NORM SATURATION EXPLAINS LAST-ITERATE CONVERGENCE IN REGRET MATCHING

Anonymous authors
Paper under double-blind review

ABSTRACT

Regret matching is the update at the core of counterfactual regret minimization; its guarantees concern the average of all strategies played, and for regret matching the last iterate may not converge at all. Scale-invariant extragradient regret matching (IREG-PRM$^+$) normalizes the cumulative regret vector by its own norm and attains optimal regret without knowledge of the payoff scale. Run unmodified on zero-sum matrix games, it converges linearly in the last iterate, and no analysis explains why. The obstacle is that the algorithm has no fixed step size to analyze: the step size is a state variable, the inverse of a regret norm that the trajectory itself moves. Every proved linear rate for regret-matching dynamics comes from restarting or modifying the update. We identify the mechanism as *norm saturation*: the regret norm rises to a finite limit and freezes the step size. We prove that it always does, with an explicit bound, and that saturation forces the last-iterate Nash gap to vanish on every matrix game; pointwise convergence follows whenever the equilibrium is unique. Near a unique strictly complementary equilibrium the active support freezes in one step, and the one-round Jacobian on that support has a closed form. The last-iterate then converges linearly at a closed-form rate, provided one scale-invariant quantity stays below one: the saturated step size times the largest singular value of the value-centered payoff submatrix on the support. On the 216-instance testbed, the 184 instances with a resolvable limit all satisfy it. The same analysis gives a ratio certificate: observable norm-increment ratios bound the unobservable Nash-gap ratio up to a constant that enters once and does not accumulate with the iteration count. Its slope-two law holds on 96.1% of the instances where the slope is measurable, and the same increment monitors progress in extensive-form games, where best-response passes can be scheduled sparsely.

1 INTRODUCTION

Regret matching is the update inside counterfactual regret minimization, the method that solved heads-up limit hold'em, which powers the strongest poker agents (Tammelin, 2014; Bowling et al., 2015; Brown & Sandholm, 2019). It is parameter-free, and its theory rests on no-regret learning: it connects to follow-the-regularized-leader and Blackwell approachability (Zinkevich, 2003; McMahan, 2011; Shalev-Shwartz, 2011; Abernethy et al., 2011), its averaged iterates reach the game value at a predictable rate (Hart & Mas-Colell, 2000; Zinkevich et al., 2007; Greenwald et al., 2006; Marden et al., 2007), and its regret guarantee holds under full-information feedback.

Every one of those guarantees is about the *average* of the strategies played so far. In a large game the average is a second copy of the strategy that has to be stored and updated. The last-iterate is a separate problem (Lee et al., 2021): on a 3×3 matrix game, regret matching$^+$ and its predictive variant do not converge (Cai et al., 2025). The linear last-iterate rates proved for regret-matching-type dynamics come from restarting the algorithm or modifying its update (Cai et al., 2025; Meng et al., 2025; Perolat et al., 2021; Liu et al., 2023). A lower bound shows that an algorithm which never forgets its past regrets admits no game-independent last-iterate rate (Cai et al., 2024), so an explanation has to use game-dependent structure.

IREG-PRM$^+$ normalizes the cumulative regret vector by its own norm and attains optimal regret without knowledge of the payoff scale (Zhang et al., 2025a). Run as stated in the extra-gradient setup on zero-sum matrix games, it converges linearly in the last iterate, an observation its authors recorded

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.

and left unexplained. The algorithm keeps its full regret memory and has no restart, no smoothing and no step-size parameter. Three features of the update block the standard arguments. The step size is a state variable: the inverse of a regret norm that the trajectory itself moves, while contraction arguments need a step size fixed in advance. The update clips at zero, so the one-round map is not smooth. The fixed points form a cone, one ray per equilibrium, because the regret norms are free to take any positive value; the map therefore contracts toward a set, never toward a point.

We find the mechanism in the normalization. Scale invariance turns the cumulative regret norm into a control variable: the norm sets the step size of the round, and the step size determines how far the state can travel. If the norm rises to a finite limit, the step size freezes and the dynamics settle. We call this event *norm saturation*, and we prove that it always occurs. Three steps turn saturation into rates. Saturation is equivalent to a finite second-order path length, which gives the norm limit an explicit bound. Near a strictly complementary equilibrium the clipping pattern locks after one round, so the non-smooth map becomes real-analytic on the support face, and on that face the differentials of the value-centering shift cancel exactly, which leaves a Jacobian in closed form. The squared-norm increment that the algorithm computes in every round is two-sidedly proportional to the squared Nash gap, which turns an internal scalar into a progress certificate.

We make five contributions.

1. **We identify norm saturation as the mechanism behind last-iterate convergence.** We prove that the cumulative regret norm saturates on every zero-sum matrix game with an explicit bound, and that saturation alone forces the last-iterate Nash gap to zero (Theorems 1 and 2). Under a unique equilibrium the iterates converge pointwise, with no spectral or strict-complementarity hypothesis (Theorem 3).
2. **We derive a closed-form linear rate for a regret-matching dynamic with no restart, no smoothing and no step-size parameter.** Near a unique strictly complementary equilibrium the support freezes in one step (Lemma 6), and the last iterate contracts at a rate written in closed form from the payoff matrix restricted to the frozen support (Theorem 4). One scale-invariant quantity decides it: the saturated step size times the largest singular value of that submatrix. A limit that violates the condition is linearly unstable (Proposition 2).
3. **We turn a scalar the algorithm already computes into a progress certificate.** The squared-norm increment is two-sidedly proportional to the squared Nash gap, so gap ratios are certified from increment ratios up to a single condition-number factor that enters once and never accumulates with the iteration count (Theorem 5).
4. **We confirm every prediction on** 216 **random matrix games.** The closed-form spectral radius matches finite-difference Jacobians to 2.4×10^{-10}, the predicted rate matches measured slopes at median ratio 1.000, the certificate exponent measures 1.986 against a predicted 2, and the scale-invariant quantity never exceeds one on all 184 instances with a resolvable limit (Table 1).
5. **We carry the certificate to extensive-form poker.** On Kuhn and Leduc the same increment tracks the NashConv gap with log–log correlation 0.998 and 0.997, using $126\times$ and $64\times$ fewer best-response passes than per-iteration monitoring (Section 5).

2 Related Work

Regret matching and scale-invariant regret matching. Regret matching (RM, Hart & Mas-Colell, 2000) and its tree-form extension, counterfactual regret minimization (CFR, Zinkevich et al., 2007), drive equilibrium computation, through RM$^+$/CFR$^+$ (Tammelin, 2014; Bowling et al., 2015; Burch et al., 2019) and optimistic variants (Farina et al., 2019a;b). Abstraction and pruning refinements scale these methods to large extensive-form games (Li et al., 2024; Li & Huang, 2025; 2026b;c). The *last-iterate* was a standing gap: CFR diverges on rock–paper–scissors (Lee et al., 2021), and RM$^+$ and its predictive variant PRM$^+$ can oscillate forever (Farina et al., 2023). Zhang et al. (2025a) observed IREG-PRM$^+$'s linear last-iterate convergence, which they left open. Their bounded second-order path length already bounds the regret norm and supplies the saturation equivalence of Lemma 1; Theorem 1 restates it with an explicit constant, and the local theory below builds on that limit. Concurrent parameter-free regret matching (Meng et al., 2026) also uses a nondecreasing regret norm to reach $O(1/T)$ average-iterate rates; in our analysis the norm's *limit* selects the local step

size. Analyses of regret-matching dynamics now reach potential games and constrained optimization (Anagnostides et al., 2025b).

Linear last-iterate convergence. Extragradient descent (Korpelevich, 1976) and its optimistic variants (Daskalakis & Panageas, 2018; Mokhtari et al., 2020) achieve linear last-iterate convergence on polytopes under saddle-point metric-subregularity (Wei et al., 2021), and in extensive-form games (Lee et al., 2021). For regret-matching-type dynamics the linear last-iterate theorems of Cai et al. (2025); Meng et al. (2025); Perolat et al. (2021); Liu et al. (2023) restart or modify the update, and the unmodified dynamics were open; Theorem 4 closes that case with a rate read off the frozen-support spectrum, game-dependent as the lower bound of Cai et al. (2024) requires.

The closest computable metric is the normalized duality gap of Applegate et al. (2023): a Lagrangian gap, computable in linear time, that upper-bounds KKT residuals and drives restarts. It evaluates the current point through a dedicated gap computation. Our monitor uses the squared-norm growth the algorithm already computes, which Theorem 5 converts into a two-sided gap ratio on the last-iterate. Error-bound, active-set, and computable-metric comparisons are in Appendix F.

3 PRELIMINARIES

Notation. For $v \in \mathbb{R}^n$, $[v]_+$ is the entrywise positive part, $\|\cdot\|$ and $\|\cdot\|_1$ the Euclidean and ℓ_1 norms, Δ_n the probability simplex, $\mathbf{1}$ the all-ones vector, $\mathrm{supp}(x)$ the support of a nonnegative vector, and $\mathrm{Fix}(\Phi)$ the fixed-point set of a map Φ. For a matrix M, $\sigma_{\max}(M)$ is its largest singular value and $\sigma^+_{\min}(M)$ its smallest nonzero singular value.

Zero-sum matrix games. Consider a zero-sum game with two players and payoff matrix $A \in \mathbb{R}^{n_1 \times n_2}$: player 1 chooses $x \in \Delta_{n_1}$ to maximize $x^\top A y$, player 2 $y \in \Delta_{n_2}$ to minimize it. (x^*, y^*) is a *Nash equilibrium* (NE) when $x^\top A y^* \le (x^*)^\top A y^* \le (x^*)^\top A y$ for all x, y; the equilibrium payoff v^*, the *value*, is common to all NE (minimax theorem Von Neumann & Morgenstern, 2007). The (duality) gap of a strategy pair, $\mathrm{gap}(x,y) := \max_{i\in[n_1]}(Ay)_i - \min_{j\in[n_2]}(x^\top A)_j \ge 0$, vanishes iff (x,y) is an NE. For an NE (x^*, y^*), $S_1 := \mathrm{supp}(x^*)$ and $S_2 := \mathrm{supp}(y^*)$ are the supports, $\tilde{A} := A[S_1, S_2]$ the support submatrix, and $B := \tilde{A} - v^* \mathbf{1}\mathbf{1}^\top$ the *value-centered support matrix*, which governs the local spectrum (Section 4.2); the *support face* is the subspace $\mathcal{M}_S := \{(\tilde{r}_1, \tilde{r}_2) : \tilde{r}_1|_{S_1^c} = 0,\ \tilde{r}_2|_{S_2^c} = 0\}$. The NE is *strictly complementary* if off-support actions are strictly suboptimal:

$$\delta \;:=\; \min\Big\{ \min_{i \notin S_1} \big(v^* - (Ay^*)_i\big),\ \min_{j \notin S_2} \big(((x^*)^\top A)_j - v^*\big) \Big\} \;>\; 0, \tag{1}$$

where the minimum over an empty set is $+\infty$, as for a fully mixed player.

The IREG-PRM$^+$ algorithm. IREG-PRM$^+$ is IR-PRM$^+$ (Zhang et al., 2025a, Algorithm 2) run in the *extra-gradient* setup, one instance per player (pseudocode in Algorithm 1, Appendix A.1), in the strict form: the shifted vector $r^{(t)} = \tilde{r}^{(t)} + m^{(t)} - \gamma^{(t)}\mathbf{1}$ is unclipped and the update is $\tilde{r}^{(t+1)} = \big[r^{(t)} + g^{(t)}\big]_+$. The *regret* of an action is the payoff it would have earned against the opponent's strategy minus the player's actual payoff; in the *extra-gradient* setup each round first plays against a prediction of the opponent's payoff vector, then corrects with the observed one. Player 1 maintains a nonnegative cumulative regret vector $\tilde{r}_1^{(t)}$ (initialized at $\mathbf{0}$, as scale invariance requires) and the *pre-iterate* strategy $\tilde{x}^{(t)} = \tilde{r}_1^{(t)}/\|\tilde{r}_1^{(t)}\|_1$; player 2 symmetrically maintains $\tilde{r}_2^{(t)}$ and $\tilde{y}^{(t)}$, with payoff operator $-A^\top$. We write $N_i(t) := \|\tilde{r}_i^{(t)}\|_2$ for the regret norm. The key ingredient is the norm-preserving γ-shift:

Definition 1 (γ-shift)**.** Given $v \in \mathbb{R}^n$ and target $\tau > 0$, the map $\gamma \mapsto \|[v - \gamma\mathbf{1}]_+\|_2$ decreases strictly to zero, so

$$\big\| [v - \gamma\mathbf{1}]_+ \big\|_2 \;=\; \tau \tag{2}$$

has a unique solution $\gamma(v, \tau)$ (Zhang et al., 2025a, Appendix B); invoked with $v = \tilde{r}^{(t)} + m^{(t)}$ and $\tau = \|\tilde{r}^{(t)}\|_2$, it preserves the ℓ_2 norm of the prediction step exactly.

Dynamics, limits, and non-degeneracy. One round of IREG-PRM$^+$ defines a memoryless map Φ on the positive-norm states. Its restriction to the feasible part $\mathcal{M}_S \cap \{N_1, N_2 > 0\}$ of the

support face, written $\Phi_S := \Phi|_{\mathcal{M}_S}$, is a self-map near the fixed-point cone (Lemma 6). We write $z = (\tilde{r}_1, \tilde{r}_2)$ for the joint state and $N_i(z) := \|\tilde{r}_i\|_2$, so that along a trajectory $N_i(z^{(t)}) = N_i(t)$; $D\Phi_S$ is the Jacobian of the face map and $\mathrm{spec}(\cdot)$ its spectrum read as a multiset. On a bounded trajectory Ω is the ω-limit set, the set of subsequential limits, and $\mathrm{dist}(\cdot,\cdot)$ the Euclidean distance to a set. A strictly complementary equilibrium is *non-degenerate* if $\ker B = \mathrm{span}\{y^*|_{S_2}\}$ and $\ker B^\top = \mathrm{span}\{x^*|_{S_1}\}$, which forces $|S_1| = |S_2|$. For a strictly complementary equilibrium, non-degeneracy is equivalent to uniqueness (Lemma 11); Assumption 3 states it as a hypothesis. We use the two forms interchangeably.

Saturation. The dynamics are governed by the per-player regret norms; their convergence is the structural event of the theory:

Definition 2 (Norm saturation; per-player inverse norm)**.** Player i's trajectory *saturates* if $N_i(t) \uparrow N_{i,\infty} < \infty$. The *per-player inverse norm* at round t is $\eta_i(t) := 1/N_i(t)$ (defined once $N_i(t) > 0$); under saturation with $N_{i,\infty} > 0$ it converges to $\eta_{i,\infty} := 1/N_{i,\infty} > 0$. The symbol η_∞ is reserved for the joint ℓ_1-scale step size of the local theory (Section 4.2).

Let $t_i^* := \min\{t : N_i(t) > 0\}$ be the first nonzero round ($t_i^* = \infty$ iff $N_i \equiv 0$, a constant payoff); before it the algorithm does not move, and after it $N_i(t) > 0$ (Lemma 5). For $t \geq t_i^*$ the structure supplies three elementary facts (player 1): $\langle g^{(t)}, [r^{(t)}]_+ \rangle = 0$ (F1); $\|g^{(t)}\| \leq (1+\sqrt{n_1})\,\|A\|_{\mathrm{op}}\|y^{(t)} - \tilde{y}^{(t)}\|$ (F2; only the linear dependence is used); and by (F1), $N(t+1)^2 = N(t)^2 + \|g^{(t)}\|^2 - D(t)$ with $0 \leq D(t) \leq \|g^{(t)}\|^2$ (F3). The machinery of Zhang et al. (2025a) gives the following (proof in Appendix B.2, where Remark 3 records the tighter constant).

Lemma 1 (Saturation $\Leftrightarrow$ finite second-order path length)**.** *If both players run* IREG-PRM$^+$ *against each other (self-play) on A from $\tilde{r}_i^{(1)} = \mathbf{0}$, then both saturate, $N_i(t) \uparrow N_{i,\infty} < \infty$ for $i \in \{1,2\}$, if and only if the second-order path length is finite:*

$$\sum_{t=1}^{\infty} \left(\|x^{(t)} - \tilde{x}^{(t)}\|^2 + \|y^{(t)} - \tilde{y}^{(t)}\|^2 \right) \;<\; \infty.$$

Moreover, under saturation, for each player i with $t_i^* < \infty$ (if $t_i^* = \infty$ the left side below is 0),

$$\sum_{t=1}^{\infty} \|x_i^{(t)} - \tilde{x}_i^{(t)}\|^2 \;\leq\; \frac{8n_i\left(N_{1,\infty} + N_{2,\infty}\right)}{N_i(t_i^*)} \;<\; \infty, \tag{3}$$

and the internally available scalars $\beta_t^i := \langle m_i^{(t)}, x_i^{(t)} \rangle - \gamma_i^{(t)}$ ($\beta_t^i := 0$ for $t < t_i^*$) satisfy

$$\beta_t^i \;\geq\; \frac{1}{8n_i}\, N_i(t)\, \|x_i^{(t)} - \tilde{x}_i^{(t)}\|^2 \;\geq\; 0, \qquad \sum_{t=1}^{\infty} \left(\beta_t^1 + \beta_t^2\right) \;\leq\; N_{1,\infty} + N_{2,\infty}. \tag{4}$$

Theorem 1 (Saturation is unconditional)**.** *For every A and both players, $N_i(t) \uparrow N_{i,\infty} < \infty$. If $\tau := \max(t_1^*, t_2^*) < \infty$ then, with $K_i := (1 + \sqrt{n_i})\|A\|_{\mathrm{op}}$, $\alpha_i := N_i(\tau)/(8n_i)$, $M := \max\{K_1^2/\alpha_2, K_2^2/\alpha_1\}$, $Q_\tau := N_1(\tau)^2 + N_2(\tau)^2$,*

$$N_1(t) + N_2(t) \;\leq\; M + \sqrt{M^2 + 2Q_\tau} \qquad \textit{for all } t; \tag{5}$$

if $t_i^* = \infty$ then $N_i \equiv 0$ and the opponent's norm is eventually constant.

The proof (Appendix B.2) is four lines from (F2), (F3) and the finite-horizon form of (4); it restates the bounded second-order path length of Zhang et al. (2025a, Theorem 5.5) with an explicit constant. Boundedness alone does not give convergence, because square-summable residuals admit slow drift; Sections 4.1–4.2 supply the ω-limit argument and the rate.

Proposition 1 (Fixed points $\Leftrightarrow$ Nash equilibria)**.** *Let $z = (\tilde{r}_1, \tilde{r}_2)$ with $\tilde{r}_1 \neq \mathbf{0} \neq \tilde{r}_2$, and let Φ be the one-round map of Algorithm 1. Then $\Phi(z) = z$ if and only if $(x, y) := \left(\tilde{r}_1/\|\tilde{r}_1\|_1,\ \tilde{r}_2/\|\tilde{r}_2\|_1\right)$ is a Nash equilibrium of A. In that case the shifts equal the game value: $\gamma_1 = x^\top A y = v^*$ and $\gamma_2 = -v^*$. Consequently,*

$$\mathrm{Fix}(\Phi) \cap \{\tilde{r}_1 \neq \mathbf{0} \neq \tilde{r}_2\} \;=\; \left\{(c_1 x^*,\, c_2 y^*) \;:\; c_1, c_2 > 0,\ (x^*, y^*) \textit{ is an NE of } A\right\},$$

the positive cone over the NE directions: regret directions are pinned to equilibrium strategies while the scales c_1, c_2 that set the norms stay free. The proof is in Appendix B.1.

4 MAIN RESULTS

Saturation (Theorem 1) forces the last-iterate gap to vanish on every matrix game (Theorem 2), and to vanish pointwise under a unique equilibrium (Theorem 3). The linear rate is local: near a unique strictly complementary equilibrium that meets the scale-invariant step-size condition of Assumption 2, the rate has a closed form (Theorem 4). The same theory gives a ratio certificate (Theorem 5). Proofs are in Appendices B–D, sketches below.

4.1 GLOBAL THEORY: SATURATION IMPLIES CONVERGENCE

By Theorem 1, $N_i(t) \uparrow N_{i,\infty} < \infty$; we assume $t_+ := \max(t_1^*, t_2^*) < \infty$, so $N_{i,\infty} > 0$ (else one player never moves). For $t \geq t_+$ one round of IREG-PRM$^+$ acts on the positive-norm states through the memoryless map Φ of Section 3:

$$\mathcal{Z} \;:=\; \big\{\, z : \tilde{r}_1 \geq 0,\; \tilde{r}_2 \geq 0,\; N_1(z) > 0,\; N_2(z) > 0 \,\big\}, \qquad \Phi : \mathcal{Z} \to \mathcal{Z}. \tag{6}$$

Lemma 2 (Norm-preserving pairing)**.** *Let $a, b \in \mathbb{R}^n_{\geq 0} \setminus \{0\}$ with $\|a\|_2 = \|b\|_2 = N$, and let $x := a/\|a\|_1$, $\tilde{x} := b/\|b\|_1$. Then*

$$a - b \;=\; N\Big(\tfrac{x}{\|x\|_2} - \tfrac{\tilde{x}}{\|\tilde{x}\|_2}\Big), \qquad \textit{hence} \qquad \|a - b\|_2 \;\leq\; 2\sqrt{n}\, N\, \|x - \tilde{x}\|_2. \tag{7}$$

Applied with $a = [r_1(z)]_+$, $b = \tilde{r}_1$ (equal norms by the γ-shift), it bounds the prediction-step displacement by the pre-iterate movement, which is square-summable (Lemma 1; Appendix C.2, (27)*).*

Theorem 2 (Global asymptotics; no spectral condition)**.** *Let $t_+ < \infty$, so that both norms saturate at positive limits $N_{i,\infty}$ (Theorem 1). Then:*

(i) The tail trajectory $\{z^{(t)}\}_{t \geq t_+}$ has nonempty, compact ω-limit set Ω with $\Phi(\Omega) = \Omega$, $\operatorname{dist}(z^{(t)}, \Omega) \to 0$*, and*

$$\Omega \;\subseteq\; \operatorname{Fix}(\Phi) \,\cap\, \{\, z :\; \tilde{r}_i \geq 0,\; \|\tilde{r}_i\|_2 = N_{i,\infty} \,\}.$$

Every point of Ω is a Nash ray $(c_1 x^*, c_2 y^*)$ (Proposition 1), where $\gamma_1 = v^*$.

(ii) Consequently the last-iterate Nash gap vanishes:

$$\operatorname{gap}\big(x^{(t)}, y^{(t)}\big) \;=\; \max_i\, (Ay^{(t)})_i - \min_j\, (x^{(t)^\top} A)_j \;\longrightarrow\; 0,$$

and $\operatorname{dist}\big((x^{(t)}, y^{(t)}), \mathrm{NE}\big) \to 0$*, for the pre-iterates too.*

Only $t_+ < \infty$ is assumed; no spectral, strict-complementarity or uniqueness condition enters. If $t_+ = \infty$ the gap is 0 *in finitely many rounds (Lemma 18), so (ii) holds on every matrix game.*

Proof sketch (Appendix C.2). The displacement is $O(\|x - \tilde{x}\| + \|y - \tilde{y}\|)$, ℓ^2-summable (Lemma 1). Limit points are fixed points: continuity of the γ-shift gives $\|X(\bar{z}) - \tilde{X}(\bar{z})\| = 0$ on Ω, and (7) upgrades strategy equality to regret-vector equality (La Salle, 1976); part (ii) is a compactness bootstrap.

Theorem 3 (Point convergence under NE uniqueness; no spectral condition)**.** *In addition to the hypotheses of Theorem 2, suppose the Nash equilibrium (x^*, y^*) is unique. Then the whole state trajectory converges to a single point,*

$$z^{(t)} \;\longrightarrow\; \bar{z} := \Big(\tfrac{N_{1,\infty}}{\|x^*\|_2}\, x^*,\; \tfrac{N_{2,\infty}}{\|y^*\|_2}\, y^*\Big), \qquad x^{(t)}, \tilde{x}^{(t)} \to x^*, \quad y^{(t)}, \tilde{y}^{(t)} \to y^*.$$

The statement needs uniqueness alone.

