weighted by off-support mass, which the payoff slacks convert into dissipation. In the next step, the on-support multiplier activity cancels against the threshold shift. Identity (A.39) leaves only the *centered* movement, with the dimension-free factor $\sqrt{1-1/k}$, where k is a support size.

The resulting pointwise bound $r_t \le B_t - B_{t+1} + \frac{1}{2}\mathscr{D}_t$, with bounded storage B_t, combines with the exact energy balance to give the residual budget $\sum_{t<T} r_t \le 4 + (2+8\sqrt{2})\,\eta L$ in Section A.6. The next step uses two shifted telescopings of the projection inequalities to convert that budget into the gap bound in Section A.7.

Notation. The vector $\mathbf{1}$ has every entry one, with dimension determined by context; I_d is the $d \times d$ identity; e_i is the ith standard basis vector. Vector inequalities are coordinatewise, and $\odot$ is the coordinatewise product. For an index set S, we write S^c for its complement, $|S|$ for its cardinality, $\mathbf{1}_S$ for its indicator vector, and g_S for the corresponding subvector of g; $\mathrm{supp}(z)$ is the set of indices of nonzero coordinates of z. The norm $\|\cdot\|_2$ is Euclidean on vectors and spectral on matrices, and for a symmetric positive semidefinite matrix G we write $\|z\|_G^2 = z^\top G z$.

A.2 The Goldman–Tucker saddle point and the main theorem

The proof uses one strictly complementary saddle point at which supports and payoff slacks separate cleanly. Writing ν for the game value, the x-player's payoff slack on coordinate i is $(A^\top y^\star)_i - \nu$. At this selected saddle point, on-support coordinates ($x_i^\star > 0$) have zero payoff slack, i.e., $(A^\top y^\star)_i = \nu$, whereas off-support coordinates ($x_i^\star = 0$) have strictly positive payoff slack, i.e., $(A^\top y^\star)_i > \nu$.

The existence of such a saddle point can be shown by writing the matrix games as linear programs. The bilinear two player zero sum games, written as linear programs in the strategy and a value variable, are

$$\min_{x,\theta} \quad \theta \qquad \text{subject to} \quad Ax \le \theta\mathbf{1}, \quad \mathbf{1}^\top x = 1, \quad x \ge 0, \tag{A.5}$$

$$\max_{y,\omega} \quad \omega \qquad \text{subject to} \quad A^\top y \ge \omega\mathbf{1}, \quad \mathbf{1}^\top y = 1, \quad y \ge 0. \tag{A.6}$$

These constitute a primal-dual linear programming pair, and the strict-complementarity theorem for general-form programs (Williams, 1970) applies; equivalently, the Goldman-Tucker theorem (Goldman & Tucker, 1956) applies after a routine conversion of the free value variables to standard form. Either theorem yields an optimal pair $(x^\star, y^\star)$ with common value ν whose slack vectors

$$s_x = A^\top y^\star - \nu\mathbf{1}, \qquad s_y = \nu\mathbf{1} - Ax^\star \tag{A.7}$$

are nonnegative and satisfy, coordinatewise,

$$x_i^\star s_{x,i} = 0, \qquad x_i^\star + s_{x,i} > 0, \qquad y_j^\star s_{y,j} = 0, \qquad y_j^\star + s_{y,j} > 0. \tag{A.8}$$

In words, on each coordinate, either the equilibrium strategy places mass or the payoff slack is positive, and never both. In our proof of the main theorem, we first fix one such pair and call it the *Goldman-Tucker saddle point*. In the remainder of this appendix, we use $(x^\star, y^\star)$ to denote the selected Goldman-Tucker saddle point.

This saddle point determines the index sets (on-support and off-support sets of coordinates) and separation constants of the analysis. Let $I = \mathrm{supp}(x^\star)$ and $J = \mathrm{supp}(y^\star)$, with counts $k_x = |I|$, $\ell_x = |I^c|$, $k_y = |J|$, and $\ell_y = |J^c|$. Note that the support sizes k_x, k_y are nonzero. We define the smallest on-support probabilities:

$$\underline{x} = \min_{i \in I} x_i^\star, \qquad \underline{y} = \min_{j \in J} y_j^\star,$$

and the smallest off-support payoff slacks:

$$\sigma_x = \min_{i \in I^c} s_{x,i}, \qquad \sigma_y = \min_{j \in J^c} s_{y,j},$$

where σ_x is defined when $I^c \neq \varnothing$ and σ_y is defined when $J^c \neq \varnothing$. Thus, $\underline{x}, \underline{y} > 0$, and every defined payoff slack minimum is positive. Let $L > 0$ satisfy $\|A\|_2 \leq L$, and set

$$\delta = \min\left\{\underline{x},\ \underline{y},\ \frac{\sigma_x}{L} \text{ if } I^c \neq \varnothing,\ \frac{\sigma_y}{L} \text{ if } J^c \neq \varnothing\right\}, \tag{A.9}$$

where terms corresponding to empty off-support index sets (i.e., $|I| = n$ or $|J| = m$) are omitted. Then, we have $\delta > 0$ and $\delta \leq 1$ because $\underline{x} \leq 1$ and $\underline{y} \leq 1$. We call this scalar the *Goldman-Tucker separation constant* associated with the selected Goldman-Tucker saddle point and the chosen bound L. Informally, it records how decisively the Goldman-Tucker saddle point separates its supports from their off-support coordinates, which can be measured by the minimum of *the least probability assigned to an on-support action* and *the least normalized payoff slack of an off-support action*.

The following theorem restates Theorem 1.

Theorem A.1. *Let $A \in \mathbb{R}^{m \times n}$, let $L > 0$ satisfy $\|A\|_2 \leq L$, and define δ from a selected Goldman–Tucker strictly complementary saddle point as in* (A.9)*. For every initialization $(x^0, y^0) \in \Delta_n \times \Delta_m$, every $T \geq 1$, and every fixed stepsize satisfying*

$$0 < \eta \leq \frac{\delta}{2\sqrt{2}\, L}, \tag{A.10}$$

the averages (A.4) *of the AltGDA iterates* (A.3) *satisfy, in the duality gap* (A.2)*,*

$$\boxed{\mathrm{Gap}(\overline{x}^T, \overline{y}^T) \leq \frac{15}{2\eta T}.} \tag{A.11}$$

In the subsequent sections, we assume a nonzero $A \neq 0$ (otherwise the matrix game is trivial). Note that we can choose an arbitrary $L > 0$ with $\|A\|_2 \leq L$, and in practice this can be simply done by setting $L = \|A\|_2$.

A.3 Projection and residual identities

The first-order description of the two projections provides upper bounds for the following "per-step iterate movement":

$$\Delta x_t = x^{t+1} - x^t, \qquad \Delta y_t = y^{t+1} - y^t, \qquad v_t = -A^\top y^t, \qquad u_t = A x^{t+1} \tag{A.12}$$

Here, v_t and u_t are shorthand notations for the signed payoff vectors that drive the movement of the iterates. We use Δ to denote the forward difference at some time index: for any time-indexed quantity a_t, we write $\Delta a_t = a_{t+1} - a_t$.

The first-order description also provides a classic and useful residual representation, which we call the projection KKT representation of the AltGDA algorithm.

Lemma A.1 (Projection KKT representation of AltGDA)**.** *For every $t \geq 0$, there are scalars γ_t, λ_t and vectors $\mu_t \in \mathbb{R}^n_+, \rho_t \in \mathbb{R}^m_+$ such that*

$$\Delta x_t = \eta(v_t - \gamma_t \mathbf{1} + \mu_t), \qquad \mu_t \odot x^{t+1} = 0, \tag{A.13}$$

$$\Delta y_t = \eta(u_t - \lambda_t \mathbf{1} + \rho_t), \qquad \rho_t \odot y^{t+1} = 0. \tag{A.14}$$

Moreover, we have

$$\|\Delta x_t\|_2 \leq \eta L, \qquad \|\Delta y_t\|_2 \leq \eta L. \tag{A.15}$$

Proof. By the KKT conditions for the Euclidean projection $p^+ = \Pi_{\Delta_d}(p + \eta g)$, we have

$$p^+ - p = \eta(g - \gamma \mathbf{1} + \mu), \qquad \mu \geq 0, \qquad \mu \odot p^+ = 0, \tag{A.16}$$

where γ is the (scalar) multiplier of the sum constraint and μ is a vector corresponding to the multipliers of the sign constraints. Applying Eq. (A.16) with $(p, g) = (x^t, v_t)$ and $(p, g) = (y^t, u_t)$, we have Eqs. (A.13) and (A.14).

For the per-step iterate movement bound, we use the nonexpansive property of the projections and the definition of the induced 2-norm:

$$\|\Delta x_t\|_2 = \left\|\Pi_{\Delta_n}(x^t - \eta A^\top y^t) - \Pi_{\Delta_n}(x^t)\right\|_2 \le \eta \|A^\top y^t\|_2 \le \eta L \|y^t\|_2 \le \eta L,$$

where the last step uses $\|y^t\|_2 \le \mathbf{1}^\top y^t = 1$. The y-iterate movement bound follows identically. □

In Eqs. (A.13) and (A.14), the multiplier μ_t is strictly positive only on those coordinates on which the next-step iterate touches the simplex boundary ($x_i^{t+1} = 0$), and likewise for ρ_t.

Next, we introduce a quantity called "residual term" in the proof:

$$r_t := \eta v_t^\top \Delta x_t + \eta u_t^\top \Delta y_t - \|\Delta x_t\|_2^2 - \|\Delta y_t\|_2^2. \tag{A.17}$$

This is a key quantity in the convergence proof for AltGDA. In the unconstrained case, each iterate movement equals η times the corresponding payoff vector and hence r_t vanishes. In the constrained case, it is much trickier to handle. The next lemma computes a more useful equivalent closed-form expression of r_t. Since $\eta > 0$ and the probabilities and projection multipliers are nonnegative, Eq. (A.18) implies that $r_t > 0$ if and only if there is a coordinate i with $x_i^t > 0$ and $\mu_{t,i} > 0$, or a coordinate j with $y_j^t > 0$ and $\rho_{t,j} > 0$. By complementarity in Lemma A.1, these conditions imply $x_i^{t+1} = 0$ or $y_j^{t+1} = 0$, respectively. Thus, the residual is strictly positive exactly when at least one player has a previously positive coordinate reach the simplex boundary with a strictly positive projection multiplier; merely touching the boundary is not sufficient.

Lemma A.2 (Exact residual representation)**.** *For every $t \ge 0$,*

$$r_t = \eta \left(\mu_t^\top x^t + \rho_t^\top y^t\right) \ge 0. \tag{A.18}$$

Proof. Since $\mathbf{1}^\top \Delta x_t = 0$, by Eq. (A.13) we have

$$\begin{aligned}
\eta v_t^\top \Delta x_t - \|\Delta x_t\|_2^2 &= \eta \left(v_t - \frac{\Delta x_t}{\eta}\right)^\top \Delta x_t \\
&= \eta(\gamma_t \mathbf{1} - \mu_t)^\top \Delta x_t \\
&= -\eta \mu_t^\top (x^{t+1} - x^t) = \eta \mu_t^\top x^t,
\end{aligned}$$

where the last equality follows because $\mu_t^\top x^{t+1} = 0$. By a symmetric proof, we can show that the y-player's identity holds as well. Furthermore, r_t is nonnegative for any $t \ge 0$ because every factor in Eq. (A.18) is nonnegative. □

A.4 An energy function for AltGDA

Another important component in the proof is an energy function. A natural choice of the energy function for the minimax problems is the squared distance to a saddle point. To account for the asymmetry introduced by alternation, we add a bilinear cross term of size η. With the Goldman-Tucker saddle point $(x^\star, y^\star)$ fixed, define

$$V_t = \|x^t - x^\star\|_2^2 + \|y^t - y^\star\|_2^2 - \eta (y^t - y^\star)^\top A (x^t - x^\star). \tag{A.19}$$

Alongside the energy, the analysis tracks two nonnegative quantities built from the current iterates and the Goldman-Tucker saddle point: the payoff slack-weighted probabilities

$$P_t^x = s_x^\top x^t, \qquad P_t^y = s_y^\top y^t, \qquad P_t = P_t^x + P_t^y,$$

and the equilibrium probability-weighted multipliers

$$E_t^x = \mu_t^\top x^\star, \qquad E_t^y = \rho_t^\top y^\star, \qquad E_t = E_t^x + E_t^y.$$

In particular, taking the x-player side as an example, because s_x vanishes on I, the quantity P_t^x measures the off-support iterate probabilities, with each coordinate weighted by its payoff slack. Similarly, because $x^\star$ vanishes outside I, the quantity E_t^x measures the multipliers on the support, with each coordinate weighted by its equilibrium probability.