Proof sketch (Appendix C.2). Uniqueness makes Fix the ray family $\{(c_1 x^*, c_2 y^*) : c_i > 0\}$; the level sets $\|\tilde{r}_i\|_2 = N_{i,\infty}$ pin c_i uniquely, so the bounded trajectory has a unique limit point.

Corollary 1 (Global entry + local capture $\Rightarrow$ eventual linear convergence)**.** *Assume the hypotheses of Theorem 2, and suppose Ω contains* at least one *point $z^\dagger$ to which Theorem 4 applies. Write $z^\dagger = (a_1 x^*, a_2 y^*)$ for a strictly complementary equilibrium (x^*, y^*) (Assumption 1), so that its*

effective step size $\eta(z^\dagger) := 1/\sqrt{a_1 a_2}$ *satisfies* $\eta(z^\dagger)\,\sigma_{\max}(B) < 1$. *Then the entire trajectory converges to a single fixed point* $z_\infty \in \mathrm{Fix}(\Phi)$*, and there exist* $T^* \geq t_+$*, constants* $C, C' < \infty$ *and* $\bar\rho \in (\rho^*, 1)$*, with* ρ^* *the closed-form rate of Theorem 4, such that for all* $t \geq T^*$*,*

$$\|z^{(t)} - z_\infty\| \;\leq\; C\,\bar\rho^{\,t-T^*}, \qquad \mathrm{gap}\big(x^{(t)}, y^{(t)}\big) \;\leq\; C'\,\bar\rho^{\,t-T^*}.$$

Under Assumptions 1 and 2 with $m \geq 2$ *the hypothesis on* Ω *is automatic. For the pure saddle* $m = 1$*, Lemma 6 gives capture in finitely many rounds, so the conclusion holds for any* $\bar\rho$*.*

Proof sketch (Appendix C.2). Every neighborhood of $z^\dagger$, in particular the capture neighborhood $\mathcal{U}'$ of Theorem 4, is visited infinitely often, so entry at finite T^* captures the trajectory.

Self-stabilization. A ray with $\eta_\infty\,\sigma_{\max}(B) > 1$ is linearly unstable (Theorem 4(ii)), so a nearby trajectory is generically pushed away while its norms keep growing (Lemma 5). We expect the growth to continue until a ray satisfying the condition is reached, the hypothesis we test in Section 5.

Proposition 2 (Instability of step-size-violating rays)**.** *Let* $z^\dagger = (c_1 x^*, c_2 y^*) \in \mathrm{Fix}(\Phi)$ *correspond to a strictly complementary, nondegenerate equilibrium (Assumption 3), and suppose* $\eta(z^\dagger)\,\sigma_{\max}(B) > 1$*. Then:*

- *(a) (Full space) There is a neighborhood* $B_\dagger \ni z^\dagger$ *and a* C^1 *local center-stable manifold* $W^{cs}_{loc}(z^\dagger)$ *of dimension* $d - d_u$*, where* $d := n_1 + n_2$ *and* $d_u = 2\,\#\{k : \eta(z^\dagger)\sigma_k(B) > 1\} \geq 2$*, such that every orbit that remains in* $B_\dagger$ *forever is contained in* $W^{cs}_{loc}(z^\dagger)$*.*
- *(b) (*$\mathcal{M}_S$*-local, measure zero) Within the support subspace* $\mathcal{M}_S$*, forward-invariant near* $z^\dagger$ *by Lemma 6, the local basin of attraction of* $z^\dagger$ *has Lebesgue measure zero in* $\mathcal{M}_S$*.*

Proof sketch (Appendix C.2). The proof combines real-analyticity of Φ near the ray, a center-stable manifold theorem for C^1 maps (Hirsch et al., 1977) in the saddle-avoidance form (Lee et al., 2019), and nonsingularity of the face Jacobian $J = I + L + L^2$ (Lemma 9). The eigenvalues of J satisfy $|\mu|^2 = 1 - q + q^2 \geq 3/4$ for all $q = \eta_\infty^2\sigma^2 \geq 0$, so preimages of null sets are null.

Remark 1 (Self-stabilization). Rays with $\eta_\infty\sigma_{\max}(B) > 1$ are unstable with Lebesgue-null basin in $\mathcal{M}_S$ (Proposition 2); once the bound holds strictly, capture follows (Corollary 1). Between those endpoints sits the link that the norms keep growing while the bound is violated; we measure it on all 184 instances with a resolvable limit, where $\eta_\infty\sigma_{\max}(B)$ has median 0.790, maximum 0.9985, and zero violations, with several instances just inside the boundary (Section 5). The center-stable set lives in the full state space, while the measurement follows the single zero-initialized orbit (Remark 14).

4.2 Local theory: support freezing and a closed-form linear rate

Saturation forces convergence; the observed rate is *linear*. Near a strictly complementary equilibrium one step of IREG-PRM$^+$ lands exactly on the support face, where the face dynamics are real-analytic and the rate is governed by the local spectrum; two hypotheses are needed beyond saturation.

Assumption 1 (Unique, strictly complementary equilibrium)**.** *The game* $A \in \mathbb{R}^{n_1 \times n_2}$ *has a unique Nash equilibrium* (x^*, y^*)*, with value* v^**, and the equilibrium is* strictly complementary*: with* $S_1 := \mathrm{supp}(x^*)$*,* $S_2 := \mathrm{supp}(y^*)$*,*

$$\delta \;:=\; \min\Big\{ \min_{i \notin S_1}\big(v^* - (Ay^*)_i\big),\; \min_{j \notin S_2}\big((x^{*\top}A)_j - v^*\big)\Big\} \;>\; 0. \tag{8}$$

The proofs use uniqueness only through a kernel condition on the support matrix: non-degeneracy, which for a strictly complementary equilibrium is equivalent to uniqueness (Section 3, Lemma 11). Throughout, $\tilde A := A[S_1, S_2]$ is the support submatrix, $B := \tilde A - v^*\mathbf{1}\mathbf{1}^\top$ the *value-centered support matrix*, $\xi := x^*|_{S_1}$ and $\psi := y^*|_{S_2}$ the equilibrium in face coordinates, and $E_+ := \{(a\xi, b\psi) : a, b > 0\}$ the fixed-point cone of Proposition 1. For $a, b > 0$ the effective step size at a fixed point (ax^*, by^*) is $\eta(a, b) := 1/\sqrt{ab}$; at the reference fixed point $z^* = (c_1 x^*, c_2 y^*)$ we write $\eta_\infty := \eta(c_1, c_2)$.

Assumption 2 (Spectral step-size condition)**.** *At the reference fixed point,* $\eta_\infty\,\sigma_{\max}(B) < 1$*; for a saturated trajectory the reference point is its limit* $\bar z$ *of Theorem 3,* $c_i = N_{i,\infty}/\|x^*\|_2$ *(resp.* $\|y^*\|_2$*).*

On the support face $\mathcal{M}_S$ (for nonnegative states, $\mathrm{supp}(\tilde r_i) \subseteq S_i$), one step absorbs with exact support freezing (Lemma 6). The face Jacobian at a fixed point is the closed form $J = I + L + L^2$, where L

is the one-step *extra-gradient* operator on B at effective step size $\eta(a,b)$, selected by the saturation level (Lemma 9). The centering identity of Lemma 10 makes the square of L have real nonpositive eigenvalues, so the one-step spectrum is pinned to a single closed-form threshold. Let $\Sigma^+(B)$ denote the multiset of nonzero singular values of B (exactly $m-1$ of them, $m := |S_1| = |S_2|$, under Assumption 3), and fix Assumption 1. The pure saddle $m = 1$ is set aside: there $B = 0$, $\mathcal{M}_S$ is the fixed-point cone, and Lemma 6 gives $\Phi(\mathcal{U}_0) \subseteq \mathrm{Fix}(\Phi)$ for a neighborhood $\mathcal{U}_0$ of the fixed point, so the iteration converges in one round; Theorem 4(ii)–(iv) and Section 4.3 assume $m \geq 2$.

Theorem 4 (Closed-form local linear rate)**.** *Let Assumption 1 hold with $m \geq 2$ and let $z^* = (c_1 x^*, c_2 y^*)$, $c_1, c_2 > 0$, be a reference fixed point.*

(i) ***(Spectrum.)*** *For every fixed point $z_f = (a\xi, b\psi)$ with $\eta = \eta(a,b)$,*

$$\mathrm{spec}\big(D\Phi_S(z_f)\big) \;=\; \{1,\,1\} \;\cup\; \big\{\,1 - \eta^2\sigma^2 \pm i\,\eta\sigma \;:\; \sigma \in \Sigma^+(B)\big\} \tag{9}$$

as a multiset; the eigenvalue 1 has algebraic multiplicity exactly two, is semisimple, with eigenspace the tangent space of the fixed-point cone.

(ii) ***(Stability threshold.)*** *The non-unit spectral radius is*

$$\rho(a,b) \;=\; \max_{\sigma\in\Sigma^+(B)} \sqrt{1 - \eta^2\sigma^2\big(1-\eta^2\sigma^2\big)}\,, \tag{10}$$

and $\rho(a,b) < 1$ iff $\eta\,\sigma_{\max}(B) < 1$, while $\eta\,\sigma_{\max}(B) > 1$ makes some mode unstable.

(iii) ***(Local linear last-iterate convergence.)*** *Under Assumption 2, let $\rho^* := \rho(c_1, c_2) < 1$. For every $\varepsilon \in (0, 1-\rho^*)$ there is an open neighborhood $\mathcal{U}' \ni z^*$ and $C_\varepsilon < \infty$ such that every trajectory with $z^{(t_0)} \in \mathcal{U}'$ converges to a fixed point $z_\infty = (a_\infty\xi, b_\infty\psi) \in E_+$ with*

$$\|z^{(t)} - z_\infty\| \;\leq\; C_\varepsilon\,(\rho^* + \varepsilon)^{\,t-t_0} \qquad \textit{for all } t \geq t_0. \tag{11}$$

(iv) ***(Gap at the same rate.)*** *With the same $\mathcal{U}'$,*

$$\mathrm{gap}\big(x^{(t)}, y^{(t)}\big) \;\leq\; L_{\mathrm{gap}}\,C_\varepsilon\,(\rho^* + \varepsilon)^{\,t-t_0} \qquad \textit{for all } t \geq t_0, \tag{12}$$

where $L_{\mathrm{gap}} < \infty$ depends only on $\|A\|$, c_1, c_2, and the dimensions.

Proof sketch (Step A, Section B.3.3; Steps B–C, Appendix B.4). Elementary: linear algebra, a uniform second-order Taylor expansion, and an induction; the spectral-radius route does not apply because the unit eigenvalue of multiplicity two never leaves the spectrum. The $E \oplus W$ splitting block-diagonalizes the Jacobian, so the E-drift is $O(\|w\|^2)$ and the limit exists by telescoping.

Only $\sigma_{\max}(B)$ decides whether the iteration contracts at all. Corollary 3 identifies which extreme nonzero singular value sets the rate: a small $\sigma^+_{\min}(B)$ slows the rate without threatening stability. The rate is *spectral*, in that δ enters the constants but not ρ^*. Parts (iii)–(iv) assert an *envelope* $O((\rho^*+\varepsilon)^t)$ for every $\varepsilon > 0$, and the realized rate matches it empirically (Section 5).

4.3 A RATIO CERTIFICATE

Saturation already implies $\mathrm{gap} \to 0$ (Lemma 1); the content is *quantitative*: a *ratio* certificate in which the observable norm increments track the unobservable gap at the level of *rates*, with the unobservable constants absorbed into one condition number. The observable is

$$\dot N_i(t) \;:=\; N_i(t{+}1)^2 - N_i(t)^2 \;\geq\; 0, \qquad \dot N_{\mathrm{tot}}(t) \;:=\; \dot N_1(t) + \dot N_2(t),$$

nonnegative by Lemma 5, available online at the cost of two scalar subtractions. Throughout we assume Assumptions 1 and 2 in the local regime of Theorem 4: once the support has frozen (Lemma 6) the state decomposes as $z^{(t)} = z_f^{(t)} + w^{(t)}$ with $z_f^{(t)}$ on the fixed-point cone, so $\mathrm{dist}(z^{(t)}, \mathrm{Fix}) = \|w^{(t)}\|$. The three lemmas, and the measurement choices behind them (on-support energy, since differencing N^2 loses representability; both players aggregated), are in Appendix D (Lemmas 24, 30, 26; Remarks 18, 19).

Table 1: Headline outcomes on the 216-instance testbed. Ratios and percentage points are summarized across instances; p10–p90 is the 10th to 90th percentile range. †Heuristic of Remark 1, empirically consistent with Proposition 2.

Prediction	Source	Instances	Reading
Exact support freezing before $T/2$	Lem. 6	216	99.5%
Rate equals finite-difference radius	Thm. 4	155	median 1.000 (p10–p90 0.984–1.016)
Closed form equals finite difference	Lem. 9	32	2.4×10^{-10}
Certificate exponent $\theta = 2$	Cor. 2	129	1.986; 96.1%
$\eta_\infty \sigma_{\max}(B) \le 1$ at limits†	Rem. 1	184	max 0.9985
Ratio certificate holds	Thm. 5	124	124/124
Progress monitoring, Kuhn/Leduc	empirical	2	1.80, 3.22

Theorem 5 (Ratio certificate)**.** *Assume Assumptions 1 and 2, and let the trajectory be in the frozen local regime of Theorem 4 with $\|w^{(t)}\| \le \min(\delta_2, \delta_3)$ for all $t \ge t_{\mathrm{reg}}$, the regime-entry time, where δ_2 and δ_3 are the radii of Lemmas 30 and 26. Then for all $t_2 \ge t_1 \ge t_{\mathrm{reg}}$ with $\dot{N}_{\mathrm{tot}}(t_1) > 0$,*

$$\kappa^{-1} \sqrt{\frac{\dot{N}_{\mathrm{tot}}(t_2)}{\dot{N}_{\mathrm{tot}}(t_1)}} \;\le\; \frac{\mathrm{gap}(t_2)}{\mathrm{gap}(t_1)} \;\le\; \kappa \sqrt{\frac{\dot{N}_{\mathrm{tot}}(t_2)}{\dot{N}_{\mathrm{tot}}(t_1)}}, \qquad \kappa := \frac{L_{\mathrm{gap}}}{c_-} \cdot \sqrt{\frac{n_+}{n_-}},$$

where $L_{\mathrm{gap}}, c_-, n_\pm$ are the constants of Lemmas 30 and 26; κ is the product of a gap-sharpness condition number L_{gap}/c_- and a spectral condition number $\sqrt{n_+/n_-}$, independent of t_1, t_2 and of the trajectory amplitude; no factor $1/(1-\rho^*)$ enters.

The hypothesis $\dot{N}_{\mathrm{tot}}(t_1) > 0$ excludes exactly one degenerate case: by Lemma 26, $\dot{N}_{\mathrm{tot}}(t_1) = 0$ forces $w^{(t_1)} = 0$, the state on the fixed-point cone with $\mathrm{gap}(t_1) = 0$ already. Operationally, one calibration of the gap at a single time t_1 converts the calibrated estimate $\mathrm{gap}(t_2) = \mathrm{gap}(t_1)\sqrt{\dot{N}_{\mathrm{tot}}(t_2)/\dot{N}_{\mathrm{tot}}(t_1)}$ into a two-sided bracket at every later time. The constant κ is not observable online: it is built from equilibrium quantities (the support, the smallest on-support probabilities, the extreme nonzero singular values of B) and from regime thresholds unknown to the running algorithm. The operational content of the certificate is the slope-2 law of Corollary 2, which needs no equilibrium constant: only $\dot{N}_{\mathrm{tot}}$ and the gap on the fitted window.

Corollary 2 ($\theta = 2$ prediction)**.** *Under the hypotheses of Theorem 5, for all $t \ge t_{\mathrm{reg}}$ with $\mathrm{gap}(t) > 0$ (equivalently $\dot{N}_{\mathrm{tot}}(t) > 0$, Lemma 26),*

$$\log \dot{N}_{\mathrm{tot}}(t) \;=\; 2 \log \mathrm{gap}(t) \;+\; C(t), \qquad C(t) \in \left[\log \tfrac{n_-}{L_{\mathrm{gap}}^2},\, \log \tfrac{n_+}{c_-^2}\right],$$

an interval of fixed width $2 \log \kappa$ independent of t. Consequently, over any measurement window with $\mathrm{sd}(\log \mathrm{gap}) > 0$, the least-squares slope θ of $\log \dot{N}_{\mathrm{tot}}$ against $\log \mathrm{gap}$ satisfies $|\theta - 2| \le \log \kappa / \mathrm{sd}(\log \mathrm{gap})$, where $\mathrm{sd}(\log \mathrm{gap})$ is the standard deviation of the recorded gap logarithms. As the window spans more decades of gap decay, the fitted slope is forced to 2.

Under the corollary's hypotheses, a fitted slope differing from 2 by more than $\log \kappa / \mathrm{sd}(\log \mathrm{gap})$ would refute the local theory; on the 216-instance suite the slope has median 1.986 with 96.1% of slope-measurable instances between 1.8 and 2.2 (Section 5).

5 Experiments

The theory makes five falsifiable statements: exact support freezing (Lemma 6); a last-iterate rate bounded by the frozen-face spectral radius, closed form in B and the saturation levels (Theorem 4); the exponent $\theta = 2$ (Corollary 2); a ratio certificate with a non-accumulating constant (Theorem 5); and $\eta_\infty \sigma_{\max}(B) \le 1$ at realized limits, the heuristic of Remark 1. We test all five on 216 random matrix games and monitor extensive-form runs from the same increment. Table 1 gives the headline readings; per-protocol statistics are in Appendix E.5.

Instances cross $n \in \{5, 20, 50\}$, three payoff distributions (Gaussian; uniform on $[-1, 1]$; sparse ± 1, density 0.3), two generators (raw and antisymmetrized $(G - G^\top)/\sqrt{2}$), and 12 seeds, giving

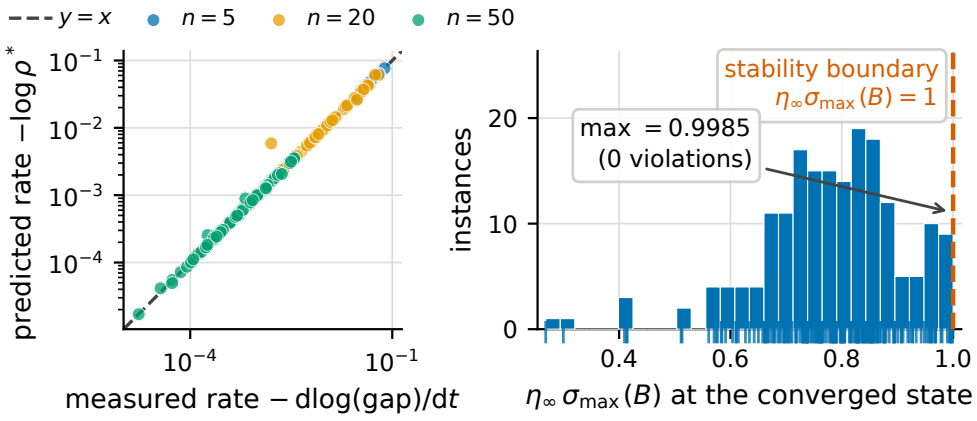


Figure 1: Left: predicted rate (frozen-face spectral radius) vs. measured semilog gap slope on the 155 instances with a fitted slope (log–log; color = n); five pre-freeze snapshots, one clearly off-line. Right: $\eta_\infty\,\sigma_{\max}(B)$ on all 184 measurable instances, none violating the boundary $= 1$ (dashed).

$3\times3\times2\times12 = 216$ instances. Each is solved numerically by linear programming for the equilibrium, supports, value v^*, and margin. IREG-PRM$^+$ runs $T = 2\times10^5$ iterations (10^6 for $n = 50$) on the *unclipped* prediction step $r^{(t)} = \tilde{r}^{(t)} + m^{(t)} - \gamma\mathbf{1}$; clipping emits off-support "dust" that breaks exact support freezing. The full protocol is in Appendix E.

1. Exact support freezing and the local rate. The exact-zero pattern of $\tilde{r}$ freezes before $T/2$ on 99.5% of instances, with median freeze time $2.7\times10^{-4}\,T$, matching Lemma 6. On the 155 instances with a fitted slope, a finite-difference Jacobian at a state snapshot reproduces the measured semilog slope at median ratio 1.000 and correlation 0.9982 (Figure 1, left); the snapshot is stored once the gap has fallen by 10^4. The envelope of Theorem 4(iii) is therefore tight.

2. The closed-form spectral radius matches finite differences. Lemmas 9 and 10 give the one-step spectrum in closed form from $B = \tilde{A} - v^*\mathbf{1}\mathbf{1}^\top$ and the saturation levels. Rebuilding B from the *converged* support and saturation levels reproduces the finite-difference value of the largest non-unit modulus ρ_2 on a 32-instance sample, median relative error 2.4×10^{-10} (Appendix E). At a degenerate instance whose Jacobian has nine eigenvalues of modulus one, outside Theorem 4's hypotheses, the formula still matches empirically.

3. The certificate exponent is $\theta = 2$. The exponent is measured on a candidate frozen window (exact-zero pattern of $\tilde{r}$ fixed), as the per-unit-time on-support energy accumulated between recorded times, an interval proxy for $\dot{N}_{\mathrm{tot}}$ (Lemma 24), against the interval's geometric-mean gap (representability guard, Appendix E). The slope has median 1.986, with 96.1% of the 129 slope-measurable instances between 1.8 and 2.2 (cross-checks in Appendix E).

4. Self-stabilization, the ratio certificate, and the point estimate. On all 184 measurable instances (the other 32 reach an exactly zero gap in finitely many rounds, Appendix E) $\eta_\infty\,\sigma_{\max}(B) \le 1$ holds: the realized dynamics satisfy the spectral condition of Theorem 4, and no instance reaches the boundary (Remark 1). An empirical interval proxy of Theorem 5's ratio law holds on $124/124$ instances against a first-order benchmark constant. Its deviation κ_{emp} has median 1.63 against an assembled κ of 1.7×10^8. Read forward as a gap estimate, the certificate's worst guarded deviation has median 1.50 (p90 1.90, max 3.9; below 2 on 91.5%).

5. Progress monitoring in extensive-form games (empirical). The accumulated update energy $G(t) = \sum_{s\le t}\sum_j \|g_j^{(s)}\|^2$ over information sets is one scalar per iteration, with no extra pass; its per-unit-time increment at recorded times is the monitored quantity. On Kuhn and Leduc poker the increment's log–log correlation with the NashConv gap (the sum of per-player best-response gains) is 0.998 and 0.997, and the point estimate calibrated at the first guarded increment has worst deviation 1.80 and 3.22 over the window. The sparse record schedule uses $126\times$ and $64\times$ fewer best-response passes than an every-iteration schedule; a NashConv pass costs 0.076/0.087 of a solver iteration.

AI USE STATEMENT

Generative AI assisted this work in three ways: polishing the language of the manuscript, drafting and debugging parts of the experimental codebase, and checking proofs during revision. The authors verified every mathematical result, ran and validated all AI-assisted code, and take responsibility for the correctness and final content of the paper.

REPRODUCIBILITY STATEMENT

Appendix E gives the full protocol: the 216-instance testbed, the LP-based equilibrium solver, and the definitions of all measured quantities. Proofs appear in Appendices B–D. The anonymized supplement reproduces every reported number on a single CPU node from the released per-instance data, and regenerates the full suite from scratch. It also provides the source code of the solver, the game generators and the measurement protocols.

REFERENCES

Jacob Abernethy, Peter L Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. In *Proceedings of the 24th Annual Conference on Learning Theory*, pp. 27–46. JMLR Workshop and Conference Proceedings, 2011.

Ioannis Anagnostides and Tuomas Sandholm. Convergence of $\log(1/\epsilon)$ for gradient-based algorithms in zero-sum games without the condition number: A smoothed analysis. *Advances in Neural Information Processing Systems*, 37:125839–125878, 2024.