We then define a quantity that combines P_t and E_t:

$$\mathscr{D}_t = \eta(P_t + P_{t+1}) + 2\eta E_t. \tag{A.20}$$

We call the above quantity the *dissipation* of the t-th step, and all of these three quantities (P_t, E_t, and $\mathscr{D}_t$) are nonnegative by definition.

The central identity of the analysis is that energy, dissipation, and residual are in exact balance.

Lemma A.3 (Exact energy identity)**.** *For every $t \geq 0$,*

$$\boxed{V_{t+1} - V_t + \mathscr{D}_t = r_t.} \tag{A.21}$$

Proof. Set

$$a = x^t - x^\star, \qquad b = y^t - y^\star, \qquad p = \Delta x_t, \qquad h = \Delta y_t.$$

Because the sums over all coordinates of both a and b are zero, the KKT conditions Eqs. (A.13) and (A.14) and the definitions of the payoff slack in Eq. (A.7) imply

$$a^\top p = -\eta P_t^x - \eta b^\top A a + \eta \mu_t^\top a, \tag{A.22}$$

$$b^\top h = -\eta P_t^y + \eta b^\top A a + \eta b^\top A p + \eta \rho_t^\top b. \tag{A.23}$$

Indeed, $a^\top A^\top y^\star = a^\top s_x = P_t^x$ because $\mathbf{1}^\top a = 0$ and $s_x^\top x^\star = 0$, and $b^\top A x^\star = -b^\top s_y = -P_t^y$ by the same arguments.

Expanding the two squared distances and the bilinear term in the definition of V_t in Eq. (A.19), and then using Eqs. (A.22) and (A.23), we have

$$\begin{aligned} V_{t+1} - V_t = & -2\eta P_t + 2\eta(\mu_t^\top a + \rho_t^\top b) + \|p\|_2^2 + \|h\|_2^2 \\ & + \eta b^\top A p - \eta h^\top A(a + p). \end{aligned} \tag{A.24}$$

Furthermore, it holds that

$$P_{t+1} = P_t + y^{\star\top} A p - h^\top A x^\star,$$

and, by the definition of r_t in Eq. (A.17),

$$-r_t = \eta y^{t\top} A p - \eta h^\top A x^{t+1} + \|p\|_2^2 + \|h\|_2^2.$$

Then, substituting $y^t = y^\star + b$ and $x^{t+1} = x^\star + a + p$ into these two expressions transforms (A.24) into

$$V_{t+1} - V_t = -\eta(P_t + P_{t+1}) + 2\eta\left[\mu_t^\top (x^t - x^\star) + \rho_t^\top (y^t - y^\star)\right] - r_t. \tag{A.25}$$

By Lemma A.2 and the definition of E_t, the terms in the bracket equal $r_t/\eta - E_t$. By substituting $r_t/\eta - E_t$ into Eq. (A.25), we prove Eq. (A.21). □

We also present two elementary bounds on the energy function. These bounds show the energy function is bounded by some finite constants and are useful to conclude the main theorem.

We use a shorthand notation

$$q = \eta L$$

to denote the per-step iterate movement scale (see Lemma A.1). Since $\delta \leq 1$ and our selected stepsize satisfies Eq. (A.10), we have:

$$q \leq \frac{1}{2\sqrt{2}}. \tag{A.26}$$

The bilinear cross term in the definition of V_t in Eq. (A.19) cannot overwhelm the squared distances: by Cauchy-Schwarz inequality and $2ab \leq a^2 + b^2$ for any scalars a, b, we have

$$V_t \geq \left(1 - \frac{q}{2}\right)\left(\|x^t - x^\star\|_2^2 + \|y^t - y^\star\|_2^2\right) \geq 0. \tag{A.27}$$

On the other hand, since the squared diameter of a simplex is at most 2, we can also upper bound the initial energy:

$$V_0 \leq 4 + 2q. \tag{A.28}$$

Summing Eq. (A.21) shows that the total dissipation equals the total residual plus two extra bounded energy terms. If each residual is bounded by a fraction of the dissipation plus a telescoping term, then we can show that the total residual is bounded by a constant. The next section establishes this relation.

A.5 RESIDUAL IS UPPER BOUNDED BY A FRACTION OF DISSIPATION

This section establishes an important relation between r_t and $\mathscr{D}_t$

$$r_t \le B_t - B_{t+1} + \tfrac{1}{2}\mathscr{D}_t,$$

where B_t is a uniformly bounded term. By Lemma A.2, the residual splits into the two player blocks $\eta\mu_t^\top x^t$ and $\eta\rho_t^\top y^t$. We next present the x-player's block in detail, and then show the results corresponding to y-player's block, which follow by the same argument as the x-player's block.

We first define some additional notations and an auxiliary lemma.

For a nonempty index set S and a vector g, define

$$(\mathsf{C}_S g)_i = \frac{1}{|S|}\sum_{j\in S} g_j - g_i, \qquad i \in S^c. \tag{A.29}$$

Lemma A.4. *Let S have cardinality k. For $i \in S^c$,*

$$|(\mathsf{C}_S g)_i| \le \sqrt{1+\frac{1}{k}}\,\|g\|_2. \tag{A.30}$$

If $h \in \mathbb{R}^k$ and $\overline{h} = k^{-1}\mathbf{1}^\top h$, then for every $i \in S$,

$$|\overline{h} - h_i| \le \sqrt{1-\frac{1}{k}}\,\|h\|_2. \tag{A.31}$$

Proof. For $i \notin S$,

$$(\mathsf{C}_S g)_i = \left\langle k^{-1}\mathbf{1}_S - e_i, g\right\rangle,$$

and the coefficient vector $k^{-1}\mathbf{1}_S - e_i$, $i \in S^c$ has squared norm $1/k+1$. For $i \in S$,

$$\overline{h} - h_i = \left\langle k^{-1}\mathbf{1}_S - e_i, h\right\rangle,$$

where the coefficient vector $k^{-1}\mathbf{1}_S - e_i$, $i \in S$ has squared norm $1/k - 2/k + 1 = 1 - 1/k$, because the entry at i is $1/k - 1$. Combining both inner products with Cauchy–Schwarz inequality yields Eqs. (A.30) and (A.31). □

A.5.1 UPPER BOUND THE x-PLAYER BLOCK

Assume first that $I^c \neq \varnothing$; the case $I^c = \varnothing$ is easier and is handled at the end of this subsection. We introduce a set of notations corresponding to the off-support coordinates:

$$O_x = I^c, \qquad z_t = x^t_{O_x}, \qquad M_t^x = \mathbf{1}^\top z_t.$$

In words, z_t is the restriction of the iterate x^t to the off-support coordinates and M_t^x is its total probability over all off-support coordinates. Furthermore, we define

$$d_t^x = \mathsf{C}_I(-A^\top y^t), \qquad e_t^x = \frac{1}{k_x}\sum_{i\in I}\mu_{t,i}, \qquad G_x = I_{\ell_x} + \frac{1}{k_x}\mathbf{1}\mathbf{1}^\top.$$

Here, d_t^x is the contrast of the payoff vectors on the off-support coordinates against the on-support average of the payoff vectors, e_t^x is the average of the multipliers over the on-support coordinates, and $G_x \succeq I$ is a positive semidefinite matrix.

The following lemma provides a useful decomposition of the off-support component $\eta(\mu_t)_{O_x}^\top z_t$.

Lemma A.5. *For every $t \ge 0$,*

$$G_x \Delta z_t = -\eta d_t^x - \eta e_t^x \mathbf{1} + \eta(\mu_t)_{O_x}, \tag{A.32}$$

$$\eta(\mu_t)_{O_x}^\top z_t = -\|\Delta z_t\|_{G_x}^2 - \eta(d_t^x)^\top \Delta z_t - \eta e_t^x \Delta M_t^x. \tag{A.33}$$

Proof. By summing Eq. (A.13) over I, using $\mathbf{1}^\top \Delta x_t = 0$, and rearranging terms, we have

$$\gamma_t = \frac{1}{k_x} \sum_{i \in I} v_{t,i} + e_t^x + \frac{\Delta M_t^x}{\eta k_x}.$$

Substituting the above expression into the KKT representations in Lemma A.1 proves Eq. (A.32). By taking the inner product of Eq. (A.32) with Δz_t, and using $(\mu_t)_{O_x}^\top z_{t+1} = 0$, we obtain Eq. (A.33). □

Set $d_{-1}^x = d_0^x$ and define the storage

$$B_{x,t} = \eta (d_{t-1}^x)^\top z_t. \tag{A.34}$$

Then, we have

$$-\eta (d_t^x)^\top \Delta z_t = B_{x,t} - B_{x,t+1} + \eta (d_t^x - d_{t-1}^x)^\top z_t. \tag{A.35}$$

For $t \geq 1$, linearity of C_I gives

$$d_t^x - d_{t-1}^x = \mathsf{C}_I(-A^\top \Delta y_{t-1}).$$

This difference is zero at $t = 0$ by the convention above. Combining the iterate movement bound Eq. (A.15) with the bound Eq. (A.30) on C_I yields

$$\|d_t^x - d_{t-1}^x\|_\infty \leq \sqrt{1 + \frac{1}{k_x}}\, \eta L^2. \tag{A.36}$$

We then connect the ℓ_1 norm of z_t to the dissipation term P_t^x by leveraging the definition of σ_x: since $s_{x,i} \geq \sigma_x$ for $i \in O_x$ and s_x vanishes on I, strict complementarity gives us that

$$M_t^x \leq \frac{P_t^x}{\sigma_x}. \tag{A.37}$$

Combining Eqs. (A.36) and (A.37) with Eq. (A.35), we obtain an upper bound for $-\eta (d_t^x)^\top \Delta z_t$.

Next, we turn to the on-support component. Let

$$h_t^x = (\Delta x_t)_I, \qquad \overline{h}_t^x = \frac{1}{k_x} \mathbf{1}^\top h_t^x$$

denote the on-support iterate movement and its mean. Since $\sum_{i \in I} (\Delta x_t)_i + \sum_{i \in I^c} (\Delta x_t)_i = 0$ and $(\mu_t)_I^\top x_I^{t+1} = 0$, we have

$$\Delta M_t^x = -\mathbf{1}^\top h_t^x, \qquad (\mu_t)_I^\top x_I^t = -(\mu_t)_I^\top h_t^x. \tag{A.38}$$

Combining them yields the exact identity

$$(\mu_t)_I^\top x_I^t - e_t^x \Delta M_t^x = (\mu_t)_I^\top \left(\overline{h}_t^x \mathbf{1} - h_t^x \right). \tag{A.39}$$

By the nonnegativity of $(\mu_t)_I \geq 0$, the per-step iterate movement bound Eq. (A.15), the bound Eq. (A.31) for $(\overline{h}_t^x \mathbf{1} - h_t^x)$, and the lower bound $E_t^x = \mu_t^\top x^\star \geq \underline{x} \sum_{i \in I} \mu_{t,i}$, we have

$$\begin{aligned} (\mu_t)_I^\top \left(\overline{h}_t^x \mathbf{1} - h_t^x \right) &\leq \sqrt{1 - \frac{1}{k_x}}\, \|h_t^x\|_2 \sum_{i \in I} \mu_{t,i} \\ &\leq \sqrt{1 - \frac{1}{k_x}}\, \eta L\, \frac{E_t^x}{\underline{x}}. \end{aligned} \tag{A.40}$$

Finally, we combine both off-support (Eq. (A.33)) and on-support (Eq. (A.39)) components. Note that the term $\eta e_t^x \Delta M_t^x$ in Eqs. (A.33) and (A.39) can be exactly canceled, and $-\|\Delta z_t\|_{G_x}^2 \leq 0$ because of the positive semidefiniteness. By using the equations and inequalities mentioned above, in particular, Eqs. (A.35) to (A.37) and (A.40), we attain an upper bound for the x-player block of r_t:

$$\boxed{\begin{aligned} \eta \mu_t^\top x^t \leq\ & B_{x,t} - B_{x,t+1} \\ & + \sqrt{1 + \frac{1}{k_x}}\, \frac{\eta L^2}{\sigma_x} (\eta P_t^x) + \sqrt{1 - \frac{1}{k_x}}\, \frac{\eta L}{\underline{x}} (\eta E_t^x). \end{aligned}} \tag{A.41}$$

At the end of this part, we consider $I^c = \varnothing$. In this case, we retain the definitions

$$e_t^x = \frac{1}{k_x}\sum_{i\in I}\mu_{t,i}, \qquad h_t^x = (\Delta x_t)_I, \qquad \overline{h}_t^x = \frac{1}{k_x}\mathbf{1}^\top h_t^x,$$

and set

$$M_t^x = \Delta M_t^x = B_{x,t} = 0.$$

In this case, only the last term in Eq. (A.41) is nonzero, and one can verify that Eq. (A.41) still holds by the same arguments as those showing Eq. (A.40).

A.5.2 Upper bound the y-player block

The proof for the y-player block follows the same route as that for the x-player block. Because the y-player's update at the t-th iteration is driven by Ax^{t+1}, there is a time-index change across the arguments. Again, we first assume $J^c \neq \varnothing$, and define

$$O_y = J^c, \qquad w_t = y_{O_y}^t, \qquad M_t^y = \mathbf{1}^\top w_t,$$

and

$$d_t^y = \mathsf{C}_J(Ax^t), \qquad e_t^y = \frac{1}{k_y}\sum_{j\in J}\rho_{t,j}, \qquad G_y = I_{\ell_y} + \frac{1}{k_y}\mathbf{1}\mathbf{1}^\top.$$

By eliminating λ_t exactly as in Lemma A.5, we have

$$G_y\Delta w_t = -\eta d_{t+1}^y - \eta e_t^y\mathbf{1} + \eta(\rho_t)_{O_y}, \tag{A.42}$$

$$\eta(\rho_t)_{O_y}^\top w_t = -\|\Delta w_t\|_{G_y}^2 - \eta(d_{t+1}^y)^\top\Delta w_t - \eta e_t^y\Delta M_t^y. \tag{A.43}$$

Considering the time-index changes, we define the telescoping term for the y-player:

$$B_{y,t} = \eta(d_t^y)^\top w_t. \tag{A.44}$$

Then, we can again write a decomposition of the term $-\eta(d_{t+1}^y)^\top\Delta w_t$:

$$-\eta(d_{t+1}^y)^\top\Delta w_t = B_{y,t} - B_{y,t+1} + \eta(d_{t+1}^y - d_t^y)^\top w_t, \tag{A.45}$$

and

$$d_{t+1}^y - d_t^y = \mathsf{C}_J(A\Delta x_t).$$

Hence, as in Eq. (A.36),

$$\|d_{t+1}^y - d_t^y\|_\infty \le \sqrt{1+\frac{1}{k_y}}\,\eta L^2. \tag{A.46}$$

For the on-support component, we symmetrically define

$$h_t^y = (\Delta y_t)_J, \qquad \overline{h}_t^y = \frac{1}{k_y}\mathbf{1}^\top h_t^y.$$

Then, we have

$$\Delta M_t^y = -\mathbf{1}^\top h_t^y, \qquad (\rho_t)_J^\top y_J^t = -(\rho_t)_J^\top h_t^y,$$

and therefore

$$(\rho_t)_J^\top y_J^t - e_t^y\Delta M_t^y = (\rho_t)_J^\top\left(\overline{h}_t^y\mathbf{1} - h_t^y\right). \tag{A.47}$$

Then, we can show the following two upper bounds for the y-player block:

$$M_t^y \le \frac{P_t^y}{\sigma_y}, \qquad E_t^y \ge \underline{y}\sum_{j\in J}\rho_{t,j}.$$

Combining these two bounds with Eqs. (A.43) and (A.45) to (A.47) yields

$$\boxed{\begin{aligned}\eta\rho_t^\top y^t \le\ & B_{y,t} - B_{y,t+1} \\ & + \sqrt{1+\frac{1}{k_y}}\,\frac{\eta L^2}{\sigma_y}(\eta P_t^y) + \sqrt{1-\frac{1}{k_y}}\,\frac{\eta L}{\underline{y}}(\eta E_t^y).\end{aligned}} \tag{A.48}$$

When $J^c = \varnothing$, we define

$$e_t^y = \frac{1}{k_y}\sum_{j\in J}\rho_{t,j}, \qquad h_t^y = (\Delta y_t)_J, \qquad \overline{h}_t^y = \frac{1}{k_y}\mathbf{1}^\top h_t^y,$$

and set $B_{y,t} = 0$. Again, it can be verified that Eq. (A.48) still holds by the same arguments as those showing the upper bound involving E_t^y.

A.5.3 Combining the two blocks

By choosing an appropriate stepsize satisfying Eq. (A.10), we can make the coefficients before the dissipation terms be at most $\frac{1}{2}$.

Whenever the set of the off-support coordinates is nonempty, the definition of δ in Eq. (A.9) gives $\eta L^2/\sigma_x \le 1/(2\sqrt{2})$, and $\sqrt{1+1/k_x} \le \sqrt{2}$. The same bounds hold on the y-player block. Hence, we have

$$\sqrt{1+\frac{1}{k_x}}\,\frac{\eta L^2}{\sigma_x} \le \frac{1}{2}, \qquad \sqrt{1+\frac{1}{k_y}}\,\frac{\eta L^2}{\sigma_y} \le \frac{1}{2}. \tag{A.49}$$

Similarly, $\eta L \le \delta/(2\sqrt{2})$ and $\delta \le \min\{\underline{x}, \underline{y}\}$ give

$$\sqrt{1-\frac{1}{k_x}}\,\frac{\eta L}{\underline{x}} \le \frac{1}{2\sqrt{2}} < \frac{1}{2}, \qquad \sqrt{1-\frac{1}{k_y}}\,\frac{\eta L}{\underline{y}} \le \frac{1}{2\sqrt{2}} < \frac{1}{2}. \tag{A.50}$$

Let

$$B_t := B_{x,t} + B_{y,t}$$

be the combined telescoping term. Combining the block bounds Eqs. (A.41) and (A.48) and applying Eqs. (A.49) and (A.50) provides the important bound we need:

$$\begin{aligned} r_t &\le B_t - B_{t+1} + \frac{\eta}{2}P_t + \frac{\eta}{2}E_t \\ &\le B_t - B_{t+1} + \frac{1}{2}\mathscr{D}_t, \end{aligned} \tag{A.51}$$

where the second inequality holds by the definition of the dissipation in Eq. (A.20) and the nonnegativity of P_{t+1} and E_t:

$$\frac{1}{2}\mathscr{D}_t - \left(\frac{\eta}{2}P_t + \frac{\eta}{2}E_t\right) = \frac{\eta}{2}P_{t+1} + \frac{\eta}{2}E_t \ge 0.$$

The term B_t is uniformly bounded: Since each contrast payoff vector satisfies $\|d_t^x\|_\infty \le \sqrt{1+1/k_x}\,L$ by (A.30), and each off-support probability is at most one, we have

$$|B_{x,t}| \le \sqrt{1+\frac{1}{k_x}}\,q \le \sqrt{2}\,q,$$

and similarly $|B_{y,t}| \le \sqrt{2}\,q$. Hence,

$$|B_t| \le 2\sqrt{2}\,q. \tag{A.52}$$

Remark A.1. Note that the fixed index sets I and J are supports of the Goldman-Tucker saddle point, not of the current iterates. The arguments in this section hold when the iterate coordinates repeatedly leave and reenter the simplex boundary.

A.6 Finite residual budget

By combining the per-step bound Eq. (A.51) with the exact energy balance Eq. (A.21), we can show that the total budget for the residual term is finite, i.e., the residual term is summable. Let

$$R_T = \sum_{t=0}^{T-1} r_t, \qquad \mathcal{D}_T = \sum_{t=0}^{T-1} \mathscr{D}_t$$

denote the accumulated residual and dissipation. Summing up the per-step bound Eq. (A.51) and then applying the bound Eq. (A.52) yields

$$R_T \le B_0 - B_T + \frac{1}{2}\mathcal{D}_T \le 4\sqrt{2}\,q + \frac{1}{2}\mathcal{D}_T. \tag{A.53}$$

The exact energy balance Eq. (A.21), summed over the same range, provides

$$\mathcal{D}_T = R_T - V_T + V_0. \tag{A.54}$$

By substituting Eq. (A.54) into Eq. (A.53) and solving for R_T, we have

$$R_T \le 8\sqrt{2}\, q + V_0 - V_T.$$

Using the energy bounds Eqs. (A.27) and (A.28), we prove

$$\boxed{R_T \le 4 + (2 + 8\sqrt{2})q.} \tag{A.55}$$

This inequality shows that the total contribution of r_t or the "projection collisions" is bounded by an absolute constant, for every horizon T and every initialization (x^0, y^0).

A.7 From finite residual budget to the rate on the averaged gap

It remains to bound the duality gap of the averages of the iterates by the accumulated residual. This can be done by using the standard projection inequality: if $p^+ = \Pi_C(p + \eta g)$ for a closed convex set C, then for every $z \in C$,

$$\eta g^\top (z - p^+) \le \frac{1}{2}\left(\|p - z\|_2^2 - \|p^+ - z\|_2^2 - \|p^+ - p\|_2^2\right). \tag{A.56}$$

Fix arbitrary comparators $x \in \Delta_n$ and $y \in \Delta_m$, and let

$$g_t := y^\top A x^t - (y^t)^\top A x.$$

Averaging g_t over $t = 1, \dots, T$ and taking the supremum over comparators produces $\mathrm{Gap}(\overline{x}^T, \overline{y}^T)$.