Ioannis Anagnostides, Gabriele Farina, Ioannis Panageas, and Tuomas Sandholm. Optimistic mirror descent either converges to nash or to strong coarse correlated equilibria in bimatrix games. *Advances in Neural Information Processing Systems*, 35:16439–16454, 2022a.

Ioannis Anagnostides, Ioannis Panageas, Gabriele Farina, and Tuomas Sandholm. On last-iterate convergence beyond zero-sum games. In *International Conference on Machine Learning*, pp. 536–581. PMLR, 2022b.

Ioannis Anagnostides, Gabriele Farina, Tuomas Sandholm, and Brian Hu Zhang. A polynomial-time algorithm for variational inequalities under the minty condition. *arXiv preprint arXiv:2504.03432*, 2025a.

Ioannis Anagnostides, Emanuel Tewolde, Brian Hu Zhang, Ioannis Panageas, Vincent Conitzer, and Tuomas Sandholm. Convergence of regret matching in potential games and constrained optimization. *arXiv preprint arXiv:2510.17067*, 2025b.

David Applegate, Oliver Hinder, Haihao Lu, and Miles Lubin. Faster first-order primal-dual methods for linear programming using restarts and sharpness. *Mathematical Programming*, 201(1):133–184, 2023.

HF Bohnenblust, S Karlin, and LS Shapley. Solutions of discrete, two-person games. *Contributions to the Theory of Games*, 1:51–72, 1950.

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. Solving imperfect-information games via discounted regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1829–1836, 2019.

Neil Burch, Matej Moravcik, and Martin Schmid. Revisiting cfr+ and alternating updates. *Journal of Artificial Intelligence Research*, 64:429–443, 2019.

James V Burke and Jorge J Moré. On the identification of active constraints. *SIAM Journal on Numerical Analysis*, 25(5):1197–1211, 1988.

Yang Cai, Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Weiqiang Zheng. Fast last-iterate convergence of learning in games requires forgetful algorithms. *Advances in Neural Information Processing Systems*, 37:23406–23434, 2024.

Yang Cai, Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Weiqiang Zheng. Last-iterate convergence properties of regret-matching algorithms in games. In *The Thirteenth International Conference on Learning Representations*, 2025.

Constantinos Daskalakis and Ioannis Panageas. The limit points of (optimistic) gradient descent in min-max optimization. *Advances in neural information processing systems*, 31, 2018.

Gabriele Farina, Christian Kroer, Noam Brown, and Tuomas Sandholm. Stable-predictive optimistic counterfactual regret minimization. In *International conference on machine learning*, pp. 1853–1862. PMLR, 2019a.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Optimistic regret minimization for extensive-form games via dilated distance-generating functions. *Advances in neural information processing systems*, 32, 2019b.

Gabriele Farina, Christian Kroer, and Tuomas Sandholm. Faster game solving via predictive blackwell approachability: Connecting regret matching and mirror descent. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 5363–5371, 2021.

Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, and Haipeng Luo. Regret Matching+: (In)Stability and Fast Convergence in Games. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL `https://openreview.net/forum?id=nYgs0qZJ97`.

Amy Greenwald, Zheng Li, and Casey Marks. Bounds for regret-matching algorithms. In *AI&M*, 2006.

Warren L Hare and Adrian S Lewis. Identifying active constraints via partial smoothness and prox-regularity. *Journal of Convex Analysis*, 11(2):251–266, 2004.

Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000.

M Hirsch, C Pugh, and Michael Shub. Invariant manifolds (lecture notes in mathematics, 583), 1977.

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

Joseph P La Salle. *The stability of dynamical systems*. SIAM, 1976.

Chung-Wei Lee, Christian Kroer, and Haipeng Luo. Last-iterate convergence in extensive-form games. *Advances in Neural Information Processing Systems*, 34:14293–14305, 2021.

Jason D Lee, Ioannis Panageas, Georgios Piliouras, Max Simchowitz, Michael I Jordan, and Benjamin Recht. First-order methods almost always avoid strict saddle points: Jd lee et al. *Mathematical programming*, 176(1):311–337, 2019.

Boning Li and Longbo Huang. Efficient online pruning and abstraction for imperfect information extensive-form games. In *The Thirteenth International Conference on Learning Representations*, 2025.

Boning Li and Longbo Huang. Agents that certify their own exploits: Confidence-scheduled restricted responses for safe opponent exploitation. *arXiv preprint arXiv:2607.28520*, 2026a.

Boning Li and Longbo Huang. Effective, efficient, and general information abstraction for imperfect-information extensive-form games. *arXiv preprint arXiv:2605.10900*, 2026b.

Boning Li and Longbo Huang. Real-time parallel counterfactual regret minimization. *arXiv preprint arXiv:2605.19928*, 2026c.

Boning Li, Zhixuan Fang, and Longbo Huang. Rl-cfr: improving action abstraction for imperfect information extensive-form games with reinforcement learning. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 27752–27770, 2024.

Boning Li, Yu Chen, and Longbo Huang. Correlated chance sampling for monte carlo counterfactual regret minimization. *arXiv preprint arXiv:2607.27035*, 2026.

Jingwei Liang, Jalal Fadili, and Gabriel Peyré. Activity identification and local linear convergence of forward–backward-type methods. *SIAM Journal on Optimization*, 27(1):408–437, 2017.

Mingyang Liu, Asuman E. Ozdaglar, Tiancheng Yu, and Kaiqing Zhang. The power of regularization in solving extensive-form games. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=bPiHuNUNv_R.

Jason R Marden, Gürdal Arslan, and Jeff S Shamma. Regret based dynamics: convergence in weakly acyclic games. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, pp. 1–8, 2007.

Brendan McMahan. Follow-the-regularized-leader and mirror descent: Equivalence theorems and l1 regularization. In Geoffrey Gordon, David Dunson, and Miroslav Dudík (eds.), *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *Proceedings of Machine Learning Research*, pp. 525–533, Fort Lauderdale, FL, USA, 11–13 Apr 2011. PMLR.

Linjian Meng, Youzhi Zhang, Zhenxing Ge, Tianyu Ding, Shangdong Yang, Zheng Xu, Wenbin Li, and Yang Gao. Last-iterate convergence of smooth regret matching$^+$ variants in learning nash equilibria. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=v771lscnlS.

Linjian Meng, Youzhi Zhang, Shangdong Yang, Wenbin Li, Tianyu Ding, and Yang Gao. A faster parameter-free regret matching algorithm. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=JLllvi7dsg.

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 Silvia Chiappa and Roberto Calandra (eds.), *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1497–1507. PMLR, 26–28 Aug 2020. URL https://proceedings.mlr.press/v108/mokhtari20a.html.

Julien Perolat, Remi Munos, Jean-Baptiste Lespiau, Shayegan Omidshafiei, Mark Rowland, Pedro Ortega, Neil Burch, Thomas Anthony, David Balduzzi, Bart De Vylder, Georgios Piliouras, Marc Lanctot, and Karl Tuyls. From poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 8525–8535. PMLR, 18–24 Jul 2021. URL https://proceedings.mlr.press/v139/perolat21a.html.

Thomas Pethick, Ioannis Mavrothalassitis, and Volkan Cevher. Efficient interpolation between extragradient and proximal methods for weak mvis. In *The Thirteenth International Conference on Learning Representations*, 2025.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Machine Learning*, 4(2): 107–194, 2011.

Oskari Tammelin. Solving large imperfect information games using cfr+. *arXiv preprint arXiv:1407.5042*, 2014.

Paul Tseng. On linear convergence of iterative methods for the variational inequality problem. *Journal of Computational and Applied Mathematics*, 60(1-2):237–252, 1995.

John Von Neumann and Oskar Morgenstern. Theory of games and economic behavior: 60th anniversary commemorative edition. In *Theory of games and economic behavior*. Princeton university press, 2007.

Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=dx11_7vm5_r.

Brian Hu Zhang and Tuomas Sandholm. Finding and certifying (near-) optimal strategies in black-box extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pp. 5779–5788, 2021.

Brian Hu Zhang, Ioannis Anagnostides, and Tuomas Sandholm. Scale-invariant regret matching and online learning with optimal convergence: Bridging theory and practice in zero-sum games. *arXiv preprint arXiv:2510.04407*, 2025a.

Brian Hu Zhang, Ioannis Anagnostides, Emanuel Tewolde, Ratip Emin Berker, Gabriele Farina, Vincent Conitzer, and Tuomas Sandholm. Expected variational inequalities. In *International Conference on Machine Learning*, pp. 74422–74446. PMLR, 2025b.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th international conference on machine learning (icml-03)*, pp. 928–936, 2003.

Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20, 2007.

A THEOREM LEDGER: CLAIMS, HYPOTHESES, AND STATUS

Hypotheses at a glance. Saturation is a theorem (Theorem 1), and the gap conclusion of Theorem 2 is unconditional: the case $t_+ = \infty$ is settled by Lemma 18, while the ω-limit set and point convergence assume $t_+ < \infty$. The linear rate and the certificate require strictly more: a unique strictly complementary equilibrium with $m \geq 2$ and the post-saturation step-size condition (Assumptions 1 and 2). Table 2 records each claim with its hypotheses, epistemic status, and verification.

Table 2: Ledger of headline claims. A1 is Assumption 1 (unique strictly complementary equilibrium), A2 is Assumption 3 (strictly complementary, non-degenerate equilibrium ray), A3 is Assumption 2 ($\eta_\infty \sigma_{\max}(B) < 1$). "Ver." points to the numbered verification of Appendix E.3.

Claim	Hypotheses	Status	Ver.
Norm saturation, explicit bound (Thm. 1)	none	proved	n/a (bound loose by 10^3–10^6, Rem. 4)
Gap $\to$ 0, ω-limit in Fix (Thm. 2)	t_+ (i); none (ii)	proved	n/a (216 inst. converge)
Point convergence (Thm. 3)	unique NE, t_+ (else Lem. 18)	proved	n/a
One-step support freezing (Lem. 6)	strict complementarity, local	proved	1 (99.5%)
Closed-form spectrum $J = I + L + L^2$ (Lem. 9, 10)	A2, frozen face	proved	3 (32 inst.)
Linear rate envelope $(\rho^* + \varepsilon)^t$ (Thm. 4)	A1 + A3, local entry	proved (upper envelope; realized rate $= \rho^*$ is empirical, §4.2)	2 (FD Jacobian, 155 inst.); 3 (closed form, 32)
Instability of limits with $\eta_\infty \sigma_{\max}(B) > 1$ (Prop. 2)	A2 (strict compl., nondegenerate), $\eta_\infty \sigma_{\max}(B) > 1$	proved (thin basins; "observed limits satisfy A3" is a genericity argument, Rem. 1)	5 (184, 0 violations)
Ratio certificate κ-sandwich (Thm. 5)	A1 + A3, frozen regime	proved (pointwise)	6 (124/124, empirical interval proxy)
Slope-two law $\theta = 2$ (Cor. 2)	A1 + A3, frozen regime	proved (asymptotic in window width)	4 (96.1%)
$\dot{N}_{\text{tot}}$ = on-support energy (Lem. 24)	frozen regime	proved (identity)	4, 6 (proxy: pattern not logged)
Self-stabilization narrative (norm grows until A3 holds)	n/a	heuristic (instability proved; growth and the zero-initialized orbit not; Rem. 1)	5 (consistent)
Degenerate-instance sandwich	outside A1	empirical only (κ assembled for 3 of 8)	App. E
EFG monitoring (Kuhn, Leduc)	outside matrix theory	empirical only	8

A.1 PSEUDOCODE OF IREG-PRM$^+$

Algorithm 1 lists one round of the strict extra-gradient form described in Section 3; player 2 runs the same steps with payoff operator $-A^\top$.

B PROOFS FOR THE LOCAL THEORY

Throughout this appendix, one round of IREG-PRM$^+$ (Algorithm 2 of Zhang et al., 2025a, strict version: the shifted regret vector is kept unclipped and clipping is applied once, in the update) from a state $z = (\tilde{r}_1, \tilde{r}_2)$ with $\tilde{r}_1, \tilde{r}_2 \neq 0$ reads, for player 1,

$$\begin{aligned}
&\tilde{x} = \tilde{r}_1/\|\tilde{r}_1\|_1, \qquad \tilde{y} = \tilde{r}_2/\|\tilde{r}_2\|_1, \qquad m_1 = A\tilde{y}, \\
&\gamma_1 \text{ the unique solution of } \|[\tilde{r}_1 + m_1 - \gamma_1 \mathbf{1}]_+\|_2 = \|\tilde{r}_1\|_2 =: N_1, \\
&r_1 := \tilde{r}_1 + m_1 - \gamma_1 \mathbf{1} \ \text{(kept unclipped)}, \qquad p_1 := [r_1]_+, \qquad x := p_1/\|p_1\|_1, \\
&\text{res}_1 := A(y - \tilde{y}), \qquad g_1 := \text{res}_1 - \langle \text{res}_1, x \rangle \mathbf{1}, \qquad \tilde{r}_1' = [r_1 + g_1]_+ ,
\end{aligned} \tag{13}$$

Algorithm 1 IREG-PRM$^+$ (strict Algorithm 2 of Zhang et al., 2025a, extra-gradient setup): round t, player 1's updates.

Require: state $\tilde{r}^{(t)} \in \mathbb{R}^{n_1}_{\geq 0}$, default $\tilde{x}^{(1)} \leftarrow \mathbf{1}/n_1$; if $\tilde{r}^{(t)} = \mathbf{0}$: keep $\tilde{x}^{(t)} \leftarrow \tilde{x}^{(t-1)}$ for $t \geq 2$ (the default at $t = 1$), set $m^{(t)} \leftarrow \mathbf{0}$, $r^{(t)} \leftarrow \mathbf{0}$, play $x^{(t)} \leftarrow \tilde{x}^{(t)}$, and go to line 6 (no shift)

1: $\tilde{x}^{(t)} \leftarrow \tilde{r}^{(t)}/\|\tilde{r}^{(t)}\|_1$ ▷ pre-iterate; $\tilde{y}^{(t)}$ symmetric
2: $m^{(t)} \leftarrow A\tilde{y}^{(t)}$ ▷ prediction: payoff at the opponent's pre-iterate
3: $v^{(t)} \leftarrow \tilde{r}^{(t)} + m^{(t)}$; $\gamma^{(t)} \leftarrow$ solution of $\|\left[v^{(t)} - \gamma\mathbf{1}\right]_+\|_2 = \|\tilde{r}^{(t)}\|_2$
4: $r^{(t)} \leftarrow \tilde{r}^{(t)} + m^{(t)} - \gamma^{(t)}\mathbf{1}$ ▷ unclipped; norm-preserving
5: $x^{(t)} \leftarrow \left[r^{(t)}\right]_+ / \|\left[r^{(t)}\right]_+\|_1$ ▷ main iterate
6: observe $u^{(t)} \leftarrow Ay^{(t)}$; $\mathrm{res}^{(t)} \leftarrow u^{(t)} - m^{(t)}$
7: $g^{(t)} \leftarrow \mathrm{res}^{(t)} - \langle \mathrm{res}^{(t)}, x^{(t)} \rangle \mathbf{1}$ ▷ centered instantaneous regret
8: $\tilde{r}^{(t+1)} \leftarrow \left[r^{(t)} + g^{(t)}\right]_+$

and identically for player 2 with A replaced by $-A^\top$ (so $m_2 = -A^\top\tilde{x}$, $\mathrm{res}_2 = -A^\top(x - \tilde{x})$, and $y := p_2/\|p_2\|_1$). Recall that $\varphi(\gamma) := \|[v - \gamma\mathbf{1}]_+\|_2$ is continuous and nonincreasing in γ, strictly decreasing wherever it is positive, with $\varphi(\gamma) \to \infty$ as $\gamma \to -\infty$ and $\varphi(\gamma) \to 0$ as $\gamma \to +\infty$; since $N_1 > 0$, the defining equation for γ_1 has a unique solution, and by construction the clipped shifted vector is norm-preserving: $\|p_1\|_2 = N_1$. Two elementary facts are used repeatedly: $\langle [u]_+, u \rangle = \|[u]_+\|_2^2$ for every u, and the orthogonality

$$\langle g_1, p_1 \rangle = \langle \mathrm{res}_1, p_1 \rangle - \langle \mathrm{res}_1, x \rangle \, \|p_1\|_1 = 0, \tag{14}$$

which holds by the definition $x = p_1/\|p_1\|_1$ and $\langle \mathbf{1}, p_1 \rangle = \|p_1\|_1$ (as $p_1 \geq 0$).

B.1 Fixed points: an algebraic proof of Proposition 1

We first isolate two statewise lemmas. Both concern a single round at a single state; no property of the trajectory is used anywhere in this subsection.

Lemma 3 (γ-shift energy identity)**.** *Let $\tilde{r} \in \mathbb{R}^n_{\geq 0} \setminus \{0\}$ and $m \in \mathbb{R}^n$, put $N := \|\tilde{r}\|_2$, let γ be the unique solution of $\|[\tilde{r} + m - \gamma\mathbf{1}]_+\|_2 = N$, and set $r := \tilde{r} + m - \gamma\mathbf{1}$, $p := [r]_+$, $x := p/\|p\|_1$, and $\tilde{x} := \tilde{r}/\|\tilde{r}\|_1$. Then*

$$\langle m, x \rangle - \gamma \;=\; \frac{\|p - \tilde{r}\|_2^2}{2\,\|p\|_1}\,. \tag{15}$$

In particular $\langle m, x \rangle \geq \gamma$, with equality if and only if $p = \tilde{r}$; and

$$\langle m, x \rangle - \gamma \;\geq\; \frac{N}{8n}\,\|x - \tilde{x}\|_2^2\,. \tag{16}$$

Proof. Since $p = [r]_+$ we have $\langle p, r \rangle = \|p\|_2^2 = N^2$. Substituting $r = \tilde{r} + m - \gamma\mathbf{1}$ and $\langle p, \mathbf{1} \rangle = \|p\|_1$,

$$\langle p, \tilde{r} \rangle + \langle p, m \rangle - \gamma\|p\|_1 \;=\; N^2.$$

Because $\|p\|_2 = \|\tilde{r}\|_2 = N$,

$$N^2 - \langle p, \tilde{r} \rangle = \tfrac{1}{2}\big(\|p\|_2^2 + \|\tilde{r}\|_2^2 - 2\langle p, \tilde{r} \rangle\big) = \tfrac{1}{2}\|p - \tilde{r}\|_2^2,$$

and dividing by $\|p\|_1 > 0$ (note $\|p\|_1 \geq \|p\|_2 = N > 0$) gives (15). Nonnegativity and the equality case are immediate. For (16), write

$$x - \tilde{x} = \frac{p - \tilde{r}}{\|p\|_1} + \tilde{x}\,\frac{\|\tilde{r}\|_1 - \|p\|_1}{\|p\|_1},$$

and use $\big|\|\tilde{r}\|_1 - \|p\|_1\big| \leq \|p - \tilde{r}\|_1 \leq \sqrt{n}\,\|p - \tilde{r}\|_2$ together with $\|\tilde{x}\|_2 \leq \|\tilde{x}\|_1 = 1$ to obtain $\|x - \tilde{x}\|_2 \leq (1 + \sqrt{n})\|p - \tilde{r}\|_2/\|p\|_1 \leq 2\sqrt{n}\,\|p - \tilde{r}\|_2/\|p\|_1$. Hence $\|p - \tilde{r}\|_2^2 \geq \|p\|_1^2\|x - \tilde{x}\|_2^2/(4n)$, and (15) together with $\|p\|_1 \geq N$ gives (16). □

Lemma 4 (Clipping rigidity)**.** *Let $r, g \in \mathbb{R}^n$ with $\langle g, [r]_+ \rangle = 0$. Then $\|[r+g]_+\|_2 \geq \|[r]_+\|_2$, and if equality holds then*

$$[r+g]_+ \;=\; [r]_+\,, \qquad \textit{and} \qquad g_i = 0 \;\textit{for every } i \textit{ with } r_i > 0.$$

Proof. By the orthogonality hypothesis, $\|[r]_+ + g\|_2^2 = \|[r]_+\|_2^2 + \|g\|_2^2$. Define

$$D := \|[r]_+ + g\|_2^2 - \|[r+g]_+\|_2^2 = \sum_i d_i, \qquad d_i := \big([r_i]_+ + g_i\big)^2 - [r_i + g_i]_+^2\,,$$

so that $\|[r+g]_+\|_2^2 = \|[r]_+\|_2^2 + \|g\|_2^2 - D$. We claim $0 \leq d_i \leq g_i^2$ for every i, and we characterize the equality cases of the upper bound.

- *Case $r_i \geq 0$, $r_i + g_i \geq 0$:* then $d_i = (r_i+g_i)^2 - (r_i+g_i)^2 = 0$, and $g_i^2 - d_i = g_i^2 \geq 0$, with equality iff $g_i = 0$.
- *Case $r_i \geq 0$, $r_i + g_i < 0$:* then $d_i = (r_i+g_i)^2 \geq 0$ and $g_i^2 - d_i = -r_i(r_i + 2g_i)$. Here $g_i < -r_i \leq 0$, hence $r_i + 2g_i = (r_i+g_i) + g_i < 0$, so $g_i^2 - d_i = r_i\big(-(r_i+2g_i)\big) \geq 0$, with equality iff $r_i = 0$.
- *Case $r_i < 0$:* then $[r_i]_+ = 0$, so $d_i = g_i^2 - [r_i+g_i]_+^2$ and $g_i^2 - d_i = [r_i+g_i]_+^2 \geq 0$, with equality iff $r_i + g_i \leq 0$. Also $d_i \geq 0$: if $r_i + g_i \leq 0$ this is clear, and if $r_i + g_i > 0$ then $0 < r_i + g_i < g_i$, so $[r_i+g_i]_+^2 < g_i^2$.

Thus $0 \leq D \leq \|g\|_2^2$, which proves the inequality. If $\|[r+g]_+\|_2 = \|[r]_+\|_2$, then $D = \|g\|_2^2$, i.e. $g_i^2 - d_i = 0$ for every i. By the case analysis this forces: $g_i = 0$ whenever $r_i \geq 0$ and $r_i + g_i \geq 0$; $r_i = 0$ whenever $r_i \geq 0$ and $r_i + g_i < 0$; and $r_i + g_i \leq 0$ whenever $r_i < 0$. In the first case $[r_i+g_i]_+ = r_i = [r_i]_+$; in the second, $[r_i+g_i]_+ = 0 = r_i = [r_i]_+$; in the third, $[r_i+g_i]_+ = 0 = [r_i]_+$. Hence $[r+g]_+ = [r]_+$ componentwise. Finally, if $r_i > 0$ the second case is excluded (it forces $r_i = 0$), so the first case applies and $g_i = 0$. □

Proof of Proposition 1. Let $z = (\tilde r_1, \tilde r_2)$ with $\tilde r_1, \tilde r_2 \neq 0$ and let $(\tilde x, \tilde y) = (\tilde r_1/\|\tilde r_1\|_1,\ \tilde r_2/\|\tilde r_2\|_1)$. We prove: *(a)* if $(\tilde x, \tilde y)$ is a Nash equilibrium then $\Phi(z) = z$, with $\gamma_1 = v^*$ and $\gamma_2 = -v^*$; *(b)* conversely, if $\Phi(z) = z$ then $(\tilde x, \tilde y)$ is a Nash equilibrium, $\gamma_1 = v^*$, $\gamma_2 = -v^*$, and each $\tilde r_i$ is a positive multiple of the corresponding equilibrium strategy (in particular $\mathrm{supp}(\tilde r_1) = \mathrm{supp}(\tilde x)$ and $\mathrm{supp}(\tilde r_2) = \mathrm{supp}(\tilde y)$). Together, (a) and (b) give $\mathrm{Fix}(\Phi) \cap \{\tilde r_1, \tilde r_2 \neq 0\} = \{(c_1 x^\dagger, c_2 y^\dagger) : (x^\dagger, y^\dagger)$ a Nash equilibrium, $c_1, c_2 > 0\}$, which is the two-parameter cone E_+ when the equilibrium is unique.