Because the two players update in alternation, we use the telescoping potential functions that pair the current x-iterate *once with the current* and *once with the preceding y-iterate*: for $t \ge 1$, define

$$\phi_t = \frac{1}{2}\|x^t - x\|_2^2 + \frac{1}{2}\|y^t - y\|_2^2 + \eta (y^t)^\top A x, \tag{A.57}$$

$$\psi_t = \frac{1}{2}\|x^t - x\|_2^2 + \frac{1}{2}\|y^{t-1} - y\|_2^2 - \frac{1}{2}\|y^t - y^{t-1}\|_2^2, \tag{A.58}$$

$$\Phi_t = \phi_t + \psi_t. \tag{A.59}$$

Applying Eq. (A.56) to the x-update with comparator $x \in \Delta_n$ and to the y-update with comparator $y \in \Delta_m$, and assembling the bilinear terms, it follows that

$$\eta g_{t+1} \le \phi_t - \phi_{t+1} + \eta u_t^\top \Delta y_t - \frac{1}{2}\left(\|\Delta x_t\|_2^2 + \|\Delta y_t\|_2^2\right). \tag{A.60}$$

Instead, pairing the current x-update with the preceding y-update leads to

$$\eta g_t \le \psi_t - \psi_{t+1} + \eta v_t^\top \Delta x_t - \frac{1}{2}\left(\|\Delta x_t\|_2^2 + \|\Delta y_t\|_2^2\right). \tag{A.61}$$

Adding up Eqs. (A.60) and (A.61) yields the residual term in Eq. (A.17):

$$\eta(g_t + g_{t+1}) \le \Phi_t - \Phi_{t+1} + r_t, \qquad t \ge 1. \tag{A.62}$$

For $T \ge 2$, summing (A.62) from $t = 1$ to $T - 1$ counts every interior g_t twice and the endpoints once, so

$$2\eta \sum_{t=1}^{T} g_t \le \Phi_1 - \Phi_T + \sum_{t=1}^{T-1} r_t + \eta(g_1 + g_T). \tag{A.63}$$

The following bounds show that the endpoint quantities are bounded uniformly over the comparators:

$$\Phi_1 \le 4 + q, \qquad -\Phi_T \le q + \frac{q^2}{2}, \qquad \eta(g_1 + g_T) \le 4q, \tag{A.64}$$

where the first bound follows from the simplex squared-diameter bound 2 together with $|(y^1)^\top A x| \le L$, the second inequality holds because the only possibly negative terms in Φ_T: $-\|\Delta y_{T-1}\|_2^2/2$ and $\eta(y^T)^\top A x$ are bounded below by $-q^2/2$ and $-q$ respectively, and the third bound uses $|g_t| \le 2L$.

Because every residual is nonnegative by Lemma A.2, we have

$$\sum_{t=1}^{T-1} r_t \le R_T,$$

and then we can apply the finite budget bound Eq. (A.55). Combining it with the endpoint bounds Eq. (A.64) in Eq. (A.63), and then using $q \le 1/(2\sqrt{2})$ from Eq. (A.26), we obtain

$$\begin{aligned} 2\eta \sum_{t=1}^{T} g_t &\le 8 + (8 + 8\sqrt{2})q + \frac{q^2}{2} \\ &\le 12 + 2\sqrt{2} + \frac{1}{16} = \frac{193}{16} + 2\sqrt{2} < 15. \end{aligned} \tag{A.65}$$

By dividing both sides of Eq. (A.65) by $2\eta T$, taking the supremum over the comparators, and using the bilinearity of the pairing g_t, we have

$$\mathrm{Gap}(\overline{x}^T, \overline{y}^T) \le \frac{15}{2\eta T},$$

which completes Theorem A.1 for $T \ge 2$.

For $T = 1$, the average is just the first iterate, and the direct bound suffices:

$$\mathrm{Gap}(x^1, y^1) \le 2L = \frac{2q}{\eta} \le \frac{1}{\sqrt{2}\,\eta} < \frac{15}{2\eta}.$$

Therefore, Theorem A.1 is proven for every $T \ge 1$.

B Lyapunov-LMI Details

In this appendix we describe the Lyapunov LMI that is at the core of the PEP framework which motivated the convergence analysis of Theorem 1. Also we present a proof of Theorem 2, which is conditional on having a exactly feasible solution to the LMI in consideration. Throughout, the setting consists of compact convex sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ of radius D, a payoff matrix normalized by $\|A\|_2 \le 1$, and AltGDA with projections onto $\mathcal{X}$ and $\mathcal{Y}$.

Local notation. Notation that we use in this appendix is as follows. $\mathbb{S}^d$ (respectively $\mathbb{S}^d_+$) denotes the symmetric (respectively symmetric positive semidefinite) $d \times d$ matrices, and $M \succeq 0$ means $M \in \mathbb{S}^d_+$. The symmetric outer product of $u, v \in \mathbb{R}^d$ is $\boxdot(u, v) = \frac{1}{2}\left(uv^\top + vu^\top\right)$. For a closed convex set C, ι_C is its indicator function, so $\partial\iota_C(z)$ is the normal cone to C at z. Several symbols reuse letters used elsewhere in the paper, always with a new, locally defined meaning. The starred pair $(x^\star, y^\star)$ is an arbitrary comparator in $\mathcal{X} \times \mathcal{Y}$, not the strictly complementary saddle point used to prove Theorem 1. The symbols μ and λ denote LMI variables, *not* the projection multipliers of Appendix A. The vectors $\bar{\delta}, \underline{\delta}$ are projection residuals and they are not related to the separation parameter δ.

B.1 Interpolation results

In this section and the next, we will be using the following interpolation results from Bousselmi et al., 2024 and Taylor et al., 2017a, Theorem 3.6, respectively.

Lemma B.1 (Interpolation of a matrix with bounded singular value, Bousselmi et al., 2024)**.** *Consider the sets of pairs $\{(x^i, p^i)\}_{i\in\{1,\dots,T_1\}} \subseteq \mathbb{R}^n \times \mathbb{R}^m$ and $\{(y^j, q^j)\}_{j\in\{1,\dots,T_2\}} \subseteq \mathbb{R}^m \times \mathbb{R}^n$, and define the matrices $X = [x^1 \mid \cdots \mid x^{T_1}] \in \mathbb{R}^{n\times T_1}$, $P = [p^1 \mid \cdots \mid p^{T_1}] \in \mathbb{R}^{m\times T_1}$, $Y = [y^1 \mid \cdots \mid y^{T_2}] \in \mathbb{R}^{m\times T_2}$, and $Q = [q^1 \mid \cdots \mid q^{T_2}] \in \mathbb{R}^{n\times T_2}$. Then there exists a matrix $A \in \mathbb{R}^{m\times n}$ with maximum singular value $\sigma_{\max}(A) \le L$ such that $p^i = Ax^i$ for all $i \in \{1, \dots, T_1\}$ and $q^j = A^\top y^j$ for all $j \in \{1, \dots, T_2\}$ if and only if*

$$\begin{aligned} &X^\top Q = P^\top Y, \\ &L^2 X^\top X - P^\top P \succeq 0, \\ &L^2 Y^\top Y - Q^\top Q \succeq 0. \end{aligned}$$

Lemma B.2 (Interpolation of a compact convex set with bounded radius, Taylor et al., 2017a, Theorem 3.6)**.** *Let $\mathcal{I}$ be a finite index set and $\{(z^i, g^i)\}_{i\in\mathcal{I}} \subseteq \mathbb{R}^d \times \mathbb{R}^d$. There exists a compact convex set $C \subseteq \mathbb{R}^d$ contained in the ball of radius R about the origin with $z^i \in C$ and $g^i \in \partial\iota_C(z^i)$ for all $i \in \mathcal{I}$ if and only if*

$$(g^j)^\top (z^i - z^j) \le 0, \quad i, j \in \mathcal{I},$$
$$\|z^i\|_2^2 \le R^2, \quad i \in \mathcal{I}.$$

The bounded-radius assumption is not part of the LMI itself; it is used only at the end of the proof to bound the initial endpoint potential.

B.2 Derivation of the Lyapunov LMI

Window notation. Fix a base index k, an integer $h \ge 0$, and a normalized matrix satisfying $\|A\|_2 \le 1$. Let $p^i = Ax^i$ and $q^i = A^\top y^i$. For the projection residuals, write $\bar{\delta}^{i+1} = x^i - \eta q^i - x^{i+1} \in \partial\iota_{\mathcal{X}}(x^{i+1})$ and $\underline{\delta}^{i+1} = y^i + \eta p^{i+1} - y^{i+1} \in \partial\iota_{\mathcal{Y}}(y^{i+1})$, with $\bar{\delta}^0 = 0$ and $\underline{\delta}^0 = 0$. These two zero residuals are padding terms: they make the $k = 0$ window match the shifted-window notation used below. For a comparator $(x^\star, y^\star) \in \mathcal{X} \times \mathcal{Y}$, set $p^\star = Ax^\star$, $q^\star = A^\top y^\star$, and $\bar{\delta}^\star = \underline{\delta}^\star = 0$. Write $I_r^\star(k) = \{\star, k, k+1, \dots, k+r\}$.

Potential states. The numerical certificate is designed to search over Lyapunov matrices $Q^x \in \mathbb{S}^{2h+6}$ and $Q^y \in \mathbb{S}^{2h+7}$; these are required to be symmetric but need not be positive semidefinite a priori. For a fixed comparator, define

$$H_k^x = \left[x^\star \mid q^\star \mid x^k \mid \bar{\delta}^k \mid \cdots \mid \bar{\delta}^{k+h+1} \mid q^k \mid \cdots \mid q^{k+h}\right],$$

$$H_k^y = \left[y^\star \mid p^\star \mid y^k \mid \underline{\delta}^k \mid \cdots \mid \underline{\delta}^{k+h+1} \mid p^k \mid \cdots \mid p^{k+h+1}\right],$$

and define H_{k+1}^x and H_{k+1}^y by shifting the iterate, residual, and gradient indices forward by one. Thus H_k^x and H_{k+1}^x have $2h+6$ columns and H_k^y, H_{k+1}^y have $2h+7$; the shifted pair repeats the column pattern of the unshifted pair with base index $k+1$. The potential is

$$\mathcal{V}_k = \mathbf{tr}\left(Q^x (H_k^x)^\top H_k^x\right) + \mathbf{tr}\left(Q^y (H_k^y)^\top H_k^y\right),$$

written $\mathcal{V}_k(x^\star, y^\star)$ when the dependence on the comparator columns of H_k^x, H_k^y matters.

Full-state selectors. Let $N_x = 2h+8$ and $N_y = 2h+9$. Define the full-state data matrices

$$F_k^x = \left[x^\star \mid q^\star \mid x^k \mid \bar{\delta}^k \mid \cdots \mid \bar{\delta}^{k+h+2} \mid q^k \mid \cdots \mid q^{k+h+1}\right] \in \mathbb{R}^{n \times N_x},$$

and

$$F_k^y = \left[y^\star \mid p^\star \mid y^k \mid \underline{\delta}^k \mid \cdots \mid \underline{\delta}^{k+h+2} \mid p^k \mid \cdots \mid p^{k+h+2}\right] \in \mathbb{R}^{m \times N_y}.$$

Their Gram matrices are $G^x = (F_k^x)^\top F_k^x$ and $G^y = (F_k^y)^\top F_k^y$. Hatted symbols denote the corresponding coordinate selectors in these full states: a hatted selector $\widehat{z} \in \mathbb{R}^{N_x}$ represents the actual vector $F_k^x \widehat{z}$ on the x-side, and analogously on the y-side. For the x-side, set $\widehat{x}^\star = e_1^{N_x}$, $\widehat{q}^\star = e_2^{N_x}$, $\widehat{x}^k = e_3^{N_x}$, $\widehat{\bar{\delta}}^\star = 0^{N_x}$, $\widehat{\bar{\delta}}^{k+i} = e_{4+i}^{N_x}$ for $i \in \{0, \dots, h+2\}$, $\widehat{q}^{k+i} = e_{h+7+i}^{N_x}$ for $i \in \{0, \dots, h+1\}$, and

$$\widehat{x}^{k+i} = \widehat{x}^k - \sum_{t=1}^{i} \widehat{\bar{\delta}}^{k+t} - \eta \sum_{t=0}^{i-1} \widehat{q}^{k+t}, \qquad i \in \{1, \dots, h+2\}.$$

For the y-side, set $\widehat{y}^\star = e_1^{N_y}$, $\widehat{p}^\star = e_2^{N_y}$, $\widehat{y}^k = e_3^{N_y}$, $\widehat{\underline{\delta}}^\star = 0^{N_y}$, $\widehat{\underline{\delta}}^{k+i} = e_{4+i}^{N_y}$ for $i \in \{0, \dots, h+2\}$, $\widehat{p}^{k+i} = e_{h+7+i}^{N_y}$ for $i \in \{0, \dots, h+2\}$, and

$$\widehat{y}^{k+i} = \widehat{y}^k - \sum_{t=1}^{i} \widehat{\underline{\delta}}^{k+t} + \eta \sum_{t=1}^{i} \widehat{p}^{k+t}, \qquad i \in \{1, \dots, h+2\}.$$

Note that, for each potential-state matrix, the hatted version denotes the selector matrix that is obtained columnwise from the corresponding hatted vectors, e.g., $H_k^x = F_k^x \widehat{H}_k^x$ and $H_k^y = F_k^y \widehat{H}_k^y$, with the same convention for the shifted matrices. Let us now define

$$\widehat{X} = \left[\widehat{x}^\star \mid \widehat{x}^k \mid \widehat{x}^{k+1} \mid \cdots \mid \widehat{x}^{k+h+2}\right], \qquad \widehat{P} = \left[\widehat{p}^\star \mid \widehat{p}^k \mid \widehat{p}^{k+1} \mid \cdots \mid \widehat{p}^{k+h+2}\right],$$

$$\widehat{Y} = \left[\widehat{y}^\star \mid \widehat{y}^k \mid \widehat{y}^{k+1} \mid \cdots \mid \widehat{y}^{k+h+1}\right], \qquad \widehat{Q} = \left[\widehat{q}^\star \mid \widehat{q}^k \mid \widehat{q}^{k+1} \mid \cdots \mid \widehat{q}^{k+h+1}\right].$$

The symbol $\widehat{Q}$ here is a selector matrix for the q-vectors and is unrelated to the Lyapunov matrices Q^x and Q^y.

Generic condition. We now build the algebraic certificate from a generic trace condition. For symmetric target matrices $M_x \in \mathbb{S}^{N_x}$ and $M_y \in \mathbb{S}^{N_y}$, we can write the desired trace condition $\mathbf{tr}(M_x G^x) + \mathbf{tr}(M_y G^y) \le 0$ for Gram matrices G^x, G^y associated with an admissible AltGDA window. The worst-case problem computes the largest value of this target over all AltGDA trajectories in the admissible window and all normalized matrices compatible within this window:

$$\mathsf{p}_{\mathrm{op}}(M_x, M_y) := \left(\begin{array}{l} \underset{\substack{F_k^x \in \mathbb{R}^{n\times N_x},\, F_k^y \in \mathbb{R}^{m\times N_y},\\ A\in\mathbb{R}^{m\times n},\, m,n\in\mathbb{N}}}{\text{maximize}} \ \mathbf{tr}\big(M_x (F_k^x)^\top F_k^x\big) + \mathbf{tr}\big(M_y (F_k^y)^\top F_k^y\big) \\ \text{subject to} \\ \|A\|_2 \le 1, \\ F_k^y \widehat{p}^i = A F_k^x \widehat{x}^i, \quad i \in I^\star_{h+2}(k), \\ F_k^x \widehat{q}^j = A^\top F_k^y \widehat{y}^j, \quad j \in I^\star_{h+1}(k), \\ \left\langle F_k^x \widehat{\tilde{\delta}}^j, F_k^x(\widehat{x}^i - \widehat{x}^j) \right\rangle \le 0, \quad i,j \in I^\star_{h+2}(k), \quad j \notin \{i,\star\}, \\ \left\langle F_k^y \widehat{\underline{\delta}}^j, F_k^y(\widehat{y}^i - \widehat{y}^j) \right\rangle \le 0, \quad i,j \in I^\star_{h+2}(k), \quad j \notin \{i,\star\}. \end{array} \right)$$

In the problem above, the last two constraints model convexity of the sets in consideration. The hatted selectors already encode the AltGDA recursions, so we do not need any separate update equations for them. The first two constraint lines impose $p^i = Ax^i$ and $q^j = A^\top y^j$ and the last two are the normal-cone inequalities for the two projections (with zero cases omitted).

By Lemma B.1 (with $L = 1$), the explicit matrix A can be eliminated: the problem is equivalently written with the interpolation constraints $(F_k^x \widehat{X})^\top (F_k^x \widehat{Q}) = (F_k^y \widehat{P})^\top (F_k^y \widehat{Y})$, $(F_k^y \widehat{P})^\top (F_k^y \widehat{P}) \preceq (F_k^x \widehat{X})^\top (F_k^x \widehat{X})$, and $(F_k^x \widehat{Q})^\top (F_k^x \widehat{Q}) \preceq (F_k^y \widehat{Y})^\top (F_k^y \widehat{Y})$ in place of the first three constraint lines. Passing to $G^x = (F_k^x)^\top F_k^x$ and $G^y = (F_k^y)^\top F_k^y$ and dropping the implicit rank restrictions gives the homogeneous Gram relaxation

$$\mathsf{p}(M_x, M_y) := \left(\begin{array}{l} \underset{\substack{G^x \in \mathbb{S}_+^{N_x} \\ G^y \in \mathbb{S}_+^{N_y}}}{\text{maximize}} \ \mathbf{tr}(M_x G^x) + \mathbf{tr}(M_y G^y) \\ \text{subject to} \\ \mathbf{tr}\big(\boxdot(\widehat{x}^i, \widehat{q}^j) G^x\big) - \mathbf{tr}\big(\boxdot(\widehat{p}^i, \widehat{y}^j) G^y\big) = 0, \quad i \in I^\star_{h+2}(k), \quad j \in I^\star_{h+1}(k), \\ \widehat{P}^\top G^y \widehat{P} \preceq \widehat{X}^\top G^x \widehat{X}, \\ \widehat{Q}^\top G^x \widehat{Q} \preceq \widehat{Y}^\top G^y \widehat{Y}, \\ \mathbf{tr}\left(\boxdot(\widehat{\tilde{\delta}}^j, \widehat{x}^i - \widehat{x}^j) G^x \right) \le 0, \quad i,j \in I^\star_{h+2}(k), \quad j \notin \{i,\star\}, \\ \mathbf{tr}\left(\boxdot(\widehat{\underline{\delta}}^j, \widehat{y}^i - \widehat{y}^j) G^y \right) \le 0, \quad i,j \in I^\star_{h+2}(k), \quad j \notin \{i,\star\}, \end{array} \right)$$

where the first three constraints use the Lemma B.1 last two constraints use Lemma B.2. The feasible set of this Gram relaxation contains every Gram pair generated by the preceding vector formulations, so nonpositivity of $\mathsf{p}(M_x, M_y)$ is a sufficient certificate for the generic trace condition.

Compact dual certificate for a generic condition. The maximization problem above is homogeneous, so a zero-objective dual feasible point certifies $\mathsf{p}(M_x, M_y) \le 0$. Concretely, take PSD blocks $U \in \mathbb{S}_+^{h+4}$, $V \in \mathbb{S}_+^{h+3}$, $Z^x \in \mathbb{S}_+^{N_x}$, and $Z^y \in \mathbb{S}_+^{N_y}$; free multipliers $\mu_{i,j}$ for $i \in I^\star_{h+2}(k)$ and

$j \in I^\star_{h+1}(k)$; and nonnegative multipliers $\lambda^x_{i,j}, \lambda^y_{i,j} \geq 0$ for $i, j \in I^\star_{h+2}(k)$ with $j \notin \{i, \star\}$. It is enough to find such data satisfying

$$R^\lambda_x = \sum_{\substack{i,j \in I^\star_{h+2}(k) \\ j \notin \{i,\star\}}} \lambda^x_{i,j} \boxdot \left(\widehat{\delta}^j, \widehat{x}^i - \widehat{x}^j \right), \qquad R^\lambda_y = \sum_{\substack{i,j \in I^\star_{h+2}(k) \\ j \notin \{i,\star\}}} \lambda^y_{i,j} \boxdot \left(\underline{\widehat{\delta}}^j, \widehat{y}^i - \widehat{y}^j \right),$$

$$R^\mu_x = \sum_{\substack{i \in I^\star_{h+2}(k) \\ j \in I^\star_{h+1}(k)}} \mu_{i,j} \boxdot (\widehat{x}^i, \widehat{q}^j), \qquad R^\mu_y = \sum_{\substack{i \in I^\star_{h+2}(k) \\ j \in I^\star_{h+1}(k)}} \mu_{i,j} \boxdot (\widehat{p}^i, \widehat{y}^j),$$

$$-M_x + R^\lambda_x + R^\mu_x - \widehat{X} U \widehat{X}^\top + \widehat{Q} V \widehat{Q}^\top = Z^x,$$

$$-M_y + R^\lambda_y - R^\mu_y + \widehat{P} U \widehat{P}^\top - \widehat{Y} V \widehat{Y}^\top = Z^y.$$

We now take trace inner products of the last two identities with any feasible G^x, G^y. Note that the operator-equality constraints cancel the μ terms, the normal-cone inequalities make the λ terms nonpositive, the two matrix inequalities do the same for the U, V terms, and, finally, the slack terms satisfy $\mathbf{tr}(Z^x G^x) + \mathbf{tr}(Z^y G^y) \geq 0$. As a result, we have $\mathbf{tr}(M_x G^x) + \mathbf{tr}(M_y G^y) \leq 0$ for every feasible G^x, G^y.

Instantiating the three conditions. The full C_1-C_2-C_3 certificate searches over symmetric auxiliary matrices $S^x \in \mathbb{S}^{N_x}$ and $S^y \in \mathbb{S}^{N_y}$. Let us now define $\mathcal{R}^S_k = \mathbf{tr}(S^x G^x) + \mathbf{tr}(S^y G^y)$ and

$$\mathcal{Q}^x_0 = \widehat{H}^x_k Q^x (\widehat{H}^x_k)^\top, \qquad \mathcal{Q}^x_1 = \widehat{H}^x_{k+1} Q^x (\widehat{H}^x_{k+1})^\top,$$
$$\mathcal{Q}^y_0 = \widehat{H}^y_k Q^y (\widehat{H}^y_k)^\top, \qquad \mathcal{Q}^y_1 = \widehat{H}^y_{k+1} Q^y (\widehat{H}^y_{k+1})^\top.$$

The three target pairs can be written as:

$$M^{C_1}_x = \mathcal{Q}^x_1 - \mathcal{Q}^x_0 + S^x, \qquad M^{C_1}_y = \mathcal{Q}^y_1 - \mathcal{Q}^y_0 + S^y,$$
$$M^{C_2}_x = -\mathcal{Q}^x_0, \qquad M^{C_2}_y = -\mathcal{Q}^y_0,$$
$$M^{C_3}_x = \boxdot(\widehat{q}^\star, \widehat{x}^k) - S^x, \qquad M^{C_3}_y = -\boxdot(\widehat{p}^\star, \widehat{y}^k) - S^y.$$

For each condition $r \in \{C_1, C_2, C_3\}$, take an independent copy of the generic dual variables, writing $R^\lambda_{x,r}, R^\lambda_{y,r}, R^\mu_{x,r}, R^\mu_{y,r}$ for the generic sums with $\lambda^x, \lambda^y, \mu$ replaced by $\lambda^{r,x}, \lambda^{r,y}, \mu^r$. The *Lyapunov LMI* is the collection of the six affine identities

$$-M^r_x + R^\lambda_{x,r} + R^\mu_{x,r} - \widehat{X} U^r \widehat{X}^\top + \widehat{Q} V^r \widehat{Q}^\top = Z^{r,x}, \tag{B.1}$$

$$-M^r_y + R^\lambda_{y,r} - R^\mu_{y,r} + \widehat{P} U^r \widehat{P}^\top - \widehat{Y} V^r \widehat{Y}^\top = Z^{r,y}, \tag{B.2}$$

for $r \in \{C_1, C_2, C_3\}$, in the variables

$$Q^x \in \mathbb{S}^{2h+6},\ Q^y \in \mathbb{S}^{2h+7},\ S^x \in \mathbb{S}^{N_x},\ S^y \in \mathbb{S}^{N_y},$$
$$U^r \in \mathbb{S}^{h+4}_+,\ V^r \in \mathbb{S}^{h+3}_+,\ Z^{r,x} \in \mathbb{S}^{N_x}_+,\ Z^{r,y} \in \mathbb{S}^{N_y}_+, \quad r \in \{C_1, C_2, C_3\},$$
$$\mu^r_{i,j} \in \mathbb{R}, \quad i \in I^\star_{h+2}(k),\ j \in I^\star_{h+1}(k),\ r \in \{C_1, C_2, C_3\},$$
$$\lambda^{r,x}_{i,j} \geq 0,\ \lambda^{r,y}_{i,j} \geq 0, \quad i, j \in I^\star_{h+2}(k),\ j \notin \{i, \star\},\ r \in \{C_1, C_2, C_3\}.$$

This is the combined feasibility problem solved in the numerical diagnostics of Section 5.