We use the standard complementary-slackness facts for zero-sum games: if $(x^\dagger, y^\dagger)$ is a Nash equilibrium with value $v^* = x^{\dagger\top} A y^\dagger$, then, because $x^\dagger$ best-responds to $y^\dagger$, $\max_i (A y^\dagger)_i = v^*$ and every $i \in \mathrm{supp}(x^\dagger)$ attains this maximum; symmetrically $\min_j (x^{\dagger\top} A)_j = v^*$ and every $j \in \mathrm{supp}(y^\dagger)$ attains this minimum.

(a) Suppose $(\tilde x, \tilde y)$ is a Nash equilibrium; write $c_1 := \|\tilde r_1\|_1$, $c_2 := \|\tilde r_2\|_1$, so $\tilde r_1 = c_1 \tilde x$, $\tilde r_2 = c_2 \tilde y$. Then $m_1 = A\tilde y$, and $\gamma = v^*$ solves the γ-shift equation: for $i \in \mathrm{supp}(\tilde x)$ we have $(A\tilde y)_i = v^*$, so $(\tilde r_1 + m_1 - v^* \mathbf{1})_i = \tilde r_{1,i} > 0$; for $i \notin \mathrm{supp}(\tilde x)$ we have $\tilde r_{1,i} = 0$ and $(A\tilde y)_i \leq v^*$, so $(\tilde r_1 + m_1 - v^*\mathbf{1})_i \leq 0$. Hence $[\tilde r_1 + m_1 - v^*\mathbf{1}]_+ = \tilde r_1$, whose ℓ_2 norm is N_1; by uniqueness of the solution, $\gamma_1 = v^*$ and $p_1 = \tilde r_1$, so $x = \tilde x$. Symmetrically (with $m_2 = -A^\top \tilde x$ and the minimum conditions), $\gamma_2 = -v^*$, $p_2 = \tilde r_2$, and $y = \tilde y$. Then $\mathrm{res}_1 = A(y - \tilde y) = 0$, hence $g_1 = 0$ and $\tilde r_1' = [r_1 + 0]_+ = [r_1]_+ = p_1 = \tilde r_1$; likewise $\tilde r_2' = \tilde r_2$. So $\Phi(z) = z$.

(b) Suppose $\Phi(z) = z$, i.e. $\tilde r_1' = \tilde r_1$ and $\tilde r_2' = \tilde r_2$.

Step 1 (rigidity: the update returns the clipped shifted vector). Since $\tilde r_1' = \tilde r_1$, we have $\|[r_1 + g_1]_+\|_2 = \|\tilde r_1\|_2 = N_1 = \|p_1\|_2$. By (14), Lemma 4 applies (with $r = r_1$, $g = g_1$, noting $[r_1]_+ = p_1$) and its equality case yields $[r_1 + g_1]_+ = p_1$. Hence $\tilde r_1 = \tilde r_1' = p_1$. Symmetrically $\tilde r_2 = p_2$.

Step 2 (energy identity: pre-iterate equals main iterate). By Lemma 3 applied to player 1, $p_1 = \tilde{r}_1$ forces $\langle m_1, x\rangle - \gamma_1 = 0$, and $x = p_1/\|p_1\|_1 = \tilde{r}_1/\|\tilde{r}_1\|_1 = \tilde{x}$. Symmetrically $y = \tilde{y}$.

Step 3 (residuals vanish). From $y = \tilde{y}$: $\mathrm{res}_1 = A(y - \tilde{y}) = 0$, hence $g_1 = 0$; symmetrically $g_2 = 0$.

Step 4 (componentwise reading of the fixed-point equation). By Step 1, $\tilde{r}_1 = p_1 = [\tilde{r}_1 + m_1 - \gamma_1 \mathbf{1}]_+$. Componentwise: if $\tilde{r}_{1,i} > 0$ then $\tilde{r}_{1,i} + m_{1,i} - \gamma_1 = \tilde{r}_{1,i}$, i.e. $m_{1,i} = \gamma_1$; if $\tilde{r}_{1,i} = 0$ then $\tilde{r}_{1,i} + m_{1,i} - \gamma_1 \le 0$, i.e. $m_{1,i} \le \gamma_1$. Since $y = \tilde{y}$, $m_1 = A\tilde{y} = Ay$; therefore, for every strategy $x' \in \Delta_{n_1}$, $x'^\top A\tilde{y} = \langle m_1, x'\rangle \le \gamma_1$, with equality at $x' = \tilde{x}$ (which is supported on $\{i : \tilde{r}_{1,i} > 0\}$): $\tilde{x}$ is a best response to $\tilde{y}$ and $\gamma_1 = \tilde{x}^\top A\tilde{y}$. Symmetrically, for player 2, $(-A^\top \tilde{x})_j = \gamma_2$ on $\mathrm{supp}(\tilde{r}_2)$ and $\le \gamma_2$ elsewhere, so $\tilde{y} \in \arg\max_{y'} (-\tilde{x}^\top A y') = \arg\min_{y'} \tilde{x}^\top A y'$ and $\gamma_2 = -\tilde{x}^\top A\tilde{y}$. Hence $(\tilde{x}, \tilde{y})$ is a Nash equilibrium; since all equilibria of a zero-sum game share the value, $\gamma_1 = \tilde{x}^\top A\tilde{y} = v^*$ and $\gamma_2 = -v^*$. Finally $\tilde{r}_1 = \|\tilde{r}_1\|_1\, \tilde{x}$ and $\tilde{r}_2 = \|\tilde{r}_2\|_1\, \tilde{y}$ by definition of the pre-iterates, which is the claimed cone structure, and $\mathrm{supp}(\tilde{r}_i)$ equals the support of the corresponding equilibrium strategy. □

Remark 2 (A trajectory-free proof). The proof above is purely algebraic: it inspects a single round at a single state and never invokes summability along trajectories, whereas the telescoping machinery of Zhang et al. (2025a) yields $\langle m, x\rangle - \gamma \to 0$ only along a trajectory. The exact identity (15) is of independent interest. The bound (16) is exactly the per-round scalar lower bound (4) on which the saturation analysis rests, and it is used verbatim in the proof of Lemma 1 below. Its equality case pinpoints the fixed points without any limiting argument.

B.2 Saturation equivalence: proof of Lemma 1

Remark 3 (Tighter constant). Zhang et al. (2025a, Theorem 3.2)'s element-wise inequality yields (4) and (3) with $1/(2n_i)$; we keep the self-contained $1/(8n_i)$ of Appendix B.2.

The saturation analysis uses one elementary fact.

Lemma 5 (Norm monotonicity)**.** *Let $r \in \mathbb{R}^n$ and $g \in \mathbb{R}^n$ satisfy $\langle g, [r]_+\rangle = 0$. Then $\|[r+g]_+\|_2 \ge \|[r]_+\|_2$. Consequently, along any trajectory of* IREG-PRM$^+$*, $N(t{+}1) \ge N(t)$.*

Proof. The inequality is the first claim of Lemma 4. Along one round of IREG-PRM$^+$, $\tilde{r}_i^{(t+1)} = \left[r_i^{(t)} + g_i^{(t)}\right]_+$ with $\langle g_i^{(t)}, \left[r_i^{(t)}\right]_+\rangle = 0$ (F1), and the γ-shift is defined so that $\|\left[r_i^{(t)}\right]_+\|_2 = \|\tilde{r}_i^{(t)}\|_2$ (Definition 1); the inequality therefore reads $N_i(t{+}1) \ge N_i(t)$. □

The proof below is self-contained: it uses only the energy identity (Lemma 3), the structural facts (F2)–(F3), norm monotonicity (Lemma 5), and the elementary nonnegativity of the summed external regrets in a zero-sum game (cf. Zhang et al., 2025a, Fact 2.2), which we reprove in one line.

Proof of Lemma 1. We write the argument for player 1; player 2 is symmetric (replace A by $-A^\top$). Set $u^{(t)} := Ay^{(t)}$, the payoff vector at the opponent's main iterate, and $\beta_t^1 = \langle m_1^{(t)}, x^{(t)}\rangle - \gamma_1^{(t)}$. For the trivial rounds $t < t_1^*$ (where $\tilde{r}_1^{(t)} = 0$) the algorithm sets $m_1^{(t)} = 0$, applies no shift ($\gamma_1^{(t)} = 0$), and plays $x^{(t)} = \tilde{x}^{(t)}$; hence $\beta_t^1 = 0$ and $\|x^{(t)} - \tilde{x}^{(t)}\| = 0$ on those rounds, consistently with the convention in the statement.

Step 1 (per-round lower bound; first part of (4)*).* For $t \ge t_1^*$, apply Lemma 3 at the state $\tilde{r} = \tilde{r}_1^{(t)}$ with $m = m_1^{(t)}$ and $n = n_1$: by (15) and (16),

$$\beta_t^1 \;=\; \frac{\|p_1^{(t)} - \tilde{r}_1^{(t)}\|_2^2}{2\|p_1^{(t)}\|_1} \;\ge\; \frac{N_1(t)}{8n_1}\, \|x^{(t)} - \tilde{x}^{(t)}\|_2^2 \;\ge\; 0.$$

Step 2 (entrywise telescoping). Since $[v]_+ \ge v$ entrywise, the update in (13) satisfies, entrywise,

$$\tilde{r}_1^{(t+1)} = \left[r_1^{(t)} + g_1^{(t)}\right]_+ \;\ge\; \tilde{r}_1^{(t)} + m_1^{(t)} + g_1^{(t)} - \gamma_1^{(t)}\mathbf{1}.$$

Substituting $m_1^{(t)} + g_1^{(t)} = u^{(t)} - \langle \mathrm{res}_1^{(t)}, x^{(t)} \rangle \mathbf{1}$ and $\gamma_1^{(t)} + \langle \mathrm{res}_1^{(t)}, x^{(t)} \rangle = \langle u^{(t)}, x^{(t)} \rangle - \beta_t^1$ (both identities are immediate from $\mathrm{res}_1^{(t)} = u^{(t)} - m_1^{(t)}$ and the definition of β_t^1),

$$\tilde{r}_1^{(t+1)} \;\geq\; \tilde{r}_1^{(t)} + \big(u^{(t)} - \langle u^{(t)}, x^{(t)} \rangle \mathbf{1}\big) + \beta_t^1 \, \mathbf{1};$$

the same inequality holds on the trivial rounds (there $\tilde{r}_1^{(t+1)} = \Big[g_1^{(t)}\Big]_+ \geq u^{(t)} - \langle u^{(t)}, x^{(t)} \rangle \mathbf{1}$ and $\beta_t^1 = 0$). Telescoping from $\tilde{r}_1^{(1)} = \mathbf{0}$ and evaluating at the coordinate a attaining the external regret $\mathrm{Reg}_1^{(T)} := \max_{a'} \sum_{t \leq T} \big(u_{a'}^{(t)} - \langle u^{(t)}, x^{(t)} \rangle\big)$, and using $\tilde{r}_1^{(T+1)}[a] \leq \|\tilde{r}_1^{(T+1)}\|_2 = N_1(T{+}1)$,

$$N_1(T{+}1) \;\geq\; \mathrm{Reg}_1^{(T)} + \sum_{t \leq T} \beta_t^1. \tag{17}$$

Step 3 (summing the players; second part of (4)*).* With $u_2^{(t)} = -A^\top x^{(t)}$, the two external regrets add up to

$$\mathrm{Reg}_1^{(T)} + \mathrm{Reg}_2^{(T)} = \max_{x'} \sum_{t \leq T} x'^\top A y^{(t)} - \min_{y'} \sum_{t \leq T} x^{(t)\top} A y' \;\geq\; T v^* - T v^* = 0,$$

where any equilibrium (x°, y°) can be used in the two optimizations ($x^{\circ\top} A y \geq v^*$ for every y, and $x^\top A y^\circ \leq v^*$ for every x); this is Fact 2.2 of Zhang et al. (2025a). Adding (17) for the two players and rearranging, for every T,

$$\sum_{t \leq T} \big(\beta_t^1 + \beta_t^2\big) \;\leq\; N_1(T{+}1) + N_2(T{+}1) - \big(\mathrm{Reg}_1^{(T)} + \mathrm{Reg}_2^{(T)}\big) \;\leq\; N_{1,\infty} + N_{2,\infty}$$

under two-sided saturation; since every term is nonnegative (Step 1), the series converges, proving (4).

Step 4 (second-order path length; (3) *and* $\Rightarrow$*).* For $t \geq t_1^*$ with $t_1^* < \infty$, Lemma 5 gives $N_1(t) \geq N_1(t_1^*) > 0$, so Step 1 yields $\|x^{(t)} - \tilde{x}^{(t)}\|^2 \leq 8 n_1 \, \beta_t^1 / N_1(t_1^*)$; rounds before t_1^* contribute zero. Summing and using $\sum_t \beta_t^1 \leq N_{1,\infty} + N_{2,\infty}$ (Step 3, with $\beta_t^2 \geq 0$) gives (3), and in particular the finiteness of the second-order path length under saturation.

Step 5 ($\Leftarrow$). If $t_1^* = \infty$ then $N_1 \equiv 0$ saturates trivially. Otherwise, for $t \geq t_1^*$ the prediction is $m_1^{(t)} = A\tilde{y}^{(t)}$ and (F2) gives $\|g_1^{(t)}\| \leq (1 + \sqrt{n_1}) \|A\|_{\mathrm{op}} \|y^{(t)} - \tilde{y}^{(t)}\|$ (on the finitely many rounds $t < t_1^*$ the prediction is 0 and (F2) is not used), so $\sum_{t \geq t_1^*} \|g_1^{(t)}\|^2 < \infty$ when the second-order path length is finite; by (F3), valid for $t \geq t_1^*$, and $D(t) \geq 0$, $N_1(T{+}1)^2 = N_1(t_1^*)^2 + \sum_{t_1^* \leq t \leq T} \big(\|g_1^{(t)}\|^2 - D(t)\big) \leq N_1(t_1^*)^2 + \sum_{t \geq t_1^*} \|g_1^{(t)}\|^2 < \infty$, and the nondecreasing (Lemma 5) bounded sequence $N_1(t)$ saturates; the same holds for player 2. □

Proof of Theorem 1. Write $e_1(t) := \|x^{(t)} - \tilde{x}^{(t)}\|$, $e_2(t) := \|y^{(t)} - \tilde{y}^{(t)}\|$, $S(t) := N_1(t) + N_2(t)$ and $Q(t) := N_1(t)^2 + N_2(t)^2 \geq S(t)^2/2$. Three of the inequalities above hold for every finite horizon, without any saturation hypothesis. *(i)* For $t \geq \tau$, Step 1 and Lemma 5 ($N_i(t) \geq N_i(\tau)$) give $\beta_t^i \geq \alpha_i \, e_i(t)^2$. *(ii)* The first inequality of Step 3 holds for every T: $\sum_{t \leq T} (\beta_t^1 + \beta_t^2) \leq S(T{+}1)$, since $\mathrm{Reg}_1^{(T)} + \mathrm{Reg}_2^{(T)} \geq 0$. *(iii)* For $t \geq \tau$ both predictions are $m_i^{(t)} = A\tilde{y}^{(t)}$, $-A^\top \tilde{x}^{(t)}$, so (F3) with $D(t) \geq 0$ and (F2) give $N_1(t{+}1)^2 - N_1(t)^2 \leq \|g_1^{(t)}\|^2 \leq K_1^2 \, e_2(t)^2$ and symmetrically $N_2(t{+}1)^2 - N_2(t)^2 \leq K_2^2 \, e_1(t)^2$. Summing (iii) from τ to T and inserting (i) then (ii),

$$\begin{aligned} Q(T{+}1) \;&\leq\; Q_\tau + K_1^2 \sum_{t=\tau}^{T} e_2(t)^2 + K_2^2 \sum_{t=\tau}^{T} e_1(t)^2 \\ &\leq\; Q_\tau + \frac{K_1^2}{\alpha_2} \sum_{t \leq T} \beta_t^2 + \frac{K_2^2}{\alpha_1} \sum_{t \leq T} \beta_t^1 \;\leq\; Q_\tau + M \, S(T{+}1). \end{aligned}$$

With $Q(T{+}1) \geq S(T{+}1)^2/2$ this is the quadratic inequality $S^2 - 2MS - 2Q_\tau \leq 0$ in $S = S(T{+}1)$, whence $S(T{+}1) \leq M + \sqrt{M^2 + 2Q_\tau}$ for every $T \geq \tau$; for $t \leq \tau$ the bound holds by monotonicity.

Each N_i is nondecreasing (Lemma 5) and bounded, hence convergent. If $t_1^* = \infty$, then $\tilde{r}_1^{(t)} = \mathbf{0}$ and $x^{(t)} = \tilde{x}^{(t)}$ for all t; hence $\mathrm{res}_2^{(t)} = -A^\top(x^{(t)} - \tilde{x}^{(t)}) = 0$ and $g_2^{(t)} = 0$ for $t \geq t_2^*$, so by (F3) N_2 is constant from t_2^* on (and $N_2 \equiv 0$ if also $t_2^* = \infty$). □

Remark 4 (The bound is loose but structural). On the testbed of Section 5 the right-hand side of (5) exceeds the realized $N_{1,\infty} + N_{2,\infty}$ by 10^3–10^6 (it grows like $n^2\|A\|^2/N_i(\tau)$), and none of the three finite-horizon inequalities (i)–(iii) is violated at any round on 36 instances spanning all cells. Its role is qualitative: it removes saturation from the hypothesis list of every result in this paper.

B.3 Proofs

This subsection proves the three results of this section: the absorption lemma, the closed-form Jacobian together with the value-centering identities, and the spectral geometry that produces the closed form (9) of Theorem 4. The two remaining steps of the proof of Theorem 4, a uniform second-order Taylor expansion and the bootstrap that converts the spectral picture into the rate (11), are standard in technique and are carried out in Appendix B.4; the algebraic proof of Proposition 1 and the saturation equivalence of Lemma 1 are in Appendix B.

B.3.1 One-step absorption: proof of Lemma 6

Lemma 6 (One-step absorption and local analyticity)**.** *Let (x^*, y^*) be a strictly complementary equilibrium with margin $\delta > 0$ as in* (8)*, and let $z_f = (c_1 x^*, c_2 y^*)$, $c_1, c_2 > 0$, be any point of the corresponding ray (a fixed point of Φ by Proposition 1). There is an open neighborhood $\mathcal{U}_0 \subset \mathbb{R}^{n_1+n_2}$ of z_f such that:*

- *(i) $\Phi(\mathcal{U}_0 \cap \mathcal{Z}) \subseteq \mathcal{M}_S \cap \mathcal{Z}$; more precisely, for every $z \in \mathcal{U}_0 \cap \mathcal{Z}$ the off-support coordinates of $\Phi(z)$ are* exactly *zero;*
- *(ii) on $\mathcal{U}_0 \cap \mathcal{M}_S \cap \mathcal{Z}$ the clipping pattern of every $[\cdot]_+$ operation in one round is constant, and $\Phi_S := \Phi|_{\mathcal{M}_S}$ is real-analytic there (so $\mathcal{M}_S$ is forward-invariant near z_f; it need not be invariant globally).*

Proof of Lemma 6. Let $z_f = (c_1 x^*, c_2 y^*)$ with $c_1, c_2 > 0$. By Proposition 1, $\Phi(z_f) = z_f$ with $\gamma_1 = v^*$, $\gamma_2 = -v^*$, and at z_f we have $x = \tilde{x} = x^*$, $y = \tilde{y} = y^*$, $p_1 = \tilde{r}_1$, $p_2 = \tilde{r}_2$, and $g_1 = g_2 = 0$.

Step 1 (continuity of the round at z_f). On a neighborhood of z_f the pre-iterate normalizations are continuous (denominators $\|\tilde{r}_i\|_1 \geq c_i/2$), hence so are m_1 and m_2. The multiplier γ_1 depends continuously on its inputs: the map $\varphi(\gamma; v) := \|[v - \gamma\mathbf{1}]_+\|_2$ is jointly continuous and, at the solution (where $\varphi = N_1 > 0$), strictly decreasing in γ; a standard monotone-crossing argument then gives continuity of the solution in (v, N_1). Consequently p_1, x, res_1, and g_1 are continuous near z_f (the denominator $\|p_1\|_1 \geq \|p_1\|_2 = N_1$ is bounded below), and $g_1 \to 0$, $\gamma_1 \to v^*$ as $z \to z_f$; the same holds for player 2.

Step 2 (off-support coordinates are clipped exactly). Fix $i \notin S_1$. As $z \to z_f$,

$$r_{1,i} = \tilde{r}_{1,i} + m_{1,i} - \gamma_1 \;\longrightarrow\; 0 + (Ay^*)_i - v^* \;\leq\; -\delta,$$

by strict complementarity (8). Choose the neighborhood $\mathcal{U}_0$ small enough that, for all $z \in \mathcal{U}_0$ and all $i \notin S_1$, $r_{1,i} \leq -\delta/2$ and $\|g_1\|_\infty \leq \delta/4$. Then $(r_1 + g_1)_i \leq -\delta/4 < 0$, so $\tilde{r}'_{1,i} = [(r_1 + g_1)_i]_+ = 0$ *exactly* (and also $p_{1,i} = 0$). The same holds for player 2, using the second minimum in (8). Hence $\Phi(\mathcal{U}_0) \subseteq \mathcal{M}_S$, proving (i).

Step 3 (support coordinates are never clipped). Fix $i \in S_1$. As $z \to z_f$, $r_{1,i} \to c_1 x_i^* > 0$. Shrinking $\mathcal{U}_0$, we may assume $r_{1,i} \geq \kappa := \frac{c_1}{2}\min_{i \in S_1} x_i^* > 0$ and $\|g_1\|_\infty \leq \kappa/2$ for all $z \in \mathcal{U}_0$, so $(r_1 + g_1)_i \geq \kappa/2 > 0$: neither the clipped shifted vector nor the update clips any support coordinate. Symmetrically for player 2.

Step 4 (frozen pattern and analyticity on $\mathcal{U}_0 \cap \mathcal{M}_S$). Set $\iota := \min\{\delta/8, \kappa/4\}$ and shrink $\mathcal{U}_0$ once more so that $|\gamma_1 - v^*| \leq \iota$ on $\mathcal{U}_0$ (possible by Step 1). Then, for $z \in \mathcal{U}_0 \cap \mathcal{M}_S$ and every γ with $|\gamma - v^*| \leq \iota$ (an interval that contains the solution γ_1, and $|\gamma - \gamma_1| \leq 2\iota$ throughout), the vector $\tilde{r}_1 + m_1 - \gamma\mathbf{1}$ is strictly positive exactly on S_1 (entries $\geq \kappa - 2\iota \geq \kappa/2$, by Step 3) and strictly

negative off S_1 (entries $\le -\delta/2 + 2\iota \le -\delta/4$, by Step 2); hence on $\mathcal{U}_0 \cap \mathcal{M}_S$ the defining equation for γ_1 reads

$$\sum_{i \in S_1} \big(\tilde{r}_{1,i} + m_{1,i} - \gamma_1\big)^2 \;=\; \sum_{i \in S_1} \tilde{r}_{1,i}^2\,, \tag{18}$$

a quadratic equation in γ_1 whose left-hand side has γ-derivative $-2\sum_{i\in S_1} r_{1,i} \le -2\kappa < 0$ at the solution. By the implicit function theorem, γ_1 is a real-analytic function of $(\tilde{r}_1|_{S_1}, m_1|_{S_1})$, and $m_1|_{S_1} = \tilde{A}\,\tilde{y}|_{S_2}$ is real-analytic in $\tilde{r}_2|_{S_2}$. All remaining operations of the round (normalizations with denominators bounded below, matrix products, the update with the clipping pattern frozen by Steps 2–3) are real-analytic on $\mathcal{U}_0 \cap \mathcal{M}_S$. Hence $\Phi_S = \Phi|_{\mathcal{M}_S}$ is a composition of real-analytic maps on $\mathcal{U}_0 \cap \mathcal{M}_S$, proving (ii). □

B.3.2 The Jacobian on the frozen face: proofs of Lemma 9 and Lemma 10

Throughout this part Assumption 3 is in force at the designated equilibrium (x^*, y^*); the computation uses only the kernel condition, not uniqueness (which it implies, Lemma 11). We work in face coordinates $z = (\tilde{r}_1, \tilde{r}_2) \in \mathbb{R}^{S_1} \times \mathbb{R}^{S_2}$ on $\mathcal{U}_0 \cap \mathcal{M}_S$, where $\mathcal{U}_0$ is the neighborhood of a fixed point $z_f = (a\xi, b\psi)$, $a, b > 0$, provided by Lemma 6. With the clipping pattern frozen, one round of the face dynamics reads

$$\tilde{r}_1' \;=\; r_1 + g_1 \;=\; \tilde{r}_1 + \tilde{A}\,y - \big(\gamma_1 + \langle \mathrm{res}_1, x\rangle\big)\mathbf{1}, \qquad \tilde{r}_2' \;=\; \tilde{r}_2 - \tilde{A}^\top x - \big(\gamma_2 + \langle \mathrm{res}_2, y\rangle\big)\mathbf{1}, \tag{19}$$

where $r_1 = \tilde{r}_1 + m_1 - \gamma_1\mathbf{1} > 0$, $x = r_1/\|r_1\|_1$, $m_1 = \tilde{A}\tilde{y}$, $\mathrm{res}_1 = \tilde{A}(y - \tilde{y})$ (so that $m_1 + \mathrm{res}_1 = \tilde{A}y$), and symmetrically $r_2 = \tilde{r}_2 + m_2 - \gamma_2\mathbf{1} > 0$, $y = r_2/\|r_2\|_1$, $m_2 = -\tilde{A}^\top\tilde{x}$, $\mathrm{res}_2 = -\tilde{A}^\top(x - \tilde{x})$. At the fixed point z_f: $\gamma_1 = v^*$, $\gamma_2 = -v^*$, $r_1 = \tilde{r}_1 = a\xi$, $r_2 = \tilde{r}_2 = b\psi$, $x = \tilde{x} = \xi$, $y = \tilde{y} = \psi$, and $\mathrm{res}_1 = \mathrm{res}_2 = 0$ (Proposition 1 restricted to the face). Recall also the differential of the normalization $u \mapsto u/\langle \mathbf{1}, u\rangle$ (for $u > 0$): at $u = a\xi$ it is $du \mapsto P_\xi\, du/a$, and at $u = b\psi$ it is $du \mapsto P_\psi\, du/b$.