Use in the proof. The Gram matrices of an actual AltGDA window are feasible for the relaxation above as the normal-cone and operator-consistency constraints hold by construction and the interpolation constraints hold by the necessity direction of Lemma B.1. Thus we can apply the preceding trace argument here to each condition r and its identities (B.1)–(B.2), leading to $\mathbf{tr}(M^r_x G^x) + \mathbf{tr}(M^r_y G^y) \leq 0$ on every actual window. Note that, while we state the identities here for fixed index k, the selector pattern depends only on offsets from k, so a single feasible certificate yields these inequalities for every base window. Finally, the target matrices for the three conditions encode the inequalities

$$\mathcal{V}_{k+1} - \mathcal{V}_k + \mathcal{R}^S_k \leq 0,$$
$$\mathcal{V}_k \geq 0,$$
$$(q^\star)^\top x^k - (p^\star)^\top y^k \leq \mathcal{R}^S_k,$$

which the proof below combines by telescoping.

B.3 Proof of the conditional theorem

Proof of Theorem 2. Let $(x^\star, y^\star) \in \mathcal{X} \times \mathcal{Y}$ be an arbitrary comparator, and set $p^\star = Ax^\star$ and $q^\star = A^\top y^\star$. Let us now fix a base window k. Due to the trace argument of the compact dual-certificate construction applied to the condition-r identities (B.1)–(B.2) on the actual Gram matrices, we have

$$\mathbf{tr}(M_x^r G^x) + \mathbf{tr}(M_y^r G^y) \le 0, \qquad r \in \{C_1, C_2, C_3\}.$$

For $r = C_1$, the definition of M^{C_1} gives $\mathcal{V}_{k+1}(x^\star, y^\star) - \mathcal{V}_k(x^\star, y^\star) + \mathcal{R}_k^S \le 0$. For $r = C_3$, the definition of M^{C_3} gives $(q^\star)^\top x^k - (p^\star)^\top y^k - \mathcal{R}_k^S \le 0$. Combining the two yields the one-step bound

$$(q^\star)^\top x^k - (p^\star)^\top y^k \le \mathcal{V}_k(x^\star, y^\star) - \mathcal{V}_{k+1}(x^\star, y^\star).$$

Summing from $k = 0$ to $T-1$ telescopes the potential. For $\bar{x}_T = \frac{1}{T}\sum_{k=0}^{T-1} x^k$ and $\bar{y}_T = \frac{1}{T}\sum_{k=0}^{T-1} y^k$, the comparator terms satisfy

$$\frac{1}{T}\sum_{k=0}^{T-1}\left((q^\star)^\top x^k - (p^\star)^\top y^k\right) = (y^\star)^\top A\bar{x}_T - \bar{y}_T^\top A x^\star,$$

and therefore

$$T\left[(y^\star)^\top A\bar{x}_T - \bar{y}_T^\top A x^\star\right] \le \mathcal{V}_0(x^\star, y^\star) - \mathcal{V}_T(x^\star, y^\star).$$

The C_2 condition applied to the terminal window gives $-\mathcal{V}_T(x^\star, y^\star) \le 0$, so

$$T\left[(y^\star)^\top A\bar{x}_T - \bar{y}_T^\top A x^\star\right] \le \mathcal{V}_0(x^\star, y^\star).$$

For any symmetric matrix Q and any matrix H, $\mathbf{tr}(QH^\top H) \le \max\{\lambda_{\max}(Q), 0\}\|H\|_F^2$. Thus

$$\mathcal{V}_0(x^\star, y^\star) \le \varsigma_x^+ \|H_0^x\|_F^2 + \varsigma_y^+ \|H_0^y\|_F^2.$$

The radius assumptions and $\|A\|_2 \le 1$ give $\|p^i\| \le D$ and $\|q^i\| \le D$. Moreover, each nonzero projection residual in the initial window has norm at most ηD, because

$$\|\bar{\delta}^{i+1}\| \le \mathbf{dist}(x^i - \eta q^i, \mathcal{X}) \le \eta\|q^i\| \le \eta D,$$

and similarly $\|\underline{\delta}^{i+1}\| \le \eta D$. The residual-padding convention sets $\bar{\delta}^0 = \underline{\delta}^0 = 0$, so only $h+1$ residual columns contribute on each side. Hence

$$\|H_0^x\|_F^2 \le (h+4)D^2 + (h+1)\eta^2 D^2, \qquad \|H_0^y\|_F^2 \le (h+5)D^2 + (h+1)\eta^2 D^2.$$

Substituting these bounds gives

$$\mathcal{V}_0(x^\star, y^\star) \le D^2\left((h+4)\varsigma_x^+ + (h+5)\varsigma_y^+ + (h+1)\eta^2(\varsigma_x^+ + \varsigma_y^+)\right),$$

and the supremum over comparators completes the proof. □

B.4 Computational setup

For each tested memory length $h \in \{2, 3, 4, 5\}$ and stepsize η, we modeled the final C_1, C_2, C_3 feasibility problem of Section B.2 using the JuMP package (Lubin et al., 2023) in Julia and then solved it using MOSEK at its default tolerances. After each LMI numerical solve in Julia, we also independently computed three diagnostic terms representing quality of the solution: the maximum affine-identity residual over (B.1)–(B.2), the minimum eigenvalue over all condition-specific PSD blocks $U^r, V^r, Z^{r,x}, Z^{r,y}$, and the minimum of the multipliers constrained to be nonnegative. A numerical solution is accepted if acceptance conditions described in Section 5 are satisfied.

C Computing the Last-Iterate Duality Gap for AltGDA via PEP

The finite-horizon PEP formulation for AltGDA presented in this section follows a framework similar to that of Nan et al. (2026) but with the last-iterate gap as its objective. For completeness, we provide the full formulation here, prove the exactness of the upper bound using the counterexample described in Example 1 over compact convex sets, and then present simplex counterexample described in Example 2.

Notation. The spaces $\mathbb{S}^d$ and $\mathbb{S}^d_+$ contain symmetric and symmetric positive semidefinite $d \times d$ matrices, respectively; $M \succeq 0$ means $M \in \mathbb{S}^d_+$. For $u, v \in \mathbb{R}^d$, set $\boxdot(u, v) = \frac{1}{2}(uv^\top + vu^\top)$, so that $\mathbf{tr}(G \boxdot (u, v)) = u^\top G v$ for symmetric G. For a closed convex set C, ι_C is its indicator and $\partial\iota_C(z)$ its normal cone at $z \in C$.

The feasible sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$ are compact convex sets contained in the Euclidean balls of radii R_x and R_y about the origin, and $\|A\|_2 \leq L$. In our numerical experiments, we set $L = R_x = R_y = 1$.

As in Appendix B, $p^i = Ax^i$ and $q^i = A^\top y^i$ are operator images. The pair $(x^\star, y^\star)$ is an arbitrary comparator in $\mathcal{X} \times \mathcal{Y}$, not the Goldman–Tucker saddle point of Theorem 1. The full-horizon Gram blocks G^x, G^y are distinct from both the window blocks of Appendix B and the reflected metrics G_x, G_y of Appendix A.

The sampled-data index sets are

$$\mathcal{J}_T = \{-3, -2, -1, 0, 1, \ldots, T\}, \qquad \mathcal{I}_T = \mathcal{J}_T \setminus \{-1\},$$

and they differ from the supports I, J of Section 3. Indices $1, \ldots, T$ label the iterates, 0 the initialization, -2 the comparator $(x^\star, y^\star)$, and -3 an index-weighted average. Index -1 is reserved for the all-ones direction; $e_j \in \mathbb{R}^{2T+11}$ denotes a standard basis vector.

The code labels the player-side columns (x, v) and (u, y). In the paper's notation, the code variable u is the maximizer iterate y, the code variable $y = Ax$ is the image p, and $v = A^\top u$ is q.

C.1 The finite-horizon worst-case problem

The modeled class allows arbitrary m, n, compact convex $\mathcal{X}, \mathcal{Y}$ with the stated radii, and $\|A\|_2 \leq L$. The AltGDA iterates satisfy (AltGDA). With $L = R_x = R_y = 1$, this class contains every matrix-game specialization of (SPP) with $\|A\|_2 \leq 1$, because each simplex is a compact convex subset of the unit ball. We provide the exact compact-convex worst case and the independent simplex example in Section C.8. The horizon-T last-iterate gap is

$$\mathcal{P}_T(\eta) = \left(\begin{array}{l} \underset{\substack{\{(x^t, y^t)\}_{0 \leq t \leq T},\ (x^\star, y^\star),\\ \mathcal{X} \subseteq \mathbb{R}^n,\ \mathcal{Y} \subseteq \mathbb{R}^m,\\ A \in \mathbb{R}^{m \times n},\ m, n \in \mathbb{N}}}{\text{maximize}} \quad (y^\star)^\top A x^T - (y^T)^\top A x^\star \\ \text{subject to} \\ \mathcal{X} \text{ compact convex with radius } R_x, \quad \mathcal{Y} \text{ compact convex with radius } R_y, \\ \|A\|_2 \leq L, \\ \{(x^t, y^t)\}_{1 \leq t \leq T} \text{ generated by AltGDA with stepsize } \eta \text{ from } (x^0, y^0), \\ (x^0, y^0) \in \mathcal{X} \times \mathcal{Y}, \quad (x^\star, y^\star) \in \mathcal{X} \times \mathcal{Y}. \end{array} \right) \tag{C.1}$$

Note that the objective is the duality gap of the last-iterates (x^T, y^T), as in (1). The problem is infinite-dimensional at this stage due to the constraints associated with the sets, matrix, and dimensions.

C.2 Span form of AltGDA

The trajectory constraint of (C.1) has an equivalent description in which the projections become normal-cone elements.

Lemma C.1 (Span form of AltGDA)**.** *Fix $\eta > 0$, closed convex sets $\mathcal{X} \subseteq \mathbb{R}^n$ and $\mathcal{Y} \subseteq \mathbb{R}^m$, a matrix $A \in \mathbb{R}^{m \times n}$, and $(x^0, y^0) \in \mathcal{X} \times \mathcal{Y}$. A sequence $\{(x^t, y^t)\}_{1 \leq t \leq T}$ satisfies $x^t = \Pi_{\mathcal{X}}(x^{t-1} - \eta A^\top y^{t-1})$ and $y^t = \Pi_{\mathcal{Y}}(y^{t-1} + \eta A x^t)$ for $t \in \{1, \ldots, T\}$ if and only if there exist $f^t \in \partial\iota_{\mathcal{X}}(x^t)$ and $h^t \in \partial\iota_{\mathcal{Y}}(y^t)$ for $t \in \{1, \ldots, T\}$ such that, with $p^t = Ax^t$ and $q^t = A^\top y^t$,*

$$x^t = x^0 - \eta \sum_{j=1}^{t} f^j - \eta \sum_{j=0}^{t-1} q^j, \qquad y^t = y^0 - \eta \sum_{j=1}^{t} h^j + \eta \sum_{j=1}^{t} p^j, \qquad t \in \{1, \ldots, T\}. \tag{C.2}$$

Proof. For a closed convex set C, a point $z \in C$, and any $w \in \mathbb{R}^d$, $z = \Pi_C(w)$ if and only if $w - z \in \partial\iota_C(z)$ (Bauschke & Combettes, 2017, Proposition 6.47). The x-update is therefore

equivalent to $x^{t-1} - \eta q^{t-1} - x^t \in \partial \iota_{\mathcal{X}}(x^t)$. Because the normal cone is a cone and $\eta > 0$, this membership holds exactly when the residual equals ηf^t for some $f^t \in \partial \iota_{\mathcal{X}}(x^t)$, or equivalently, $x^t = x^{t-1} - \eta q^{t-1} - \eta f^t$. Likewise, the y-update, which uses $p^t = Ax^t$, is equivalent to $y^t = y^{t-1} + \eta p^t - \eta h^t$ for some $h^t \in \partial \iota_{\mathcal{Y}}(y^t)$. By summing these one-step identities, we have (C.2). On the other hand, by differencing (C.2) we recover the one-step identities, and using the same characterization we can convert them back into the two projection updates. □

C.3 Interpolation argument

The formulation accesses $\mathcal{X}, \mathcal{Y}$, and A only through finitely many projections and matrix–vector products. Using Lemma C.1, Lemma B.2, and Lemma B.1, we can characterize exactly which finite collections of first-order data are realizable.

At the indices $\mathcal{I}_T$, the formulation samples the iterates, initialization, comparator $(x^{-2}, y^{-2}) = (x^\star, y^\star)$, and index-weighted averages

$$x^{-3} = \frac{\sum_{j=1}^T j\, x^j}{\sum_{j=1}^T j}, \qquad y^{-3} = \frac{\sum_{j=1}^T j\, y^j}{\sum_{j=1}^T j}, \tag{C.3}$$

which lie in $\mathcal{X} \times \mathcal{Y}$ by convexity. Each sample has normal-cone data f^i, h^i and images $p^i = Ax^i$, $q^i = A^\top y^i$. Choosing zero normal-cone elements makes the samples at $i \in \{-3, -2, 0\}$ valid. The reserved index -1 adds one formal sample per side only for the operator families.

C.4 Finite-dimensional reformulation

Collect the sampled points and images over the full index set, with columns ordered $-3, -2, -1, 0, 1, \ldots, T$:

$$X = [x^i]_{i \in \mathcal{J}_T} \in \mathbb{R}^{n \times (T+4)}, \qquad P = [p^i]_{i \in \mathcal{J}_T} \in \mathbb{R}^{m \times (T+4)},$$
$$Y = [y^i]_{i \in \mathcal{J}_T} \in \mathbb{R}^{m \times (T+4)}, \qquad Q = [q^i]_{i \in \mathcal{J}_T} \in \mathbb{R}^{n \times (T+4)}.$$

In the sampled data the objective reads $(q^{-2})^\top x^T - (y^T)^\top p^{-2}$. Replacing the trajectory and class constraints of (C.1) by the sampled conditions produces a finite-dimensional problem with the same optimal value:

$$\mathcal{P}_T(\eta) = \left(\begin{array}{l} \underset{\substack{\{(x^i, f^i, q^i)\}_{i \in \mathcal{J}_T} \subseteq \mathbb{R}^n, \\ \{(y^i, h^i, p^i)\}_{i \in \mathcal{J}_T} \subseteq \mathbb{R}^m, \\ m, n \in \mathbb{N}}}{\text{maximize}} \; (q^{-2})^\top x^T - (y^T)^\top p^{-2} \\ \text{subject to} \\ (f^j)^\top (x^i - x^j) \le 0, \quad (h^j)^\top (y^i - y^j) \le 0, \quad i, j \in \mathcal{I}_T,\ i \ne j, \\ \|x^i\|_2^2 \le R_x^2, \quad \|y^i\|_2^2 \le R_y^2, \quad i \in \mathcal{I}_T, \\ x^i = x^0 - \eta \sum_{j=1}^i f^j - \eta \sum_{j=0}^{i-1} q^j, \quad i \in \{1, \ldots, T\}, \\ y^i = y^0 - \eta \sum_{j=1}^i h^j + \eta \sum_{j=1}^i p^j, \quad i \in \{1, \ldots, T\}, \\ x^{-3} \text{ and } y^{-3} \text{ satisfy (C.3)}, \\ (x^i)^\top q^j = (p^i)^\top y^j, \quad i, j \in \mathcal{J}_T, \\ L^2 X^\top X - P^\top P \succeq 0, \qquad L^2 Y^\top Y - Q^\top Q \succeq 0. \end{array} \right) \tag{C.4}$$

Every feasible point of (C.1) corresponds to a feasible point of (C.4) with the same objective value. The constraints associated with AltGDA updates hold with the normal-cone elements supplied by Lemma C.1, the constraints modeling convex compact sets hold by Lemma B.2, and the operator constraint associated with modeling the matrix A with $\|A\|_2 \le L$ holds by Lemma B.1. The samples at the reserved index -1 may be set to zero without violating any displayed constraint.

Conversely, consider any feasible point of (C.4). By Lemma B.2, there are compact convex sets $\mathcal{X}$ and $\mathcal{Y}$ of radii R_x and R_y that contain the sampled points with the sampled normal-cone elements. By Lemma B.1, there is a matrix A with $\|A\|_2 \le L$ that matches every image pair, including the samples at indices -1 and -3. Finally, Lemma C.1 identifies the recursions with the AltGDA projections while preserving the objective value.

Applying A and $A^\top$ to the affine combinations (C.3) shows that the realized data satisfy $p^{-3} = \sum_{j=1}^T j\, p^j / \sum_{j=1}^T j$ and $q^{-3} = \sum_{j=1}^T j\, q^j / \sum_{j=1}^T j$, so the averaged sample is the weighted average of the corresponding trajectory. Thus, the two optimal values are the same, and the equality in (C.4) is exact.

C.5 Gram lift and the solved semidefinite program

Stack the variables of (C.4) into one data matrix per player side,

$$\begin{aligned} F^x &= \left[x^{-2} \mid x^{-1} \mid x^0 \mid f^{-3} \mid \cdots \mid f^T \mid q^{-3} \mid \cdots \mid q^T\right] \in \mathbb{R}^{n\times(2T+11)}, \\ F^y &= \left[y^{-2} \mid y^{-1} \mid y^0 \mid h^{-3} \mid \cdots \mid h^T \mid p^{-3} \mid \cdots \mid p^T\right] \in \mathbb{R}^{m\times(2T+11)}, \end{aligned} \tag{C.5}$$

in which the normal-cone block and the image block each run over $\mathcal{J}_T$ in the order $-3,-2,-1,0,1,\ldots,T$: three point columns, $T+4$ normal-cone columns, and $T+4$ image columns per side. The Gram blocks are $G^x = (F^x)^\top F^x$ and $G^y = (F^y)^\top F^y$, both in $\mathbb{S}_+^{2T+11}$. The Gram matrices G^x and G^y arising this way have ranks at most n and m, respectively. However, because in our worst-case optimization problem, we are also optimizing over m, n, we can drop the rank constraint without changing the optimal objective value. Hence, the Gram description and optimizing over m, n allows us to remove the ambient dimension m, n without changing the optimal value.

By the hatted symbols we denote coordinate selectors, where a selector $\widehat{z} \in \mathbb{R}^{2T+11}$ represents $F^x\widehat{z}$ on the x-side and $F^y\widehat{z}$ on the y-side. The columns of (C.5) use basis selectors, while the remaining points use affine selectors that encode the recursion and averaging:

$$\begin{aligned} &\widehat{x}^{-2} = e_1, \quad \widehat{x}^{-1} = e_2, \quad \widehat{x}^0 = e_3, \quad \widehat{f}^i = e_{7+i}, \quad \widehat{q}^i = e_{T+11+i}, \quad i \in \mathcal{J}_T, \\ &\widehat{x}^i = \widehat{x}^0 - \eta\sum_{j=1}^{i}\widehat{f}^j - \eta\sum_{j=0}^{i-1}\widehat{q}^j, \quad i \in \{1,\ldots,T\}, \qquad \widehat{x}^{-3} = \frac{\sum_{j=1}^T j\,\widehat{x}^j}{\sum_{j=1}^T j}, \end{aligned} \tag{C.6}$$

$$\begin{aligned} &\widehat{y}^{-2} = e_1, \quad \widehat{y}^{-1} = e_2, \quad \widehat{y}^0 = e_3, \quad \widehat{h}^i = e_{7+i}, \quad \widehat{p}^i = e_{T+11+i}, \quad i \in \mathcal{J}_T, \\ &\widehat{y}^i = \widehat{y}^0 - \eta\sum_{j=1}^{i}\widehat{h}^j + \eta\sum_{j=1}^{i}\widehat{p}^j, \quad i \in \{1,\ldots,T\}, \qquad \widehat{y}^{-3} = \frac{\sum_{j=1}^T j\,\widehat{y}^j}{\sum_{j=1}^T j}. \end{aligned} \tag{C.7}$$

When the recursion and averaging constraints of (C.4) hold, $x^i = F^x\widehat{x}^i$, $f^i = F^x\widehat{f}^i$, and $q^i = F^x\widehat{q}^i$ for every $i \in \mathcal{J}_T$, and likewise $y^i = F^y\widehat{y}^i$, $h^i = F^y\widehat{h}^i$, and $p^i = F^y\widehat{p}^i$; every constraint and objective of (C.4) then becomes a trace or semidefinite condition on (G^x, G^y). Define the selector matrices with columns ordered $-3,-2,-1,0,1,\ldots,T$:

$$\widehat{X} = [\widehat{x}^i]_{i\in\mathcal{J}_T}, \quad \widehat{P} = [\widehat{p}^i]_{i\in\mathcal{J}_T}, \quad \widehat{Y} = [\widehat{y}^i]_{i\in\mathcal{J}_T}, \quad \widehat{Q} = [\widehat{q}^i]_{i\in\mathcal{J}_T},$$

all in $\mathbb{R}^{(2T+11)\times(T+4)}$, so that $X = F^x\widehat{X}$, $Q = F^x\widehat{Q}$, $Y = F^y\widehat{Y}$, and $P = F^y\widehat{P}$. The resulting semidefinite program is the one solved in the study, with the objective $\mathcal{M}_G^{\mathrm{lst}}$ written out in Section C.6:

$$\mathcal{P}_T(\eta) = \left(\begin{array}{l} \underset{G^x,\, G^y\in\mathbb{S}^{2T+11}}{\text{maximize}} \ \mathcal{M}_G^{\mathrm{lst}} \\ \text{subject to} \\ \mathbf{tr}\Big(G^x \boxdot (\widehat{f}^j, \widehat{x}^i - \widehat{x}^j)\Big) \le 0, \quad i,j \in \mathcal{I}_T,\ i \ne j, \\ \mathbf{tr}\Big(G^y \boxdot (\widehat{h}^j, \widehat{y}^i - \widehat{y}^j)\Big) \le 0, \quad i,j \in \mathcal{I}_T,\ i \ne j, \\ \mathbf{tr}\big(G^x \boxdot (\widehat{x}^i, \widehat{q}^j)\big) = \mathbf{tr}\big(G^y \boxdot (\widehat{p}^i, \widehat{y}^j)\big), \quad i,j \in \mathcal{J}_T, \\ L^2\widehat{X}^\top G^x\widehat{X} - \widehat{P}^\top G^y\widehat{P} \succeq 0, \\ L^2\widehat{Y}^\top G^y\widehat{Y} - \widehat{Q}^\top G^x\widehat{Q} \succeq 0, \\ \mathbf{tr}\big(G^x \boxdot (\widehat{x}^i, \widehat{x}^i)\big) \le R_x^2, \quad i \in \mathcal{I}_T, \\ \mathbf{tr}\big(G^y \boxdot (\widehat{y}^i, \widehat{y}^i)\big) \le R_y^2, \quad i \in \mathcal{I}_T, \\ G^x \succeq 0, \qquad G^y \succeq 0 \end{array}\right). \tag{C.8}$$

Equivalence between (C.1) **and** (C.8) Now we argue why (C.1) and (C.8) are equivalent to each other with same optimal objective value. First note that every feasible point of (C.4) maps to the pair $G^x = (F^x)^\top F^x$, $G^y = (F^y)^\top F^y$, which is feasible for (C.8) with the same objective value; the two objectives correspond through the operator equality at $(i, j) = (T, -2)$.