Lemma 7 (Vanishing of the multiplier differentials)**.** *At any fixed point $z_f = (a\xi, b\psi)$, $a, b > 0$, the differentials of γ_1 and γ_2 along $\mathcal{M}_S$ satisfy $d\gamma_1 = \langle \xi, dm_1\rangle$ and $d\gamma_2 = \langle \psi, dm_2\rangle$, and both vanish: $d\gamma_1 = d\gamma_2 = 0$.*

Proof. Differentiating the frozen-pattern multiplier equation (18) gives

$$\langle r_1,\; d\tilde{r}_1 + dm_1 - d\gamma_1\mathbf{1}\,\rangle \;=\; \langle \tilde{r}_1, d\tilde{r}_1\rangle.$$

At z_f we have $r_1 = \tilde{r}_1 = a\xi$, so the terms in $d\tilde{r}_1$ cancel *exactly* (this is the cancellation between the gradient $\nabla_v \gamma = x$ of the shift in its first argument and the derivative of the norm constraint), leaving $a\langle \xi, dm_1\rangle = d\gamma_1 \cdot a\langle \xi, \mathbf{1}\rangle = a\, d\gamma_1$, i.e. $d\gamma_1 = \langle \xi, dm_1\rangle$. Now $dm_1 = \tilde{A}\, d\tilde{y}$ with $d\tilde{y} = P_\psi\, d\tilde{r}_2/b$, hence

$$d\gamma_1 = \tfrac{1}{b}\,\xi^\top \tilde{A} P_\psi\, d\tilde{r}_2 = \tfrac{1}{b}\, v^*\, \mathbf{1}^\top P_\psi\, d\tilde{r}_2 = 0,$$

using the equilibrium indifference identity $\xi^\top\tilde{A} = v^*\mathbf{1}^\top$ (Lemma 10) and $\mathbf{1}^\top P_\psi = \mathbf{1}^\top - (\mathbf{1}^\top\psi)\mathbf{1}^\top = 0$. Symmetrically, $d\gamma_2 = \langle \psi, dm_2\rangle = -\frac{1}{a}\,\psi^\top\tilde{A}^\top P_\xi\, d\tilde{r}_1 = -\frac{1}{a}\, v^*\,\mathbf{1}^\top P_\xi\, d\tilde{r}_1 = 0$. □

Lemma 8 (Vanishing of the centering differentials)**.** *At any fixed point $z_f = (a\xi, b\psi)$, $a, b > 0$, $d\langle \mathrm{res}_1, x\rangle = 0$ and $d\langle \mathrm{res}_2, y\rangle = 0$.*

Proof. On the face, $\langle \mathrm{res}_1, x\rangle = x^\top\tilde{A}\,(y - \tilde{y})$. At z_f we have $y - \tilde{y} = 0$ and $x = \xi$, so by the product rule

$$d\langle \mathrm{res}_1, x\rangle = (dx)^\top\tilde{A}\,(y - \tilde{y})\big|_{z_f} + \xi^\top\tilde{A}\,(dy - d\tilde{y}) = 0 + v^*\,\mathbf{1}^\top(dy - d\tilde{y}) = 0,$$

since y and $\tilde{y}$ are simplex-valued analytic maps, whence $\mathbf{1}^\top dy = \mathbf{1}^\top d\tilde{y} = 0$. For player 2, $\langle \mathrm{res}_2, y\rangle = -y^\top\tilde{A}^\top(x - \tilde{x})$, and the same argument with $\tilde{A}\psi = v^*\mathbf{1}$ gives $d\langle \mathrm{res}_2, y\rangle = -\psi^\top\tilde{A}^\top(dx - d\tilde{x}) = -v^*\mathbf{1}^\top(dx - d\tilde{x}) = 0$. □

Lemma 9 (Closed-form Jacobian; $J = I + L + L^2$)**.** *Let Assumption 3 hold at (x^*, y^*). At any fixed point $z_f = (a\xi, b\psi)$ with $a, b > 0$,*

$$J \;:=\; D\Phi_S(z_f) \;=\; I + L + L^2, \qquad L \;=\; \begin{pmatrix} 0 & U/b \\ -V/a & 0 \end{pmatrix}, \qquad U := \tilde{A}P_\psi, \quad V := \tilde{A}^\top P_\xi, \tag{20}$$

and consequently $L^2 = -\mathrm{diag}(UV,\, VU)/(ab)$. The derivation rests on two exact cancellations at z_f: the differentials of the shift multipliers vanish, $d\gamma_1 = d\gamma_2 = 0$, and so do the differentials of the centering terms.

Proof of Lemma 9. Differentiate (19) at z_f. By Lemma 7 and Lemma 8 the differentials of the two scalar terms vanish, so

$$d\tilde{r}_1' = d\tilde{r}_1 + \tilde{A}\, dy, \qquad d\tilde{r}_2' = d\tilde{r}_2 - \tilde{A}^\top dx.$$

For dy: at z_f, $r_2 = b\psi$, so $dy = P_\psi\, dr_2/b$ with $dr_2 = d\tilde{r}_2 + dm_2 - d\gamma_2 \mathbf{1} = d\tilde{r}_2 - \frac{1}{a}\tilde{A}^\top P_\xi\, d\tilde{r}_1$ (using $dm_2 = -\tilde{A}^\top d\tilde{x}$, $d\tilde{x} = P_\xi\, d\tilde{r}_1/a$, and $d\gamma_2 = 0$). Hence, with $U = \tilde{A}P_\psi$ and $V = \tilde{A}^\top P_\xi$,

$$d\tilde{r}_1' = d\tilde{r}_1 + \tfrac{1}{b}\, U\, d\tilde{r}_2 - \tfrac{1}{ab}\, UV\, d\tilde{r}_1.$$

For dx: $dx = P_\xi\, dr_1/a$ with $dr_1 = d\tilde{r}_1 + dm_1 - d\gamma_1 \mathbf{1} = d\tilde{r}_1 + \frac{1}{b}\,\tilde{A}P_\psi\, d\tilde{r}_2 = d\tilde{r}_1 + \frac{1}{b}\, U\, d\tilde{r}_2$, so

$$d\tilde{r}_2' = d\tilde{r}_2 - \tfrac{1}{a}\, V\, d\tilde{r}_1 - \tfrac{1}{ab}\, VU\, d\tilde{r}_2.$$

These two displays are exactly the block rows of $(I + L + L^2)\,(d\tilde{r}_1, d\tilde{r}_2)$ with L as in (20), since $L^2 = -\mathrm{diag}(UV, VU)/(ab)$. □

Lemma 10 (Value-centering identities; $U = B$, $V = B^\top$)**.** *At the equilibrium, the indifference conditions $\tilde{A}\psi = v^*\mathbf{1}$ and $\tilde{A}^\top\xi = v^*\mathbf{1}$ hold, and therefore*

$$U \;=\; \tilde{A}P_\psi \;=\; \tilde{A} - v^*\mathbf{1}\mathbf{1}^\top \;=\; B, \qquad V \;=\; \tilde{A}^\top P_\xi \;=\; B^\top. \tag{21}$$

Moreover $B\psi = 0$ and $B^\top\xi = 0$; under Assumption 3 these span the kernels.

Proof of Lemma 10. Since every $i \in S_1 = \mathrm{supp}(x^*)$ is a best response to y^*, $(Ay^*)_i = v^*$ for $i \in S_1$, i.e. $\tilde{A}\psi = v^*\mathbf{1}$; symmetrically $\tilde{A}^\top\xi = v^*\mathbf{1}$. Hence

$$U = \tilde{A}P_\psi = \tilde{A} - (\tilde{A}\psi)\mathbf{1}^\top = \tilde{A} - v^*\mathbf{1}\mathbf{1}^\top = B, \qquad V = \tilde{A}^\top P_\xi = \tilde{A}^\top - (\tilde{A}^\top\xi)\mathbf{1}^\top = B^\top.$$

Moreover $B\psi = \tilde{A}\psi - v^*(\mathbf{1}^\top\psi)\mathbf{1} = v^*\mathbf{1} - v^*\mathbf{1} = 0$ and, symmetrically, $B^\top\xi = 0$. That these vectors *span* the kernels is exactly the non-degeneracy requirement of Assumption 3; under Assumption 1 it holds automatically by Lemma 11 below. Finally, $UV/(ab) = BB^\top/(ab)$ is a Gram matrix up to the positive scalar $1/(ab)$, hence symmetric positive semidefinite with real nonnegative spectrum; the identity $B^\top\xi = 0$ places the equilibrium direction ξ in its kernel, and likewise $B\psi = 0$ places ψ in the kernel of $VU/(ab) = B^\top B/(ab)$. □

B.3.3 Spectral geometry: Step A of the proof of Theorem 4

Standing notation. Assumption 1 is in force. By Lemma 11 below, Assumption 3 then holds at (x^*, y^*), so Lemma 6, Lemma 9 and Lemma 10 apply, and $|S_1| = |S_2| =: s$; the face has dimension $n_S := 2s$. Let $B = \sum_{k=1}^{s-1} \sigma_k u_k v_k^\top$ be a reduced singular value decomposition with $\sigma_1 \geq \cdots \geq \sigma_{s-1} > 0$; write $\sigma_{\max} = \sigma_1$ and $\sigma_{\min}^+ = \sigma_{s-1}$. Set $e_1 = (\xi, 0)$, $e_2 = (0, \psi)$, $E = \mathrm{span}\{e_1, e_2\}$, $E_+ = \{ae_1 + be_2 : a, b > 0\}$, and $z^* = c_1e_1 + c_2e_2$ with $c_1, c_2 > 0$ the reference fixed point; $\eta = \eta(a,b) = 1/\sqrt{ab}$, $\eta_\infty = \eta(c_1, c_2)$, and $\rho(a,b) = \max_k \sqrt{1 - \eta^2\sigma_k^2(1 - \eta^2\sigma_k^2)}$, $\rho^* = \rho(c_1, c_2)$. Closed Euclidean balls in face coordinates are denoted $\mathbb{B}(z, r)$. By Proposition 1 and uniqueness of the equilibrium, $\mathrm{Fix}(\Phi_S)$ coincides with E_+ in a neighborhood of any of its points, and Φ_S fixes E_+ pointwise.

The proofs use uniqueness only through a kernel condition on the support matrix, stated separately:

Assumption 3 (Strictly complementary, non-degenerate equilibrium ray)**.** *The designated Nash equilibrium (x^*, y^*) is* strictly complementary *($\delta > 0$) and* non-degenerate*: $\ker B = \mathrm{span}\{y^*|_{S_2}\}$ and $\ker B^\top = \mathrm{span}\{x^*|_{S_1}\}$, equivalently $\mathrm{rank}\, B = |S_1| - 1 = |S_2| - 1$; in particular $|S_1| = |S_2|$.*

Remark 5 (One hypothesis, two forms). By Lemma 11 the two assumptions are equivalent, so a game with several equilibria has no strictly complementary non-degenerate one.

Lemma 11 (Uniqueness ⇔ non-degeneracy: Assumption 1 ⇔ Assumption 3)**.** *Let (x^*, y^*) be a strictly complementary equilibrium. Then it is non-degenerate,* $\ker B = \operatorname{span}\{\psi\}$ *and* $\ker B^\top = \operatorname{span}\{\xi\}$*, if and only if it is that game's unique equilibrium. In particular, under either assumption* $\operatorname{rank} B = |S_1| - 1 = |S_2| - 1$*, so $|S_1| = |S_2|$; and in a game with several equilibria no equilibrium is both strictly complementary and non-degenerate.*

Proof. Uniqueness ⇒ non-degeneracy. Let $d \in \ker B$, i.e. $\tilde{A}d = v^*(\mathbf{1}^\top d)\mathbf{1}$.

Case 1: $\mathbf{1}^\top d = 0$ *and* $d \notin \operatorname{span}\{\psi\}$ *(so $d \neq 0$).* Then $\tilde{A}d = 0$. For $t \in \mathbb{R}$, let y_t denote the vector $\psi + td$ on S_2, extended by zero off S_2. Since $\psi > 0$ on S_2 and $\mathbf{1}^\top(\psi + td) = 1$, for $|t|$ small enough y_t is a strategy with $\operatorname{supp}(y_t) \subseteq S_2$. We claim (x^*, y_t) is a Nash equilibrium: (i) on S_1, $\tilde{A}(\psi + td) = v^*\mathbf{1}$, while for $i \notin S_1$, $(Ay_t)_i = (Ay^*)_i + t\,(A[i, S_2]\,d) \leq v^* - \delta + O(|t|) < v^*$ for $|t|$ small; hence $\max_i (Ay_t)_i = v^*$, attained on $S_1 \supseteq \operatorname{supp}(x^*)$, so x^* is a best response to y_t. (ii) $(x^{*\top}A)_j = v^*$ for all $j \in S_2$ (indifference) and $\geq v^* + \delta$ off S_2, so $\min_j (x^{*\top}A)_j = v^*$ and every strategy supported in S_2, in particular y_t, is a best response to x^* for the minimizing player. For $t \neq 0$, $y_t \neq y^*$, contradicting uniqueness of the equilibrium.

Case 2: $\mathbf{1}^\top d \neq 0$. Normalize so that $\mathbf{1}^\top d = 1$. Then $d' := d - \psi$ satisfies $\tilde{A}d' = v^*\mathbf{1} - v^*\mathbf{1} = 0$ and $\mathbf{1}^\top d' = 0$, so $d' \in \ker B$; moreover $d' \notin \operatorname{span}\{\psi\}$ unless $d' = 0$ (any nonzero multiple of ψ has nonzero coordinate sum). If $d' \neq 0$, Case 1 applies to d' and gives a contradiction; hence $d' = 0$ and $d = \psi$.

Therefore $\ker B = \operatorname{span}\{\psi\}$; the statement for $B^\top$ and ξ is symmetric. The rank–nullity theorem gives $\operatorname{rank} B = |S_2| - 1$ and $\operatorname{rank} B^\top = |S_1| - 1$; equality of the two ranks forces $|S_1| = |S_2|$.

Non-degeneracy ⇒ uniqueness. Let $(x^\dagger, y^\dagger)$ be any equilibrium. Since x^* guarantees v^* ($x^{*\top}Ay \geq v^*$ for all y) and $y^\dagger$ guarantees v^* ($x^\top Ay^\dagger \leq v^*$ for all x), $x^{*\top}Ay^\dagger = v^*$, so $(x^*, y^\dagger)$ is an equilibrium. First, $y^\dagger$ is a best response to x^*: $\sum_j y_j^\dagger \big((x^{*\top}A)_j - v^*\big) = 0$ with every term nonnegative and, by (8), strictly positive off S_2, so $\operatorname{supp}(y^\dagger) \subseteq S_2$. Second, x^* is a best response to $y^\dagger$ with $\max_i (Ay^\dagger)_i = v^*$, so every $i \in \operatorname{supp}(x^*) = S_1$ attains the maximum: $\tilde{A}\,y^\dagger|_{S_2} = v^*\mathbf{1}$. Hence $d := y^\dagger|_{S_2} - \psi$ satisfies $\tilde{A}d = 0$ and $\mathbf{1}^\top d = 0$, so $Bd = \tilde{A}d - v^*(\mathbf{1}^\top d)\mathbf{1} = 0$ and $d \in \ker B = \operatorname{span}\{\psi\}$; as $\mathbf{1}^\top\psi = 1 \neq 0 = \mathbf{1}^\top d$, $d = 0$ and $y^\dagger = y^*$. The symmetric argument gives $x^\dagger = x^*$.

Both directions are consistent with the classical non-degeneracy theory of matrix games (Bohnenblust et al., 1950); the proofs above are self-contained. □

Lemma 12 (Orthogonal splitting)**.** *Define*

$$W_k := \operatorname{span}\big\{(u_k, 0),\, (0, v_k)\big\} \quad (k = 1, \ldots, s-1), \qquad W := \bigoplus_{k=1}^{s-1} W_k.$$

Then $\mathbb{R}^{n_S} = E \oplus W$ with $E \perp W$ (Euclidean orthogonality); the corresponding orthogonal projections P_E, P_W satisfy $\|P_E\| = \|P_W\| = 1$. For $w = (w_1, w_2) \in W$ one has $w_1 \in \operatorname{range}(B) = \operatorname{span}\{u_k\}$, $w_2 \in \operatorname{range}(B^\top) = \operatorname{span}\{v_k\}$, and $\|w\|^2 = \|w_1\|^2 + \|w_2\|^2$.

Proof. By Lemma 10, $B^\top\xi = 0$, i.e. $\xi \perp \operatorname{range}(B) = \operatorname{span}\{u_k\}$; hence $e_1 = (\xi, 0) \perp (u_k, 0)$ for all k, and $e_1 \perp (0, v_k)$ trivially. The statement for e_2 is symmetric (using $B\psi = 0$). The u_k are orthonormal, as are the v_k, so the W_k are mutually orthogonal and give orthonormal coordinates on W. Dimensions: $\dim E + \dim W = 2 + 2(s-1) = 2s = n_S$ (using Lemma 11), so the orthogonal sum exhausts the face. □

Lemma 13 (Uniform block diagonalization and closed-form modes)**.** *For every fixed point $z_f = (a\xi, b\psi) \in E_+$, the Jacobian $J(z_f) = I + L(a,b) + L(a,b)^2$ of Lemma 9 (with $U = B$, $V = B^\top$ by Lemma 10) satisfies:*

(1) $J(z_f)|_E = I_E$; in fact $J(z_f)e_1 = e_1$ and $J(z_f)e_2 = e_2$ exactly.

(2) Each W_k is $J(z_f)$-invariant, and in the basis $\{(u_k, 0), (0, v_k)\}$,

$$J(z_f)|_{W_k} = J_k(a,b) := I_2 + L_k + L_k^2, \qquad L_k = \begin{pmatrix} 0 & \sigma_k/b \\ -\sigma_k/a & 0 \end{pmatrix}.$$

(3) Consequently $P_E J(z_f) P_W = 0$ and $P_W J(z_f) P_E = 0$ for all $a, b > 0$*; write* $\mathcal{A}(a,b) := J(z_f)|_W = \bigoplus_k J_k(a,b)$.

(4) $\operatorname{spec}\big(J_k(a,b)\big) = \{\mu_k^{\pm}\}$ *with* $\mu_k^{\pm} = 1 - \eta^2\sigma_k^2 \pm i\,\eta\sigma_k$ *and* $|\mu_k^{\pm}|^2 = 1 - \eta^2\sigma_k^2(1-\eta^2\sigma_k^2)$*, where* $\eta = 1/\sqrt{ab}$*. Since* $\eta\sigma_k \neq 0$*, the pair is non-real and conjugate, so each* J_k *is diagonalizable over* $\mathbb{C}$ *and* $\mathcal{A}(a,b)$ *is semisimple. Moreover* $\rho(a,b) = \max_k |\mu_k^{\pm}(a,b)|$.

(5) The map $(a,b) \mapsto \mathcal{A}(a,b)$ *is real-analytic on* $\{a, b > 0\}$*, and on the rectangle* $R_\square := [c_1/2,\, 3c_1/2] \times [c_2/2,\, 3c_2/2]$ *it is Lipschitz:*

$$\|\mathcal{A}(a,b) - \mathcal{A}(a',b')\| \le C_J\big(|a-a'| + |b-b'|\big),$$
$$C_J := \sigma_{\max}\Big(\tfrac{4}{c_1^2} + \tfrac{4}{c_2^2}\Big)\Big(1 + 2\sigma_{\max}\big(\tfrac{2}{c_1} + \tfrac{2}{c_2}\big)\Big).$$

Proof. (1)–(2): direct verification. Using $Bv_k = \sigma_k u_k$ and $B^\top u_k = \sigma_k v_k$,

$$L(a,b)\,(u_k, 0) = (0,\, -B^\top u_k/a) = (0,\, -\sigma_k v_k/a) \in W_k,$$
$$L(a,b)\,(0, v_k) = (Bv_k/b,\, 0) = (\sigma_k u_k/b,\, 0) \in W_k,$$

so W_k is L-invariant with matrix L_k in the given basis, hence J-invariant with matrix $I_2 + L_k + L_k^2$. Similarly $L(a,b)\,e_1 = (0,\, -B^\top \xi/a) = 0$ (Lemma 10), so $Je_1 = e_1$; and $Le_2 = (B\psi/b,\, 0) = 0$, so $Je_2 = e_2$.

(3): immediate from (1)–(2) and the orthogonality of the splitting (Lemma 12): J maps E into E and W into W.

(4): the eigenvalues of L_k solve $\lambda^2 = -\sigma_k^2/(ab)$, i.e. $\lambda = \pm i\eta\sigma_k$; hence the eigenvalues of J_k are $\mu = 1 + \lambda + \lambda^2 = 1 - \eta^2\sigma_k^2 \pm i\eta\sigma_k$, and

$$|\mu_k^{\pm}|^2 = (1-\eta^2\sigma_k^2)^2 + \eta^2\sigma_k^2 = 1 - \eta^2\sigma_k^2\,(1 - \eta^2\sigma_k^2).$$

The pair is non-real and distinct because $\eta\sigma_k > 0$, so J_k is diagonalizable over $\mathbb{C}$; the direct sum $\mathcal{A}(a,b)$ is therefore semisimple, and its spectral radius is $\max_k |\mu_k^{\pm}|$, which equals $\rho(a,b)$ by definition.