On the other hand, factorizing a feasible pair of (C.8) produces interpolation data that satisfy every constraint of (C.4). To see that, first note that the algorithm update equations and averaging constraints hold because the point selectors are the stated affine combinations, and the remaining families are the displayed trace and semidefinite conditions. As a result, the optimal value of (C.8) is equal to the modeled worst case at every horizon. Hence, (C.8) is equivalent to (C.1) and our reformulation is tight.

Finally, we make some comments regarding the structure and patterns of (C.8). Note that neither m nor n appears in (C.8). Also, the normal-cone and radius families run over $\mathcal{I}_T$, and the three operator families run over all of $\mathcal{J}_T$, including the reserved index. On each player side, the optimization problem has $(T+3)(T+2)$ normal-cone inequalities and $T+3$ radius bounds; the sides share $(T+4)^2$ equalities, two semidefinite constraints of order $T+4$, and the two Gram blocks of order $2T+11$, which is 71 at the largest horizon $T = 30$. The averaging of the iterates is enforced directly through the selectors: $\widehat{x}^{-3}$ and $\widehat{y}^{-3}$ are affine combinations of iterate selectors, while $\widehat{p}^{-3}$ and $\widehat{q}^{-3}$ are basis columns. Finally, we note that $p^{-3} = \sum_{j=1}^T j\, p^j / \sum_{j=1}^T j$ and $q^{-3} = \sum_{j=1}^T j\, q^j / \sum_{j=1}^T j$ need no constraint of their own as the matrix supplied by Lemma B.1 enforces those conditions directly.

C.6 The performance measure

The measure enters (C.8) through G^y:

$$\mathcal{M}_G^{\mathrm{lst}} = \mathbf{tr}\big(G^y \boxdot (\widehat{p}^T, \widehat{y}^{-2})\big) - \mathbf{tr}\big(G^y \boxdot (\widehat{p}^{-2}, \widehat{y}^T)\big). \tag{C.9}$$

At Gram pairs arising from data as in (C.5), this evaluates to $(y^\star)^\top A x^T - (y^T)^\top A x^\star$, the objective of (C.1). The exact optimal value under this objective is $\mathcal{P}_T(\eta)$, which equals 2 when $L = R_x = R_y = 1$ by Example 1. The numerical values in Fig. 2 confirm this value under the solver tolerances in Section C.7.

C.7 Numerical PEP Setup

With $L = R_x = R_y = 1$, the PEP search evaluates each horizon $T = 5, \ldots, 30$ at 25 logarithmically spaced stepsizes $\eta_c = 1/(\eta L)$ in $[0.5, 64]$, corresponding to $\eta \in [1/64, 2]$.

For every numerical experiment for this section, we have used MOSEK 11.0.29 through JuMP Lubin et al. (2023) with `MSK_DPAR_INTPNT_CO_TOL_PFEAS` = `MSK_DPAR_INTPNT_CO_TOL_DFEAS` $= 10^{-4}$ and automatic thread selection. A solve is accepted only with termination status `OPTIMAL` and primal and dual solution statuses `FEASIBLE_POINT`; no solve in our numerical study was rejected.

C.8 Worst-case examples for last-iterate duality gap

Proof of Example 1. For any feasible instance of (C.1), any comparator point $(x^\star, y^\star)$, and any $T \geq 0$, the norm and radius bounds give us:

$$(y^\star)^\top A x^T - (y^T)^\top A x^\star \leq \|y^\star\|\|A\|_2\|x^T\| + \|y^T\|\|A\|_2\|x^\star\| \leq 2.$$

Taking the maximum over comparators and instances proves the upper bound. To attain this upper bound, let $A = I_2$ and let both feasible sets be the closed Euclidean unit ball in $\mathbb{R}^2$. Then their duality gap satisfies

$$\mathrm{Gap}(x, y) = \max_{\|v\| \leq 1} v^\top x - \min_{\|u\| \leq 1} y^\top u = \|x\| + \|y\|.$$

Now consider $0 < \eta \leq 2$, and initialize

$$y^0 = (1, 0), \text{ and } x^0 = \left(\frac{\eta}{2}, \sqrt{1 - \frac{\eta^2}{4}}\right).$$

These points satisfy $\|x^0\| = \|y^0\| = 1$ and $\langle x^0, y^0\rangle = \eta/2$. Suppose now the same identities hold at an iterate (x, y), and form the unprojected candidates $x' = x - \eta y$ and $y' = y + \eta x'$. Direct calculation gives us:

$$\|x'\|^2 = 1 - 2\eta(\eta/2) + \eta^2 = 1, \qquad \langle x', y\rangle = -\eta/2,$$
$$\|y'\|^2 = 1 + 2\eta(-\eta/2) + \eta^2 = 1, \quad \langle x', y'\rangle = -\eta/2 + \eta = \eta/2.$$

Both candidates lie in their unit balls, so the projections leave them unchanged. Using induction we can easily prove the identities and hence gap 2 at every iterate, including the endpoint $\eta = 2$.

Now consider the stepsize range $\eta > 2$, and use the same matrix and sets with $x^0 = y^0 = (1, 0)$. If $x^t = y^t = e$, where $e \in \{(1, 0), (-1, 0)\}$, the first unprojected candidate is $(1 - \eta)e$, whose projection is $-e$. The second candidate is then also $(1 - \eta)e$, with projection $-e$. Thus both players are essentially alternating between $(1, 0)$ and $(-1, 0)$, retaining gap 2. These attain the upper bound at initialization and every subsequent horizon. The first initialization depends on η, as allowed in (C.1), while both constructions are using fixed two-dimensional sets and a fixed matrix. Therefore $\mathcal{P}_T(\eta) = 2$ for every $\eta > 0$ and $T \geq 1$, and $\inf_{\eta>0} \mathcal{P}_T(\eta) = 2$. □

Proof of Example 2. Write the simplex iterates as

$$x^t = (1/2 + a_t, 1/2 - a_t), \qquad y^t = (1/2 + b_t, 1/2 - b_t),$$

where $a_t, b_t \in [-1/2, 1/2]$. The prescribed initialization gives $a_0 = 1/4$, $b_0 = 0$ for every stepsize. Since $A^\top y^t = (2b_t, -2b_t)$ and $Ax^{t+1} = (2a_{t+1}, -2a_{t+1})$, the updates reduce to

$$a_{t+1} = \mathrm{clip}(a_t - 2\eta b_t), \qquad b_{t+1} = \mathrm{clip}(b_t + 2\eta a_{t+1}),$$

where $\mathrm{clip}(u) = \min\{1/2, \max\{-1/2, u\}\}$. Note that each unprojected vector has coordinate sum one. As a result the simplex projection clips the corresponding centered coordinate to $[-1/2, 1/2]$. Also, maximizing and minimizing the payoff coordinates gives $\mathrm{Gap}(x^t, y^t) = 2(|a_t| + |b_t|)$.

First, let us consider $0 < \eta < 1$, and define $E_t = a_t^2 + b_t^2 - 2\eta a_t b_t$. If neither update is clipped, substitution in the unprojected recurrence gives $E_{t+1} = E_t$. If at least one update is clipped, the new point has $|a_{t+1}| = 1/2$ or $|b_{t+1}| = 1/2$. Completing the square gives

$$E_{t+1} = (b_{t+1} - \eta a_{t+1})^2 + (1 - \eta^2)a_{t+1}^2 \geq \frac{1 - \eta^2}{4}$$

when $|a_{t+1}| = 1/2$; interchanging the two coordinates gives the same bound when $|b_{t+1}| = 1/2$. Starting from $E_0 = 1/16$, induction therefore yields

$$E_t \geq \min\left\{\frac{1}{16}, \frac{1 - \eta^2}{4}\right\} \qquad (t \geq 0).$$

Since $\eta < 1$, we also have $E_t \leq (|a_t| + |b_t|)^2$. Thus

$$\mathrm{Gap}(x^t, y^t) \geq 2\sqrt{E_t} \geq \min\left\{\frac{1}{2}, \sqrt{1 - \eta^2}\right\} > 0 \qquad (t \geq 0).$$

For $\eta \geq 1$, the same initialization produces

$$(a_0, b_0) = (1/4, 0), \qquad (a_1, b_1) = (1/4, 1/2), \qquad (a_2, b_2) = (-1/2, -1/2).$$

The clipped recurrence then alternates between $(-1/2, -1/2)$ and $(1/2, 1/2)$. The gaps are consequently $1/2$ at $t = 0$, $3/2$ at $t = 1$, and 2 for every $t \geq 2$. These identities include the endpoint $\eta = 1$. □

Remark C.1 (Averaged convergence on the same trajectory). Note that for every fixed $0 < \eta \leq 1/2$, the original initialization $x^0 = (3/4, 1/4)$, $y^0 = (1/2, 1/2)$ yields last-iterate gaps at least $1/2$ implying non-convergence, while its uniform averages converge. To see the latter directly, we first define the unprojected recurrence $a_{t+1} = a_t - 2\eta b_t$, $b_{t+1} = b_t + 2\eta a_{t+1}$ from $a_0 = 1/4$, $b_0 = 0$. It preserves $E_t = 1/16$, and $E_t \geq (1 - \eta)(a_t^2 + b_t^2)$ gives

$$|a_t|, |b_t| \leq \frac{1}{4\sqrt{1 - \eta}} \leq \frac{1}{2\sqrt{2}} < \frac{1}{2}.$$

Using induction we can thus show that every projection leaves its candidate unchanged thus corresponding to an admissible AltGDA trajectory. Telescoping the two scalar updates over the averaging indices gives us:

$$\frac{1}{T}\sum_{t=1}^{T} a_t = \frac{b_T - b_0}{2\eta T}, \qquad \frac{1}{T}\sum_{t=1}^{T} b_t = \frac{a_1 - a_{T+1}}{2\eta T}.$$

Using $b_0 = 0$, $a_1 = 1/4$, and the coordinate bounds gives us:

$$\mathrm{Gap}(\overline{x}^T, \overline{y}^T) = \frac{|b_T - b_0| + |a_1 - a_{T+1}|}{\eta T} \leq \frac{5}{4\eta T}.$$

The last-iterate lower bound $1/2$ is uniform in both the horizon and $\eta \in (0, 1/2]$, with the matrix and initialization fixed. It therefore also applies when a different constant stepsize in this range is selected for each horizon. The matrix has eigenvalues 0 and 2, so $L = 2$, and its unique saddle point $x^\star = y^\star = (1/2, 1/2)$ has full supports, giving $\delta = 1/2$ in (4). The general ergodic theorem, Theorem 1, also applies for $0 < \eta \leq 1/(8\sqrt{2})$ and the larger range $(0, 1/2]$ proved here is specific to this example.