(5): $L(a,b)$ depends affinely on $(1/a, 1/b)$, so it is real-analytic on $\{a, b > 0\}$; hence so is $\mathcal{A}$. For the Lipschitz bound, note that for the block-antidiagonal L, $\|L(a,b)\| = \max\{\sigma_{\max}/b,\ \sigma_{\max}/a\} \le \sigma_{\max}(2/c_1 + 2/c_2)$ on $R_\square$, and

$$\|L(a,b) - L(a',b')\| \le \sigma_{\max} \max\Big\{\big|\tfrac{1}{a} - \tfrac{1}{a'}\big|,\ \big|\tfrac{1}{b} - \tfrac{1}{b'}\big|\Big\} \le \sigma_{\max}\Big(\tfrac{4}{c_1^2} + \tfrac{4}{c_2^2}\Big)\big(|a-a'| + |b-b'|\big),$$

using $a, a' \ge c_1/2$ and $b, b' \ge c_2/2$. Finally $\|L^2 - L'^2\| \le (\|L\| + \|L'\|)\|L - L'\|$, so $\|\mathcal{A} - \mathcal{A}'\| \le \big(1 + 2\sigma_{\max}(\tfrac{2}{c_1} + \tfrac{2}{c_2})\big)\|L - L'\|$, which gives C_J. □

Remark 6. Items (1)–(3) hold *simultaneously at every fixed point*: the mode planes W_k are fixed once and for all by the singular vectors of B, and only the scalars σ_k/a, σ_k/b vary along the cone. By items (1) and (4), the unit eigenvalue of $J(z_f)$ is exactly double and semisimple with eigenspace E (a unit eigenvalue inside some W_k would require $\lambda + \lambda^2 = 0$, i.e. $\lambda \in \{0, -1\}$, impossible for $\lambda = \pm i\eta\sigma_k$). Thus the usual non-degeneracy hypothesis of local analyses is a theorem here, as announced in Remark 5.

Lemma 14 (Adapted norm on W)**.** *Let* $\|w\|_*^2 := \|w_1\|^2/c_1 + \|w_2\|^2/c_2$ *for* $w = (w_1, w_2) \in W$*, and set* $m_* := 1/\sqrt{\max(c_1,c_2)}$*,* $M_* := 1/\sqrt{\min(c_1,c_2)}$*,* $\kappa_* := M_*/m_* = \sqrt{\max(c_1,c_2)/\min(c_1,c_2)}$*. Then:*

(1) $\|\mathcal{A}(c_1,c_2)\|_* = \rho^*$ *(the operator norm equals the spectral radius exactly);*

(2) $m_*\|w\| \le \|w\|_* \le M_*\|w\|$ *for all* $w \in W$*;*

(3) for all $(a,b) \in R_\square$*,* $\|\mathcal{A}(a,b)\|_* \le \rho^* + \kappa_* C_J\big(|a-c_1| + |b-c_2|\big)$.

Proof. Let $S := \mathrm{diag}(\sqrt{c_1}, \sqrt{c_2})$. In the basis $\{(u_k, 0), (0, v_k)\}$ of W_k (Lemma 13(2)),

$$S^{-1} L_k(c_1, c_2) S = \begin{pmatrix} 0 & \sigma_k \sqrt{c_2}/(c_2\sqrt{c_1}) \\ -\sigma_k \sqrt{c_1}/(c_1\sqrt{c_2}) & 0 \end{pmatrix} = \eta_\infty \sigma_k \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix},$$

a skew-symmetric matrix, so $S^{-1} J_k(c_1, c_2) S = (1 - \eta_\infty^2 \sigma_k^2) I_2 + \eta_\infty \sigma_k \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} = |\mu_k|\, R(\varphi_k)$ with $R(\varphi) = \begin{pmatrix} \cos\varphi & -\sin\varphi \\ \sin\varphi & \cos\varphi \end{pmatrix}$ and $\varphi_k = \arg \mu_k^-$; the same S serves every block. For $w \in W$ with components $w_k \in W_k$ (the W_k are mutually orthogonal), $\sum_k \|S^{-1} w_k\|^2 = \|w_1\|^2/c_1 + \|w_2\|^2/c_2 = \|w\|_*^2$. Then

$$\|\mathcal{A}(c_1, c_2)\, w\|_*^2 = \sum_k \big\| \, |\mu_k| R(\varphi_k)\, S^{-1} w_k \big\|^2 = \sum_k |\mu_k|^2 \|S^{-1} w_k\|^2 \le (\rho^*)^2 \|w\|_*^2,$$

with equality when w is supported on a maximizing block; this proves (1). (2) is immediate from $\|S^{-1}\| = M_*$ and $\|S\| = 1/m_*$. For (3), for any operator T on W one has $\|T\|_* \le \kappa_* \|T\|$ (combine the two bounds of (2)), hence

$$\|\mathcal{A}(a, b)\|_* \le \|\mathcal{A}(c_1, c_2)\|_* + \kappa_* \|\mathcal{A}(a, b) - \mathcal{A}(c_1, c_2)\| \le \rho^* + \kappa_* C_J \big(|a - c_1| + |b - c_2|\big),$$

by Lemma 13(5). □

Corollary 3 (Extreme singular values determine the rate)**.** *For $\sigma \in \Sigma^+(B)$ let $q_\sigma := \eta_\infty^2 \sigma^2$, so the corresponding mode has modulus $|\mu_\sigma| = \sqrt{1 - q_\sigma(1 - q_\sigma)}$. Then:*

- *(i) ρ^* is attained at an* extreme *nonzero singular value: $\rho^* = \max\{|\mu_{\sigma_{\max}(B)}|,\ |\mu_{\sigma^+_{\min}(B)}|\}$, where $\sigma^+_{\min}(B)$ denotes the smallest nonzero singular value.*
- *(ii) Per mode, contraction is best exactly at $q_\sigma = \frac{1}{2}$, where $|\mu_\sigma| = \sqrt{3}/2$, and on the contracting range $q_\sigma \in (0, 1)$ it degrades monotonically as $q_\sigma \to 0$ or $q_\sigma \to 1$.*
- *(iii) Under Assumption 2 (every mode contracts), $\rho^* \to 1$ as $\eta_\infty\, \sigma_{\max}(B) \uparrow 1$ or as $\eta_\infty\, \sigma^+_{\min}(B) \downarrow 0$; without the assumption an unstable mode can keep ρ^* above 1 while $\eta_\infty \sigma^+_{\min}(B) \downarrow 0$.*

B.4 Proof of Theorem 4: Steps B and C

The standing notation of this proof, and Step A (the spectral geometry of the Jacobian), are in Section B.3.3; the two steps below use that notation without change. Throughout, $m = |S_1| = |S_2| \ge 2$, so that $W \ne \{0\}$ and $\Sigma^+(B) \ne \emptyset$; in the pure-saddle case $m = 1$ one has $\mathcal{M}_S = E$, every point of $\mathcal{M}_S \cap \mathcal{Z}$ is a fixed point (Proposition 1), and Lemma 6(i) already gives $\Phi(\mathcal{U}_0) \subseteq \mathrm{Fix}(\Phi)$, which is the finite-step statement recorded before Theorem 4.

Step B: uniform Taylor expansion and bootstrap

Let $\mathcal{U}$ denote the domain of analyticity of Φ_S around z^* provided by Lemma 6 (in face coordinates). Fix $r_0 > 0$ such that $\mathbb{B}(z^*, 3r_0) \subset \mathcal{U}$, the *ambient* closed ball of radius $3r_0$ around z^* is contained in the absorption neighborhood $\mathcal{U}_0$ of Lemma 6, and

$$r_0 \ \le\ \tfrac{1}{6} \min\big\{c_1 \|\xi\|_2,\ c_2 \|\psi\|_2\big\}. \tag{22}$$

Since $e_1 \perp e_2$, a point $z_f = a e_1 + b e_2 \in E$ satisfies $|a - c_1| \le \|z_f - z^*\|/\|\xi\|_2$ and $|b - c_2| \le \|z_f - z^*\|/\|\psi\|_2$; define accordingly

$$C_E := \frac{1}{\|\xi\|_2} + \frac{1}{\|\psi\|_2}, \tag{23}$$

so that $|a - c_1| + |b - c_2| \le C_E \|z_f - z^*\|$. In particular, by (22), every $z_f \in E \cap \mathbb{B}(z^*, 3r_0)$ has coordinates $(a, b) \in [c_1/2,\ 3c_1/2] \times [c_2/2,\ 3c_2/2] \subset R_\square$; such points lie in the open cone E_+ and are fixed points of Φ_S. Finally set

$$K := \tfrac{1}{2} \sup_{z \in \mathbb{B}(z^*, 3r_0)} \|D^2 \Phi_S(z)\| \ <\ \infty, \tag{24}$$

finite because Φ_S is analytic on a neighborhood of the compact ball.

Lemma 15 (Uniform second-order Taylor expansion)**.** *For every $z_f \in E_+ \cap \mathbb{B}(z^*, r_0)$ and every $w \in W$ with $\|w\| \le r_0$,*

$$\Phi_S(z_f + w) \;=\; z_f + J(z_f)\, w + R(z_f, w), \qquad \|R(z_f, w)\| \le K\|w\|^2.$$

Proof. The segment $\{z_f + sw\}_{s\in[0,1]}$ lies in $\mathbb{B}(z^*, 2r_0)$, on which Φ_S is analytic. Since $z_f \in E \cap \mathbb{B}(z^*, 3r_0) \subset E_+$ and Φ_S fixes E_+ pointwise (Proposition 1), $\Phi_S(z_f) = z_f$; and $D\Phi_S(z_f) = J(z_f)$ by Lemma 9. Taylor's theorem with integral remainder and (24) give the claim. □

Fix $\varepsilon \in (0, 1-\rho^*)$ and set $\theta := \rho^* + \varepsilon < 1$. Define

$$r_1 := \min\Big\{r_0,\ \frac{\varepsilon}{2\,\kappa_* C_J C_E}\Big\}, \qquad \delta_1 := \min\Big\{m_* r_0,\ \frac{\varepsilon\, m_*^2}{2M_* K},\ m_*\sqrt{\frac{r_1(1-\theta^2)}{2K}}\Big\}. \tag{25}$$

Lemma 16 (Bootstrap)**.** *Let $z^0 \in \mathcal{M}_S$ satisfy $\|P_E z^0 - z^*\| \le r_1/2$ and $\|P_W z^0\|_* \le \delta_1$. Define $z^{t+1} := \Phi_S(z^t)$, $z_f^t := P_E z^t$, and $w^t := P_W z^t$. Then for all $t \ge 0$:*

(M1) $\|w^t\|_* \le \theta^t \|w^0\|_*$*;*

(M2) $\|z_f^{t+1} - z_f^t\| \le K\|w^t\|^2 \le (K/m_*^2)\,\theta^{2t}\|w^0\|_*^2$*;*

(M3) $\|z_f^t - z^*\| \le r_1$*; in particular $z_f^t \in E_+ \cap \mathbb{B}(z^*, r_0)$ with coordinates $(a_t, b_t) \in R_\square$.*

Proof. Induction on t. For $t = 0$, (M1) is trivial and (M3) holds by hypothesis. Assume (M1) and (M3) hold for all $s \le t$; we prove (M2) at t and (M1), (M3) at $t+1$.

(i) The expansion applies. By (M1) and Lemma 14(2), $\|w^t\| \le \|w^t\|_*/m_* \le \delta_1/m_* \le r_0$, and by (M3), $z_f^t \in E_+ \cap \mathbb{B}(z^*, r_1) \subseteq \mathbb{B}(z^*, r_0)$. Lemma 15 gives

$$z^{t+1} = z_f^t + J(z_f^t)\, w^t + R^t, \qquad \|R^t\| \le K\|w^t\|^2.$$

(ii) W-component. By Lemma 13(3), applied at the *current* fixed point z_f^t (block diagonality is exact for every (a_t, b_t)), $J(z_f^t)w^t \in W$ and $P_W J(z_f^t)w^t = \mathcal{A}(a_t, b_t)\, w^t$; hence $w^{t+1} = \mathcal{A}(a_t, b_t)\, w^t + P_W R^t$. Lemma 14(3) bounds $\|\mathcal{A}(a_t, b_t)\|_*$ by $\rho^* + \kappa_* C_J(|a_t - c_1| + |b_t - c_2|)$; by (23) the bracket is at most $C_E\|z_f^t - z^*\|$, which (M3) bounds by $C_E r_1$; and r_1 was chosen in (25) so that $\kappa_* C_J C_E r_1 \le \varepsilon/2$. Hence

$$\|\mathcal{A}(a_t, b_t)\|_* \le \rho^* + \kappa_* C_J\, C_E\, r_1 \le \rho^* + \tfrac{\varepsilon}{2}.$$

Using $\|P_W\| = 1$, $\|P_W R^t\|_* \le M_*\|R^t\| \le M_* K\|w^t\|^2 \le (M_* K/m_*^2)\,\|w^t\|_*^2 \le (M_* K \delta_1/m_*^2)\,\|w^t\|_*$, so

$$\|w^{t+1}\|_* \le \Big(\rho^* + \tfrac{\varepsilon}{2} + \tfrac{M_* K}{m_*^2}\,\delta_1\Big)\|w^t\|_* \le \theta\,\|w^t\|_*,$$

where the last step uses $\delta_1 \le \varepsilon m_*^2/(2M_* K)$ from (25). This proves (M1) at $t+1$.

(iii) E-component. By Lemma 13(3), $P_E J(z_f^t)w^t = 0$, hence $z_f^{t+1} = z_f^t + P_E R^t$ and $\|z_f^{t+1} - z_f^t\| \le K\|w^t\|^2$; combining with (M1) and Lemma 14(2) gives (M2) at t.

(iv) (M3) at $t+1$. Summing (M2),

$$\|z_f^{t+1} - z^*\| \le \|z_f^0 - z^*\| + \sum_{s=0}^{t} K\|w^s\|^2 \le \frac{r_1}{2} + \frac{K\,\delta_1^2}{m_*^2\,(1-\theta^2)} \le \frac{r_1}{2} + \frac{r_1}{2} = r_1,$$

where the last inequality uses $\delta_1 \le m_*\sqrt{r_1(1-\theta^2)/(2K)}$ from (25). The coordinate statement follows from (22) and (23) as before. □

STEP C: PROOF OF THE THEOREM

Proof of Theorem 4. Parts (i) and (ii). By Lemma 9 and Lemma 10, $J(z_f) = I + L + L^2$ with the blocks of L given by B/b and $-B^\top/a$. By Lemma 12 and Lemma 13, the face splits orthogonally as $E \oplus W$ with $J(z_f)|_E = I_E$ and $J(z_f)|_{W_k} = J_k(a,b)$, whose eigenvalues are $1 - \eta^2\sigma_k^2 \pm i\eta\sigma_k$. Counting dimensions gives $2 + 2(s-1) = n_S$, so the multiset union in (9) exhausts the spectrum. No block contributes a unit eigenvalue (that would require $\lambda + \lambda^2 = 0$ with $\lambda = \pm i\eta\sigma_k$, impossible), so the eigenvalue 1 has algebraic multiplicity exactly two; it is semisimple because $J = I_E \oplus \big(\bigoplus_k J_k\big)$ with each J_k diagonalizable over $\mathbb{C}$, and its eigenspace is E. Since $\mathrm{Fix}(\Phi_S)$ coincides with the open cone $E_+ \subset E$ near z_f (Proposition 1 and uniqueness), E is the tangent space of the fixed-point cone. For (ii), write $q_k := \eta^2\sigma_k^2 > 0$; then $|\mu_k^\pm|^2 = 1 - q_k(1-q_k) < 1$ iff $q_k \in (0,1)$, so $\rho(a,b) < 1$ iff $q_{\max} = \eta^2\sigma_{\max}^2 < 1$ iff $\eta\,\sigma_{\max}(B) < 1$; and if $\eta\sigma_{\max}(B) > 1$ then $|\mu_1^\pm|^2 = 1 + q_{\max}(q_{\max} - 1) > 1$.

Part (iii). Assume Assumption 2, fix $\varepsilon \in (0, 1-\rho^*)$, and let θ, r_1, δ_1 be as in (25).

Step 1 (capture). Let

$$\mathcal{V} := \big\{ z \in \mathcal{M}_S : \ \|P_E z - z^*\| < r_1/2,\ \|P_W z\|_* < \delta_1 \big\},$$

a relatively open subset of $\mathcal{M}_S$ containing z^*. The full-space map Φ is continuous at z^* with $\Phi(z^*) = z^*$ and $\Phi(\mathcal{U}_0) \subseteq \mathcal{M}_S$ (Lemma 6); hence

$$\mathcal{U}' := \Phi^{-1}(\mathcal{V}) \cap \mathcal{U}_0 \cap \mathbb{B}(z^*, r_0)^\circ$$

is an open neighborhood of z^* (here $\mathbb{B}(z^*, r_0)^\circ$ is the open ball in the ambient space). If $z^{(t_0)} \in \mathcal{U}'$, then $z^{(t_0+1)} \in \mathcal{V} \subset \mathcal{M}_S$ and $z^{(t)} = \Phi_S^{t-t_0-1}\big(z^{(t_0+1)}\big)$ for $t \geq t_0 + 1$. Apply Lemma 16 with $z^0 := z^{(t_0+1)}$ (its hypotheses are the defining inequalities of $\mathcal{V}$).

Step 2 (geometric decay of the normal component). By (M1) and Lemma 14(2), $\|w^t\| \leq (\delta_1/m_*)\,\theta^t$.

Step 3 (convergence of the tangential component). By (M2), for all $p \geq 1$,

$$\|z_f^{t+p} - z_f^t\| \leq \sum_{s \geq t} K\|w^s\|^2 \leq \frac{K\,\delta_1^2}{m_*^2(1-\theta^2)}\,\theta^{2t} \ \xrightarrow[t\to\infty]{}\ 0,$$

so (z_f^t) is Cauchy and converges to some $z_\infty \in E \cap \mathbb{B}(z^*, r_1)$. By (22) and (23), its coordinates satisfy $a_\infty \geq c_1/2 > 0$ and $b_\infty \geq c_2/2 > 0$, so $z_\infty \in E_+$ and z_∞ is a fixed point (Proposition 1).

Step 4 (rate synthesis). For the bootstrap time t,

$$\|z^t - z_\infty\| \leq \|w^t\| + \|z_f^t - z_\infty\| \leq \frac{\delta_1}{m_*}\,\theta^t + \frac{K\delta_1^2}{m_*^2(1-\theta^2)}\,\theta^{2t} \leq C'\,\theta^t,$$

$$C' := \frac{\delta_1}{m_*}\Big(1 + \frac{K\delta_1}{m_*(1-\theta^2)}\Big).$$

Returning to the original time index ($t \mapsto t - t_0 - 1$), for $t \geq t_0 + 1$ we get $\|z^{(t)} - z_\infty\| \leq (C'/\theta)\,\theta^{\,t-t_0}$; for $t = t_0$, $\|z^{(t_0)} - z_\infty\| \leq \|z^{(t_0)} - z^*\| + \|z^* - z_\infty\| \leq r_0 + r_1 \leq 2r_0$. Hence (11) holds with $C_\varepsilon := \max\{C'/\theta,\ 2r_0\}$.

Part (iv). For a state z near z^*, let $(x(z), y(z))$ denote the main-iterate strategies produced by the round at z (zero-embedded when $z \in \mathcal{M}_S$), and recall $\mathrm{gap}(x,y) = \max_i (Ay)_i - \min_j (x^\top A)_j \geq 0$. On the ambient ball $\mathbb{B}(z^*, 2r_0)$ the map $z \mapsto (x(z), y(z))$ is Lipschitz: the pre-iterate normalizations have denominators bounded below, $\|\tilde r_1\|_1 \geq \|\tilde r_1\|_2 \geq c_1\|\xi\|_2 - 2r_0 \geq \frac{2}{3}\,c_1\|\xi\|_2$ by (22) (and symmetrically for $\tilde r_2$); the multiplier γ_1 is $\sqrt{2}$-Lipschitz in (v, N_1), each partial derivative having norm at most one: wherever the active set $\{i : v_i > \gamma_1\}$ is locally constant, implicit differentiation of $\varphi(\gamma_1; v)^2 = N_1^2$ gives $\nabla_v \gamma_1 = p_1/\|p_1\|_1$ and $\partial\gamma_1/\partial N_1 = -N_1/\|p_1\|_1$, both of norm at most one (as $\|p_1\|_1 \geq \|p_1\|_2 = N_1$), and γ_1 is continuous across active-set changes (Step 1 of the proof of Lemma 6); clipping is 1-Lipschitz; and the main-iterate normalizations have denominators $\|p_i\|_1 \geq \|p_i\|_2 = N_i$ bounded below. Composing with the Lipschitz functionals $y \mapsto \max_i (Ay)_i$ and $x \mapsto \min_j (x^\top A)_j$ shows that $z \mapsto \mathrm{gap}(x(z), y(z))$ is L_{gap}-Lipschitz on $\mathbb{B}(z^*, 2r_0)$, with L_{gap} depending only on $\|A\|$, c_1, c_2, and the dimensions. Every fixed point $z_f \in E_+$ produces the strategies (x^*, y^*) (Proposition 1), so $\mathrm{gap}(x(z_\infty), y(z_\infty)) = 0$, and therefore, for all $t \geq t_0$,

$$\begin{aligned}\mathrm{gap}\big(x^{(t)}, y^{(t)}\big) &= \mathrm{gap}\big(x(z^{(t)}), y(z^{(t)})\big) - \mathrm{gap}\big(x(z_\infty), y(z_\infty)\big) \\ &\leq L_{\mathrm{gap}}\,\|z^{(t)} - z_\infty\| \leq L_{\mathrm{gap}}\,C_\varepsilon\,(\rho^* + \varepsilon)^{\,t-t_0}.\end{aligned}$$

□

Dependence of constants. For reference: r_0 depends on the absorption neighborhood of Lemma 6 (hence on δ, $\|A\|$, c_1, c_2, the dimensions, and the smallest on-support probabilities $\min_{i\in S_1} x_i^*$, $\min_{j\in S_2} y_j^*$, which enter Step 3 of its proof); K on the second derivatives of Φ_S over $\mathbb{B}(z^*, 3r_0)$; C_J and C_E are explicit in $(\sigma_{\max}, c_1, c_2, \|\xi\|_2, \|\psi\|_2)$; m_*, M_*, κ_* are the closed forms of Lemma 14, functions of (c_1, c_2) alone; r_1, δ_1, C_ε depend on all of the above and on ε.

B.5 Open constants

Remark 7 (Open: quantitative absorption radius). Lemma 6 is proved by a continuity argument, so the radius of the absorption neighborhood $\mathcal{U}_0$ (and through it r_0 and, via (25), the constants r_1, δ_1, C_ε) is *existential*. We have not derived an explicit lower bound on it in terms of $(\delta, \|A\|, c_1, c_2, \min_{i\in S_1} x_i^*, \min_{j\in S_2} y_j^*)$ and the dimensions; a fully explicit-constants version of the local theory would require quantifying this radius.

Remark 8 (Open: second-order constants). The curvature bound K of (24) is obtained by compactness and is not given in closed form; the rate ρ^* of (10), the constants C_J, C_E, and the adapted-norm constants m_*, M_*, κ_* of Lemma 14 are fully explicit. Thus in Theorem 4 the *rate* is closed-form while the multiplicative *constant* C_ε is not, through K and the absorption radius alone; closing this asymmetry is open.

C Proofs for the Global Theory

Throughout this appendix, $A \in \mathbb{R}^{n_1\times n_2}$ is the payoff matrix (row player maximizes $x^\top Ay$), $d := n_1 + n_2$, and $\|A\|$ denotes the $\ell_2 \to \ell_2$ operator norm. The state is $z = (\tilde r_1, \tilde r_2) \in \mathbb{R}^{n_1}_{\geq 0} \times \mathbb{R}^{n_2}_{\geq 0}$ with $N_i(z) := \|\tilde r_i\|_2$, and the state space is $\mathcal{Z}$ as in (6). All statements refer to Algorithm 1, the strict version ($r = \tilde r + m - \gamma\mathbf{1}$ without clipping in the prediction step; update $[r+g]_+$).

Remark 9 (Robustness to the clipped-prediction variant). A legitimate variant clips already in the prediction step ($r_{\text{clip}} = [\tilde r + m - \gamma\mathbf{1}]_+$, update $[r_{\text{clip}} + g]_+$). Every *global* argument in this appendix holds verbatim for that variant: the proofs only use (a) the norm-preservation and nonnegativity of $[r]_+$ (equal to r_{clip}), (b) the 1-Lipschitz property of $[\cdot]_+$, and (c) the fact that the update step outputs $[r]_+ = r_{\text{clip}}$ when $g = 0$. The version difference affects only the *local* smoothness analysis; the analyticity arguments of Appendix C.6 are stated for the strict version.

C.1 The round map and standing facts

The round map. On $\mathcal{Z}$ define $\Phi : \mathcal{Z} \to \mathcal{Z}$ (extragradient setup, the two players symmetric):

(1) pre-iterates $\tilde X(z) = \tilde r_1/\|\tilde r_1\|_1$, $\tilde Y(z) = \tilde r_2/\|\tilde r_2\|_1$;

(2) predictions $m_1(z) = A\tilde Y(z)$ (note: through the *opponent's* component $\tilde r_2$), $m_2(z) = -A^\top \tilde X(z)$;

(3) γ-shift: $v_1 = \tilde r_1 + m_1$, $\gamma_1(z) = \gamma(v_1, N_1(z))$ (the unique solution provided by Lemma 19 below), $r_1(z) = v_1 - \gamma_1\mathbf{1}$, and main iterate $X(z) = [r_1]_+ / \|[r_1]_+\|_1$; r_2, Y symmetric;

(4) residuals $\text{res}_1(z) = A\big(Y(z) - \tilde Y(z)\big)$ and $g_1(z) = \text{res}_1 - \langle \text{res}_1, X(z)\rangle\mathbf{1}$; g_2 symmetric;

(5) $\Phi(z) := \big([r_1 + g_1]_+ , [r_2 + g_2]_+\big)$.

Well-definedness. $\|\tilde r_i\|_1 \geq \|\tilde r_i\|_2 > 0$ on $\mathcal{Z}$; γ_i exists and is unique (Lemma 19 with target $c = N_i > 0$); $\|[r_i]_+\|_1 \geq \|[r_i]_+\|_2 = N_i > 0$ by norm preservation; and $\Phi(\mathcal{Z}) \subseteq \mathcal{Z}$ because the update step does not decrease the norm: $\|[r_i + g_i]_+\|_2 \geq \|[r_i]_+\|_2 = N_i > 0$ (Lemma 5).

Trajectory–map alignment. Let $t_+ := \min\{t : N_1(t) > 0 \text{ and } N_2(t) > 0\} = \max(t_1^*, t_2^*)$. The standing assumption $t_+ < \infty$ (Section 4.1) is exactly $N_{i,\infty} > 0$ for both players (Theorem 1 supplies finiteness), and norm monotonicity gives $N_i(t) \geq N_i(t_+) > 0$ for all $t \geq t_+$. For $t < t_i^*$ player i follows its zero-state branch (Algorithm 1: its pre-iterate is carried over from the previous round, starting from the default uniform, its prediction is $\mathbf{0}$ and its main iterate equals its pre-iterate), while a player with $t_i^* \leq t < t_+$ already runs the ordinary prediction step. For $t \geq t_+$ we have $z^{(t)} \in \mathcal{Z}$ and, round by round, $z^{(t+1)} = \Phi(z^{(t)})$, with $x^{(t)} = X(z^{(t)})$, $\tilde x^{(t)} = \tilde X(z^{(t)})$, $y^{(t)} = Y(z^{(t)})$,

$\tilde{y}^{(t)} = \tilde{Y}(z^{(t)})$, $g_i^{(t)} = g_i(z^{(t)})$. That is, from t_+ on the dynamics is a *memoryless autonomous* discrete system on $\mathcal{Z}$; we analyze only the tail $\{z^{(t)}\}_{t \geq t_+}$.

Standing facts. We collect the facts quoted in the proofs, with their sources.

(E1) $N_i(t)$ is nondecreasing along the whole trajectory (Lemma 5).

(E2) *(Saturation)* $N_i(t) \uparrow N_{i,\infty} < \infty$ (Theorem 1) and $N_{i,\infty} > 0$ (the standing assumption $t_+ < \infty$), for both players.

(E3) By Lemma 1 (first inequality of (3), established self-containedly in Appendix B.2), saturation of both players implies

$$\sum_{t \geq t_+} \|x^{(t)} - \tilde{x}^{(t)}\|_2^2 \;\leq\; \frac{8n_1\,(N_{1,\infty} + N_{2,\infty})}{N_1(t_+)} \;<\; \infty,$$

(the tail sum from t_+, which is all the analysis below uses; the full-sum version carries the norm at the first positive-norm round in the denominator), and symmetrically for player 2; hence $\|x^{(t)} - \tilde{x}^{(t)}\| \to 0$ and $\|y^{(t)} - \tilde{y}^{(t)}\| \to 0$ along the *full* sequence (summability implies vanishing terms).

(E4) *(Scalar core inequality; second inequality of* (4)*)*

$$\beta_t := \langle m^{(t)}, x^{(t)} \rangle - \gamma^{(t)} \;\geq\; \frac{1}{8n_1}\, N_1(t)\, \|x^{(t)} - \tilde{x}^{(t)}\|^2 \;\geq\; 0, \qquad \sum_t \beta_t < \infty, \tag{26}$$

so $\beta_t \to 0$ along the full sequence.

(E5) By Proposition 1: for $z \in \mathcal{Z}$, $\Phi(z) = z$ iff $(X(z), Y(z)) = (\tilde{X}(z), \tilde{Y}(z))$ is a Nash equilibrium with corresponding supports; $\mathrm{Fix}(\Phi) = \{(c_1 x^*, c_2 y^*) : c_1, c_2 > 0,\ (x^*, y^*) \in \mathrm{NE}\}$ (the positive cone of NE rays), and at fixed points $\gamma_1 = v^*$ (the game value).

(E6) *(Residual bound)* See Lemma 17 below.

(E7) By Lemmas 9, 10 and 11, all stated under Assumption 3 (nondegenerate, strictly complementary equilibrium; equivalent to Assumption 1 by Lemma 11): at such an NE ray point $z^* = (c_1 x^*, c_2 y^*)$,

$$\mathrm{spec}\big(D\Phi_S(z^*)\big) = \{1, 1\} \cup \big\{\, 1 - \eta_\infty^2 \sigma^2 \pm i\, \eta_\infty \sigma \,\big\}, \qquad \eta_\infty = 1/\sqrt{c_1 c_2},$$

with σ ranging over the nonzero singular values of $B = \tilde{A} - v^* \mathbf{1}\mathbf{1}^\top$. This is the spectrum statement of Theorem 4.

Lemma 17 (Residual bound)**.** *For all $t \geq t_+$, $\|g_1^{(t)}\|_2 \leq (1 + \sqrt{n_1})\, \|A\|\, \|y^{(t)} - \tilde{y}^{(t)}\|_2$, and symmetrically for player* 2*.*

Proof. $g_1 = \mathrm{res}_1 - \langle \mathrm{res}_1, x \rangle \mathbf{1}$ with $x \in \Delta_{n_1}$. Then $|\langle \mathrm{res}_1, x \rangle| \leq \|\mathrm{res}_1\|_\infty \|x\|_1 \leq \|\mathrm{res}_1\|_2$, $\|\mathbf{1}\|_2 = \sqrt{n_1}$, and $\|\mathrm{res}_1\|_2 = \|A(y - \tilde{y})\|_2 \leq \|A\|\, \|y - \tilde{y}\|_2$. □

C.2 Proofs

This subsection proves Theorems 2 and 3. Both rest on continuity properties of the round map Φ: joint continuity of the γ-shift across support changes, continuity of every intermediate map, the residual bound, and the two elementary facts that a bounded sequence with vanishing steps has a connected set of limit points. These are collected, with proofs, in Appendix C, and are used here as stated. The proofs of the norm-preserving pairing Lemma 2, of Corollary 1, and of Proposition 2 are also in Appendix C.

C.2.1 Proof of Theorem 2

We prove the two parts as two propositions, matching the structure of the theorem.

Proposition 3 (ω-limit set: existence, structure, invariance)**.** *Under the hypotheses of Theorem 2, the tail trajectory $\{z^{(t)}\}_{t \geq t_+}$ is contained in the compact set*

$$K := \{z :\ \tilde{r}_i \geq 0,\ \|\tilde{r}_i\|_2 \leq N_{i,\infty},\ i = 1, 2\},$$

its ω-limit set Ω is nonempty, compact, connected (the connectedness is of independent interest and not used in the sequel), satisfies $\Phi(\Omega) = \Omega$, and

$$\Omega \subseteq \mathcal{L} := \{z : \tilde{r}_i \geq 0,\ \|\tilde{r}_i\|_2 = N_{i,\infty},\ i = 1, 2\} \subset \mathcal{Z}.$$

Moreover $\operatorname{dist}(z^{(t)}, \Omega) \to 0$.

Proof. Boundedness. $\tilde{r}_i^{(t)} \geq 0$ always (output of $[\cdot]_+$), and $\|\tilde{r}_i^{(t)}\|_2 = N_i(t) \leq N_{i,\infty}$ by (E1)–(E2); hence $z^{(t)} \in K$, and K is compact.

Nonempty, compact, connected. Nonemptiness, compactness, and connectedness follow from Lemma 21 (which uses (E3) and Lemma 17) and Lemma 22.

$\Omega \subseteq \mathcal{L}$. If $\bar{z} = \lim_k z^{(t_k)}$ then $\|\tilde{r}_i(\bar{z})\|_2 = \lim_k N_i(t_k) = N_{i,\infty}$ (norms are continuous, and a monotone sequence has the same limit along every subsequence). Since $N_{i,\infty} > 0$ we get $\mathcal{L} \subset \mathcal{Z}$, so Φ is continuous on an open neighborhood of Ω (Lemma 20).

Positive invariance $\Phi(\Omega) \subseteq \Omega$. If $\bar{z} \in \Omega$ with $z^{(t_k)} \to \bar{z}$, then $z^{(t_k+1)} = \Phi(z^{(t_k)}) \to \Phi(\bar{z})$ by continuity; the left-hand side has its limit in Ω, so $\Phi(\bar{z}) \in \Omega$.

Reverse inclusion $\Omega \subseteq \Phi(\Omega)$. Let $\bar{z} = \lim_k z^{(t_k)}$ with $t_k \geq t_+ + 1$. The sequence $\{z^{(t_k-1)}\} \subset K$ has an accumulation point $w \in \Omega$; along a sub-subsequence, $\Phi(w) = \lim \Phi(z^{(t_k-1)}) = \lim z^{(t_k)} = \bar{z}$. Hence $\Phi(\Omega) = \Omega$.

$\operatorname{dist}(z^{(t)}, \Omega) \to 0$. If not, then $\operatorname{dist}(z^{(t)}, \Omega) \geq \varepsilon$ along a subsequence; by compactness pass to a convergent sub-subsequence, whose limit lies in Ω, a contradiction. □

All constants in this proof are explicit and involve no hidden game-dependent quantities: the displacement constants are $(1 + \sqrt{n_i})\|A\|$ and $2\sqrt{n_i}N_{i,\infty}$, and the summability constant is $8n_i(N_{1,\infty} + N_{2,\infty})/N_i(t_+)$, all directly computable from the game and the run.

Proposition 4 (Limit points are fixed points)**.** *Under the hypotheses of Theorem 2, $\Omega \subseteq \operatorname{Fix}(\Phi)$. Combined with (E5): every point of Ω has the form $(c_1 x^*, c_2 y^*)$ with $(x^*, y^*) \in \mathrm{NE}$, the scales determined by the level-set condition $\|c_i x^*\|_2 = N_{i,\infty}$ up to the strategy direction, and $\gamma_1 = v^*$ at each such point.*

Proof. Fix $\bar{z} \in \Omega$ with $z^{(t_k)} \to \bar{z}$. By Proposition 3, $\bar{z} \in \mathcal{L} \subset \mathcal{Z}$, so all maps of Lemma 20 are continuous at $\bar{z}$.

Step 1 (pre-iterate = main iterate). By the trajectory–map alignment, $x^{(t_k)} = X(z^{(t_k)})$ and $\tilde{x}^{(t_k)} = \tilde{X}(z^{(t_k)})$; continuity gives $x^{(t_k)} \to X(\bar{z})$ and $\tilde{x}^{(t_k)} \to \tilde{X}(\bar{z})$. Both are limits of *images of the same sequence under continuous maps*, so no extra uniformity assumption is needed; the dependence of $m_1(\bar{z})$ on the opponent component $\tilde{r}_2(\bar{z})$ through $\tilde{Y}(\bar{z})$ is already absorbed in the definition of X. By (E3) the full sequence $\|x^{(t)} - \tilde{x}^{(t)}\| \to 0$, so

$$\|X(\bar{z}) - \tilde{X}(\bar{z})\| = \lim_k \|x^{(t_k)} - \tilde{x}^{(t_k)}\| = 0,$$

that is, $X(\bar{z}) = \tilde{X}(\bar{z})$; symmetrically $Y(\bar{z}) = \tilde{Y}(\bar{z})$.

Step 2 (prediction step stationary). At $\bar{z}$, take $a = [r_1(\bar{z})]_+$ and $b = \tilde{r}_1(\bar{z})$: both are nonnegative, nonzero, and $\|a\|_2 = \|b\|_2 = N_1(\bar{z})$ (the norm-preservation equation of the γ-shift holds at $\bar{z}$ by definition). Step 1 says their ℓ_1-normalizations coincide, so the identity of Lemma 2 gives $a = b$:

$$[r_1(\bar{z})]_+ = \tilde{r}_1(\bar{z}), \qquad [r_2(\bar{z})]_+ = \tilde{r}_2(\bar{z}).$$

(Coordinates with $\tilde{r}_i = 0$ are handled automatically: Lemma 2 is a vector identity.)

Step 3 (update step stationary). $\operatorname{res}_1(\bar{z}) = A\big(Y(\bar{z}) - \tilde{Y}(\bar{z})\big) = \mathbf{0}$ by Step 1, hence $g_1(\bar{z}) = \mathbf{0}$; symmetrically $g_2(\bar{z}) = \mathbf{0}$. Therefore

$$\Phi(\bar{z})_1 = [r_1(\bar{z}) + g_1(\bar{z})]_+ = [r_1(\bar{z})]_+ = \tilde{r}_1(\bar{z}),$$

and likewise for player 2, so $\Phi(\bar{z}) = \bar{z}$. The NE characterization and $\gamma_1 = v^*$ follow from (E5). □

Proof of Theorem 2. Part (i) is Propositions 3 and 4.

Part (ii): the function gap is continuous in (x, y) and vanishes on the NE set. The sequence $\mathrm{gap}(x^{(t)}, y^{(t)})$ is bounded; take any of its limit points, realized along times t_k. By compactness pass to a sub-subsequence with $z^{(t_k)} \to \bar{z} \in \Omega$; by continuity (Lemma 20) and Proposition 4 with (E5), $(x^{(t_k)}, y^{(t_k)}) \to (X(\bar{z}), Y(\bar{z})) \in \mathrm{NE}$, so the limit point equals 0. A bounded sequence all of whose limit points are 0 converges to 0. The distance statement is identical (the NE set is compact and $\mathrm{dist}(\cdot, \mathrm{NE})$ continuous), and the pre-iterate statements follow from (E3). □

C.2.2 Proof of Theorem 3

Proof. By Theorem 2, $\Omega \subseteq \mathrm{Fix}(\Phi) \cap \mathcal{L}$. NE uniqueness gives $\mathrm{Fix}(\Phi) = \{(c_1 x^*, c_2 y^*) : c_i > 0\}$ by (E5), and the level-set condition $\|c_1 x^*\|_2 = N_{1,\infty}$ determines $c_1 = N_{1,\infty}/\|x^*\|_2$ uniquely (likewise c_2), so $\mathrm{Fix}(\Phi) \cap \mathcal{L}$ is a *single point* $\bar{z}$; hence $\Omega = \{\bar{z}\}$ (it is nonempty by Proposition 3). A sequence in a compact set with a unique limit point converges, so $z^{(t)} \to \bar{z}$. Convergence of the strategies follows from Lemma 20: $x^{(t)} = X(z^{(t)}) \to X(\bar{z}) = x^*$, and similarly for $\tilde{x}^{(t)}, y^{(t)}, \tilde{y}^{(t)}$. □

C.2.3 Never-positive players: the case $t_+ = \infty$

Theorems 2 and 3 assume $t_+ < \infty$. The complementary case is elementary and is settled by the next lemma, so the gap conclusion of Theorem 2(ii) holds on every matrix game.

Lemma 18 (Never-positive players)**.** *Suppose $t_1^* = \infty$, i.e. $\tilde{r}_1^{(t)} = \mathbf{0}$ for all t (the case $t_2^* = \infty$ follows by exchanging the two players). Write $\bar{x} := \mathbf{1}/n_1$ and $\ell := -A^\top \bar{x} \in \mathbb{R}^{n_2}$, and let $K_\ell := \{k : \ell_k = \max_{k'} \ell_{k'}\}$.*

- *(i) $x^{(t)} = \tilde{x}^{(t)} = \bar{x}$ for all t, the vector $Ay^{(t)}$ is constant across coordinates for all t, and player 1's exploitability $\max_i (Ay^{(t)})_i - \bar{x}^\top A y^{(t)}$ is 0 at every round.*
- *(ii) If $t_2^* < \infty$, then for all $t \geq t_2^*$: $N_2(t) = N_2(t_2^*)$, $\gamma_2^{(t)} \leq \max_k \ell_k$, and $\sum_{t \geq t_2^*} \big(\max_k \ell_k - \gamma_2^{(t)}\big) \leq N_2(t_2^*)$.*
- *(iii) If $t_2^* < \infty$, there is a finite T such that $y^{(t)}$ is supported on K_ℓ for all $t \geq T$; hence $\mathrm{gap}(x^{(t)}, y^{(t)}) = 0$ for all $t \geq T$.*
- *(iv) If also $t_2^* = \infty$, the uniform profile $(\mathbf{1}/n_1, \mathbf{1}/n_2)$ is a Nash equilibrium and $\mathrm{gap}(x^{(t)}, y^{(t)}) = 0$ for all t.*

In particular $\mathrm{gap}(x^{(t)}, y^{(t)}) \to 0$ without any hypothesis on t_+, and if the equilibrium is unique the strategies converge to it.

Proof. (i) With $\tilde{r}_1^{(t)} = \mathbf{0}$ the zero branch of Algorithm 1 keeps $\tilde{x}_1^{(t)} = \tilde{x}_1^{(1)} = \bar{x}$, sets $m_1^{(t)} = r_1^{(t)} = \mathbf{0}$, and plays $x^{(t)} = \bar{x}$. Then $g_1^{(t)} = u^{(t)} - \langle u^{(t)}, \bar{x} \rangle \mathbf{1}$ with $u^{(t)} = Ay^{(t)}$, and $\tilde{r}_1^{(t+1)} = \left[g_1^{(t)}\right]_+ = \mathbf{0}$ forces $u_i^{(t)} \leq \langle u^{(t)}, \bar{x} \rangle$ for every i. The mean of the $u_i^{(t)}$ equals $\langle u^{(t)}, \bar{x} \rangle$, so all $u_i^{(t)}$ are equal, and $\max_i (Ay^{(t)})_i - \bar{x}^\top A y^{(t)} = 0$.

(ii) Player 2 faces the fixed opponent $\bar{x}$, so its prediction and its observed payoff coincide: $m_2^{(t)} = u_2^{(t)} = \ell$ for all $t \geq t_2^*$ (before t_2^* the zero branch plays $y^{(t)} = \bar{y}$ and sets $m_2^{(t)} = 0$). Its update reads $\tilde{r}_2^{(t+1)} = \left[\tilde{r}_2^{(t)} + \ell - \gamma_2^{(t)} \mathbf{1}\right]_+$, where $\gamma_2^{(t)}$ is the norm-preserving shift, so $N_2(t+1) = N_2(t)$ for $t \geq t_2^*$. If $\gamma_2^{(t)} > \max_k \ell_k$, every coordinate with $\tilde{r}_{2,k}^{(t)} > 0$ strictly decreases and the others stay at 0; since some coordinate is positive, the norm would strictly decrease, a contradiction. Hence $\gamma_2^{(t)} \leq \max_k \ell_k$. For $k^* \in K_\ell$ this gives $\tilde{r}_{2,k^*}^{(t+1)} \geq \tilde{r}_{2,k^*}^{(t)} + \max_k \ell_k - \gamma_2^{(t)} \geq \tilde{r}_{2,k^*}^{(t)}$, and $\tilde{r}_{2,k^*}^{(t)} \leq N_2(t_2^*)$ for all $t \geq t_2^*$, so the nonnegative increments are summable with the stated bound.

(iii) By (ii), $\gamma_2^{(t)} \to \max_k \ell_k$. Fix $j \notin K_\ell$ and $\delta_j := \max_k \ell_k - \ell_j > 0$; for all large t, $\ell_j - \gamma_2^{(t)} \le -\delta_j/2$, so $\tilde{r}_{2,j}^{(t+1)} \le \left[\tilde{r}_{2,j}^{(t)} - \delta_j/2\right]_+$, which reaches 0 in finitely many rounds and then stays at 0. Since $y^{(t)} \propto \left[\tilde{r}_2^{(t)} + \ell - \gamma_2^{(t)}\mathbf{1}\right]_+ = \tilde{r}_2^{(t+1)}$ for $t \ge t_2^*$ (line 8 of Algorithm 1; before t_2^*, $y^{(t)} = \bar{y}$ by the zero branch), $y^{(t)}$ is supported on K_ℓ from some finite T on. Then $\bar{x}^\top A y^{(t)} = -\max_k \ell_k = \min_j (\bar{x}^\top A)_j$, and with (i) the Nash gap is 0.

(iv) If both states stay at zero, both players play uniformly and, by (i) applied to each, each player's payoff vector is constant across coordinates, which is the Nash condition for the uniform profile.

Convergence of the gap is (iii) or (iv). Under a unique equilibrium (x^*, y^*), every limit point of $(\bar{x}, y^{(t)})$ has zero gap and is therefore (x^*, y^*), so $y^{(t)} \to y^*$ and $x^* = \bar{x}$. □

C.3 Preliminary lemmas

Lemma 19 (Joint continuity of the γ-shift)**.** *Let $D := \{(v, c) : v \in \mathbb{R}^n,\ c > 0\}$. For each $(v, c) \in D$ the equation $\|[v - \gamma\mathbf{1}]_+\|_2 = c$ has a unique solution $\gamma(v, c)$, which satisfies $\gamma(v, c) \in [\max_i v_i - c,\ \max_i v_i)$. Moreover γ is continuous on D, and for fixed c the map $\gamma(\cdot, c)$ is 1-Lipschitz in the ℓ_2 sense.*

Proof. Write $f(\gamma; v) := \|[v - \gamma\mathbf{1}]_+\|_2$; it is jointly continuous in (γ, v) (composition of $[\cdot]_+$ and the norm). For $\gamma \ge \max_i v_i$ we have $f = 0 < c$; at $\gamma_0 := \max_i v_i - c$, $f(\gamma_0; v) \ge v_{i^*} - \gamma_0 = c$ where i^* is the argmax. On $\{f > 0\} = (-\infty, \max_i v_i)$, f is strictly decreasing: indeed $f(\gamma)^2 = \sum_{i: v_i > \gamma} (v_i - \gamma)^2$, and as long as the active set is nonempty (which it is whenever $f > 0$), each active term strictly decreases in γ. Hence the solution exists, is unique, and lies in the stated interval.

Continuity. Let $(v_k, c_k) \to (v, c)$ and $\gamma_k := \gamma(v_k, c_k) \in [\max_i v_{k,i} - c_k,\ \max_i v_{k,i}]$, a bounded sequence. For any accumulation point $\bar{\gamma}$, joint continuity along a subsequence gives $f(\bar{\gamma}; v) = c$, so by uniqueness $\bar{\gamma} = \gamma(v, c)$; a bounded sequence all of whose accumulation points coincide converges, so $\gamma_k \to \gamma(v, c)$.

1-*Lipschitz.* By coordinatewise monotonicity of the shift, $f(\gamma; v + \delta) \le f(\gamma - \|\delta\|_\infty; v)$ and $f(\gamma; v + \delta) \ge f(\gamma + \|\delta\|_\infty; v)$; combined with the monotonicity of f in γ and uniqueness of the solution, $|\gamma(v + \delta, c) - \gamma(v, c)| \le \|\delta\|_\infty \le \|\delta\|_2$. (This is consistent with the gradient identity $\nabla_v \gamma = x \in \Delta_n$ from the anchor analysis.) □

Continuity holds *across support changes*: the argument uses only monotonicity and uniqueness, with no piecewise case analysis.

Lemma 20 (Continuity of Φ and all intermediate maps)**.** *The maps $\tilde{X}, \tilde{Y}, m_i, \gamma_i, r_i, X, Y, g_i, \Phi$, as well as $\beta(z) := \langle m_1(z), X(z)\rangle - \gamma_1(z)$, are continuous on $\mathcal{Z}$.*

Proof. Compose step by step. $\tilde{X} = \tilde{r}_1/\|\tilde{r}_1\|_1$ is continuous (the denominator is $\ge N_1(z) > 0$); $m_1 = A\tilde{Y}$ is continuous (through the opponent component $\tilde{r}_2$); $v_1 = \tilde{r}_1 + m_1$ and $c_1 = N_1(z)$ are continuous with $c_1 > 0$, so $\gamma_1 = \gamma(v_1, c_1)$ is continuous by Lemma 19, hence so is r_1; $X = [r_1]_+ / \|[r_1]_+\|_1$ is continuous (denominator $\ge \|[r_1]_+\|_2 = c_1 > 0$ by norm preservation); Y symmetrically; therefore res_1, g_1, and finally Φ are continuous. □

Proof of Lemma 2 (norm-preserving pairing). Since $a = \|a\|_1 x$ and $N = \|a\|_2 = \|a\|_1 \|x\|_2$, we get $\|a\|_1 = N/\|x\|_2$ and hence $a = N\, x/\|x\|_2$; likewise $b = N\, \tilde{x}/\|\tilde{x}\|_2$, which is the identity in (7). For the Lipschitz part, for any $u, w \ne 0$ with (w.l.o.g.) $\|u\|_2 \ge \|w\|_2$,

$$\left\| \frac{u}{\|u\|_2} - \frac{w}{\|w\|_2} \right\|_2 \le \frac{\|u - w\|_2}{\|u\|_2} + \|w\|_2 \left(\frac{1}{\|w\|_2} - \frac{1}{\|u\|_2} \right) \le \frac{2\,\|u - w\|_2}{\max(\|u\|_2, \|w\|_2)},$$

and on the simplex $\|x\|_2 \ge 1/\sqrt{n}$, giving the constant $2\sqrt{n}$. The last claim follows from the identity: if $x = \tilde{x}$ then $a = b$ directly; the converse is trivial. Coordinates with $x_i = \tilde{x}_i = 0$ force $a_i = b_i = 0$ and are absorbed by the identity, with no separate case analysis. □

Applying Lemma 2 with $a = [r_1(z)]_+$ and $b = \tilde{r}_1$ (both nonnegative, nonzero, and of equal ℓ_2 norm $N_1(z)$ by the γ-shift's defining equation) yields the estimate used repeatedly below,

$$\big\|[r_1(z)]_+ - \tilde{r}_1\big\|_2 \;\le\; 2\sqrt{n_1}\, N_1(z)\, \|X(z) - \tilde{X}(z)\|_2, \tag{27}$$

where $X, \tilde{X}$ are the main and pre-iterate strategy maps.

Lemma 21 (Vanishing one-step displacement)**.** *Under saturation (E2), $\|z^{(t+1)} - z^{(t)}\| \to 0$, with the explicit bound (for $t \ge t_+$)*

$$\|\tilde{r}_1^{(t+1)} - \tilde{r}_1^{(t)}\|_2 \;\le\; \underbrace{(1+\sqrt{n_1})\,\|A\|\,\|y^{(t)} - \tilde{y}^{(t)}\|_2}_{\text{update step, Lemma 17}} \;+\; \underbrace{2\sqrt{n_1}\, N_{1,\infty}\,\|x^{(t)} - \tilde{x}^{(t)}\|_2}_{\text{prediction step, (27)}},$$

and symmetrically for player 2.

Proof. Decompose $\tilde{r}_1^{(t+1)} - \tilde{r}_1^{(t)} = \big([r_1 + g_1]_+ - [r_1]_+\big) + \big([r_1]_+ - \tilde{r}_1^{(t)}\big)$. The first term is bounded by $\|g_1^{(t)}\|_2$ (the map $[\cdot]_+$ is coordinatewise 1-Lipschitz), hence by Lemma 17; the second by (27) together with $N_1(t) \le N_{1,\infty}$. Both factors tend to 0 by (E3). □

Lemma 22 (Connected limit sets for vanishing-step sequences)**.** *Let $\{z^{(t)}\}$ be contained in a compact set $K \subset \mathbb{R}^d$ with $\|z^{(t+1)} - z^{(t)}\| \to 0$. Then its set of limit points Ω is nonempty, compact, and* connected*.*

Proof. Nonemptiness and compactness are standard (Bolzano–Weierstrass; $\Omega = \bigcap_T \overline{\{z^{(t)} : t \ge T\}}$ is an intersection of a nested family of compact sets). For connectedness, suppose for contradiction that $\Omega = \Omega_A \sqcup \Omega_B$ with both parts nonempty, compact, and disjoint; let $\delta := \operatorname{dist}(\Omega_A, \Omega_B) > 0$ and let U_A, U_B be the open $\delta/3$-neighborhoods of Ω_A, Ω_B. Since both parts consist of limit points, the sequence visits U_A and U_B infinitely often, so there are infinitely many "exit events": choose $t_k \uparrow \infty$ with $z^{(t_k)} \in U_A$ and $\tau_k := \min\{\tau > t_k : z^{(\tau)} \notin U_A\} < \infty$ (otherwise the tail would be trapped in U_A and never revisit U_B). For k large the step size is $< \delta/3$, so

$$\operatorname{dist}\big(z^{(\tau_k)}, \Omega_A\big) \le \|z^{(\tau_k)} - z^{(\tau_k - 1)}\| + \delta/3 < 2\delta/3 \implies \operatorname{dist}\big(z^{(\tau_k)}, \Omega_B\big) \ge \delta - 2\delta/3 = \delta/3,$$

hence $z^{(\tau_k)} \in C := K \setminus (U_A \cup U_B)$. The set C is compact and contains infinitely many $z^{(\tau_k)}$, so it contains a limit point of the sequence, i.e. $\Omega \cap C \neq \emptyset$, contradicting $\Omega \subseteq U_A \cup U_B$. □

C.4 Drift when the equilibrium is not unique

Remark 10 (Open gap: drift when the NE is not unique). When the equilibrium is not unique, Ω is a connected subset of $\mathrm{Fix}(\Phi) \cap \mathcal{L}$. Slow drift of the *state* along the NE polytope directions is not excluded (at the strategy level the trajectory is already pinned to the NE set by Theorem 2(ii)). We know of no route that closes this gap with the results of this paper. In particular, Corollary 1 cannot apply. Its hypothesis asks for a point of Ω at which Theorem 4 holds, and by Lemma 11 a strictly complementary non-degenerate equilibrium is automatically the unique one. In a game with several equilibria every equilibrium is therefore degenerate or fails strict complementarity, so the frozen-face theory of Section 4.2 applies at no point of Ω. The fixed-point set through such a point is the positive cone over the NE polytope (Proposition 1), a polyhedral set with positive-dimensional smooth strata but with corners at boundary equilibria, and the natural tool is a normally hyperbolic invariant manifold argument along a smooth stratum (Appendix G). Generic matrix games have a unique NE, in which case Theorem 3 resolves drift completely.

C.5 Proof of Corollary 1

Theorem 4 supplies the following capture property: if $z^\dagger \in \mathrm{Fix}(\Phi)$ corresponds to the strictly complementary equilibrium of Assumption 1 with $\eta(z^\dagger)\sigma_{\max}(B) < 1$ (so that $\rho^* := \max\{|\lambda| : \lambda \in \mathrm{spec}(D\Phi_S(z^\dagger)) \setminus \{1\}\} < 1$ by (E7)), then there exist an open neighborhood $U' \ni z^\dagger$, a constant $C > 0$, and $\bar{\rho} \in (\rho^*, 1)$ such that: whenever $z^{(T)} \in U'$ for some $T \ge t_+$, there is a fixed point $z_\infty \in \mathrm{Fix}(\Phi)$ with $\|z^{(t)} - z_\infty\| \le C\bar{\rho}^{\,t-T}$ for all $t \ge T$ (*entry implies capture*).

Proof of Corollary 1. Since $z^\dagger \in \Omega$, the very definition of an ω-limit point provides a time $T^* \geq t_+$ with $z^{(T^*)} \in U'$: every neighborhood of $z^\dagger$ is visited infinitely often. (Neither $\mathrm{dist}(z^{(t)}, \Omega) \to 0$ nor any hypothesis on the remaining members of Ω is needed here.) Capture then gives $z^{(t)} \to z_\infty \in \mathrm{Fix}(\Phi)$ at the linear rate $\bar{\rho}$.

For the gap decay: the map $z \mapsto (X(z), Y(z))$ is locally Lipschitz on a neighborhood of z_∞, by Lemma 6 (the map Φ and its intermediate maps are smooth on the frozen-support neighborhood) or, more directly, by the 1-Lipschitz property of the γ-shift (Lemma 19) and local Lipschitzness of each composition; and gap is globally Lipschitz in (x, y) (with constant of order $2\|A\|$, up to norm-equivalence factors) and vanishes at $(X, Y)(z_\infty) \in \mathrm{NE}$. Hence $\mathrm{gap}(x^{(t)}, y^{(t)}) \leq L \, \|z^{(t)} - z_\infty\| \leq C' \bar{\rho}^{\, t - T^*}$.

Finally, why no assumption on the rest of Ω is needed: the argument used only that $z^\dagger$ itself is an ω-limit point. Once the trajectory enters U' it is captured, and the remaining members of Ω (including any unstable ray points) are never visited again; a posteriori $\Omega = \{z_\infty\}$. Under Assumptions 1 and 2 the hypothesis holds automatically: Theorem 3 makes Ω the single point $\bar{z}$, and $\bar{z}$ satisfies the hypotheses of Theorem 4 with $m \geq 2$ precisely when $\eta(\bar{z})\sigma_{\max}(B) < 1$, i.e. $\sigma_{\max}(B) < \sqrt{N_{1,\infty} N_{2,\infty} / (\|x^*\|_2 \|y^*\|_2)}$, which is Assumption 2 evaluated at the saturation level; for $m = 1$ the capture is finite-step by Lemma 6. □

C.6 Proofs for self-stabilization (Proposition 2)

Throughout this subsection $z^\dagger = (c_1 x^*, c_2 y^*)$ is a fixed point whose equilibrium (x^*, y^*) is strictly complementary (with gap $\delta > 0$ in the notation of Lemma 6) and nondegenerate, i.e. satisfies Assumption 3 (by Lemma 11 it is then the unique equilibrium, but the arguments below use only the face-local properties); $S = S_1 \times S_2$ is the support pair and $\mathcal{M}_S = \{z : z_{S^c} = 0\}$ the support subspace. We write $\eta = \eta(z^\dagger) = 1/\sqrt{c_1 c_2}$ and $q_k = \eta^2 \sigma_k^2$ for the singular values σ_k of B.

Lemma 23 (Analytic ambient extension near the ray; strengthening of Lemma 6)**.** *The point $z^\dagger$ lies on the boundary of the nonnegative state space $\mathcal{Z}$ whenever some coordinate is off-support, so no $\mathbb{R}^{n_1+n_2}$-open neighborhood of $z^\dagger$ is contained in $\mathcal{Z}$. There exist, however, an open neighborhood $U \subset \mathbb{R}^{n_1+n_2}$ of $z^\dagger$ and a real-analytic map $\hat{\Phi} : U \to \mathbb{R}^{n_1+n_2}$ with $\hat{\Phi} = \Phi$ on $U \cap \mathcal{Z}$ (we keep writing Φ for $\hat{\Phi}$ below), such that:*

- *(a) for $i \notin S_1$: $\Phi(z)_{1,i} \equiv 0$ on U (off-support components vanish exactly in one step), and symmetrically for player 2;*
- *(b) $D\Phi(z^\dagger)$ is block-triangular, $\begin{pmatrix} J_S & * \\ 0 & 0 \end{pmatrix}$ in the S/S^c coordinate split, so $\mathrm{spec}(D\Phi(z^\dagger)) = \mathrm{spec}(J_S) \cup \{0\}$, the eigenvalue 0 having multiplicity $n_1 + n_2 - |S_1| - |S_2|$ and being absent when both players are fully mixed ($S^c = \emptyset$, so $D\Phi(z^\dagger) = J_S$), where $J_S = D\Phi_S(z^\dagger) = I + L + L^2$ as in Lemma 9;*
- *(c) J_S is* nonsingular*: its eigenvalues are $\mu = 1 + \lambda + \lambda^2$ with $\lambda \in \mathrm{spec}(L)$ purely imaginary ($\lambda = \pm i\sqrt{q}$, $q \geq 0$, by Lemma 10), and $|\mu|^2 = (1-q)^2 + q = 1 - q + q^2 \geq 3/4 > 0$ for all $q \geq 0$ (the minimum $3/4$ is attained at $q = 1/2$).*

Proof. (a) and the extension. Define $\hat{\Phi}$ by running one round of Algorithm 1 with three substitutions, each of which agrees with the original operation on feasible states near $z^\dagger$: (1) the ℓ_1 normalizations $\tilde{r}_i / \|\tilde{r}_i\|_1$ and $p_i / \|p_i\|_1$ are replaced by $\tilde{r}_i / (\mathbf{1}^\top \tilde{r}_i)$ and $p_i / (\mathbf{1}^\top p_i)$; (2) the γ-shift is replaced by the solution of the quadratic $\sum_{j \in S_1} (v_j - \gamma)^2 = \tilde{r}_1^\top \tilde{r}_1$ with the active set frozen at S_1; (3) every clipping $[\cdot]_+$ is replaced by the coordinate projection onto S_i (in-support coordinates kept, off-support coordinates set to 0). Each substituted operation is real-analytic on a neighborhood of $z^\dagger$: $\mathbf{1}^\top \tilde{r}_1 \to c_1 > 0$ and $\tilde{r}_1^\top \tilde{r}_1 \to c_1^2 \|x^*\|_2^2 > 0$, the frozen-active-set quadratic has positive discriminant and an analytic root, and projections are linear. Hence $\hat{\Phi}$ is a composition of analytic maps on some $\mathbb{R}^{n_1+n_2}$-open $U \ni z^\dagger$. It remains to check $\hat{\Phi} = \Phi$ on $U \cap \mathcal{Z}$, shrinking U if necessary. The proof of Lemma 6 (under Assumption 3) for the strict version of Algorithm 1 gives, on feasible states near $z^\dagger$: for off-support components, $r_{1,i}(z) \leq -\delta/2$ and $|g_{1,i}(z)| < \delta/4$, hence $r_{1,i} + g_{1,i} < -\delta/4 < 0$, so both $([r_1]_+)_i$ and $([r_1 + g_1]_+)_i$ equal 0, as does the projection in (3); for in-support components,

$r_{1,j}(z) \to c_1 x_j^* > 0$ and $g \to 0$, so $r_{1,j} + g_{1,j} > 0$ and clipping is the identity, again as in (3). Because a feasible $\tilde{r}_1$ near $z^\dagger$ is nonnegative, $\mathbf{1}^\top \tilde{r}_1 = \|\tilde{r}_1\|_1$, and likewise for p_1, which is nonnegative and supported on S_1; so (1) agrees with the original normalizations. Finally the active set of the original γ-shift is exactly S_1 on feasible states near $z^\dagger$ (in-support: $v_j - \gamma = r_{1,j} > 0$; off-support: $< -\delta/2$), so the original γ solves the same frozen quadratic as (2), and $\tilde{r}_1^\top \tilde{r}_1 = N_1(z)^2$. Thus $\hat{\Phi} = \Phi$ on $U \cap \mathcal{Z}$, and (a) holds for $\hat{\Phi}$ on all of U by construction. Every Φ-orbit that stays in $U \cap \mathcal{Z}$ is a $\hat{\Phi}$-orbit, which is the only property the manifold theorem below needs.

(b). By (a), the S^c-rows of $D\Phi$ vanish identically on U; $\mathcal{M}_S = \{z_{S^c} = 0\}$ is a linear subspace, and the S-block is the derivative of the restriction $\Phi_S = \Phi|_{\mathcal{M}_S}$, i.e. J_S.

(c). $\mu = 0$ would require $\lambda^2 + \lambda + 1 = 0$, i.e. $\lambda = e^{\pm 2\pi i/3}$, of modulus 1 and not purely imaginary, contradicting $\mathrm{spec}(L) \subset i\mathbb{R}$ (Lemma 10). For the quantitative bound, $\lambda = i\sqrt{q}$ gives $\mu = 1 - q + i\sqrt{q}$ and $|\mu|^2 = (1-q)^2 + q = 1 - q + q^2$, a quadratic in q minimized at $q = 1/2$ with value $3/4$. □

Remark 11 (Applicability of the center-stable manifold theorem to the non-invertible map Φ). Near the ray, $\det D\Phi = 0$ whenever some coordinate is off-support (Lemma 23(b)), so Φ is then not invertible, whereas textbook statements of the (center-)stable manifold theorem, including the form invoked by Lee et al. (2019), are typically phrased for diffeomorphisms. The version we need is the local center-stable manifold theorem for C^1 *maps* at a fixed point, with no invertibility hypothesis: a C^1 manifold W^{cs}_{loc}, tangent to the center-stable eigenspace, exists and contains *every orbit whose forward iterates remain in a sufficiently small neighborhood of the fixed point*. This map version is supplied by the graph-transform construction (Hirsch et al., 1977): both the construction of the pseudo-stable manifold and the trapping inclusion use forward iteration only and never invert Φ. The eigenvalue 0 of $D\Phi$ sits on the center-stable side of the splitting, which is exactly where the graph transform places contracted directions.

Proposition 5 (Unstable ray points have thin basins; $\mathcal{M}_S$-local version)**.** *Let $z^\dagger$ be as above with $q_{\max} := \eta^2 \sigma_{\max}(B)^2 > 1$, and set $d_u := 2\,\#\{k : \eta\sigma_k > 1\} \geq 2$ (each violating singular value contributes a conjugate eigenvalue pair with $|\mu|^2 = 1 - q + q^2 > 1$ when $q > 1$). Then:*

(1) (Full space: convergent orbits lie on a center-stable manifold.) *There exist a neighborhood $B_\dagger \ni z^\dagger$ and a C^1 local manifold $W^{cs}_{loc}(z^\dagger)$ with $\dim W^{cs}_{loc} = d - d_u$ such that*

$$\{z \in B_\dagger :\ \Phi^t(z) \in B_\dagger\ \forall t \geq 0\} \ \subseteq\ W^{cs}_{loc}(z^\dagger).$$

In particular the tail of every trajectory converging to $z^\dagger$ is contained in $W^{cs}_{loc}(z^\dagger)$.

(2) (Inside $\mathcal{M}_S$: the basin has measure zero.) *$\mathcal{M}_S$ is Φ-invariant near $z^\dagger$ (Lemma 23(a)), J_S is nonsingular (Lemma 23(c)), and $z \mapsto D\Phi_S(z)$ is continuous, so there is an $\mathcal{M}_S$-open neighborhood $V \ni z^\dagger$ on which $\Phi_S|_V$ is a local diffeomorphism. The local basin*

$$\big\{z \in V :\ \Phi_S^t(z) \in V\ \forall t,\ \Phi_S^t(z) \to z^\dagger\big\} \ \subseteq\ \bigcup_{T \geq 0} \Phi_S^{-T}\big(W^{cs}_{loc}(z^\dagger) \cap \mathcal{M}_S\big)$$

is contained in a countable union of preimages of local C^1 manifold patches of dimension $\dim \mathcal{M}_S - d_u < \dim \mathcal{M}_S$; local diffeomorphisms pull back Lebesgue-null sets to Lebesgue-null sets, so this basin has Lebesgue measure zero in $\mathcal{M}_S$. The same holds at every individual point of the step-size-violating ray set $R_{>1} := \{(c_1 x^*, c_2 y^*) : c_1, c_2 > 0,\ \eta(c_1, c_2)\,\sigma_{\max}(B) > 1\} \subset \mathcal{M}_S$. Moreover, for every compact *subsegment $R_K \subset R_{>1}$ there are an $\mathcal{M}_S$-open neighborhood $V_K \supset R_K$ and a Lebesgue-null set $E_K \subset \mathcal{M}_S$ such that any Φ_S-orbit that remains in V_K forever and converges to a point of R_K has its tail contained in E_K; in particular the local basin*

$$\big\{z \in V_K :\ \Phi_S^t(z) \in V_K\ \forall t \geq 0,\ \lim_t \Phi_S^t(z) \in R_K\big\}$$

is Lebesgue-null in $\mathcal{M}_S$. The noncompact ends of the ray set ($c_i \to 0$, $c_i \to \infty$, and the approach to the boundary $\eta\,\sigma_{\max}(B) = 1$) are excluded *from the uniform statement: the neighborhood radii below degenerate there, and the source notes leave those ends open. It does* not *bound the set of all $\mathcal{M}_S$-states whose orbits eventually converge to the segment. Pulling the null set back along the portion of an orbit prior to its entry into V_K would require Φ_S to be a local diffeomorphism away from the ray, and that is not established (the*

clipping pattern is not frozen there). This is the obstruction of Remark 13, acting inside $\mathcal{M}_S$. As discussed after Proposition 2, the fact that a convergent trajectory's tail lies in the null set is true by construction and has no excluding force for that trajectory.

Proof. Part 1. By Lemma 23, the ambient extension $\hat{\Phi}$ (written Φ) is $C^\omega \subset C^1$ on an $\mathbb{R}^{n_1+n_2}$-open neighborhood of $z^\dagger$, agrees with the algorithm on feasible states there, and $\mathrm{spec}(D\Phi(z^\dagger)) = \mathrm{spec}(J_S) \cup \{0\}$, the eigenvalue 0 absent when $S^c = \emptyset$. Split the spectrum into the center-stable part $\{|\lambda| \le 1\}$ (which contains the eigenvalue 0 of the S^c directions, if any, and all eigenvalue pairs with $q_k \le 1$) and the unstable part $\{|\lambda| > 1\}$ (the d_u eigenvalues coming from $q_k > 1$, since $|\mu|^2 = 1 - q + q^2 > 1 \iff q > 1$). The local center-stable manifold theorem for C^1 maps, in the non-invertible form discussed in Remark 11 (Hirsch et al., 1977), in the application form used for saddle avoidance of first-order methods (Lee et al., 2019, Thm. 2) (cf. Daskalakis & Panageas, 2018), yields $B_\dagger$ and $W^{cs}_{loc}(z^\dagger)$ of dimension $d - d_u$ containing all orbits that remain in $B_\dagger$ forever. A trajectory converging to $z^\dagger$ eventually remains in $B_\dagger$, so its tail lies in $W^{cs}_{loc}(z^\dagger)$. The dimension count: eigenvalues of J_S with $|\mu| > 1$ arise exactly from the conjugate pairs with $q_k > 1$ by (E7), and the S^c directions (eigenvalue 0) count as stable; thus $\dim W^{cs} \le d - 2 < d$.

Part 2, single point. As stated: invariance of $\mathcal{M}_S$ and one-step exact vanishing of off-support coordinates are Lemma 23(a); nonsingularity of J_S plus continuity of $z \mapsto D\Phi_S(z)$ give, via the inverse function theorem, an $\mathcal{M}_S$-neighborhood V on which Φ_S is a local diffeomorphism. Any point of the local basin reaches $W^{cs}_{loc} \cap \mathcal{M}_S$ (a C^1 patch of dimension $\dim \mathcal{M}_S - d_u < \dim \mathcal{M}_S$, hence Lebesgue-null in $\mathcal{M}_S$) after some finite number T of steps while remaining in V; therefore the basin is covered by $\bigcup_{T \ge 0} \Phi_S^{-T}(W^{cs}_{loc} \cap \mathcal{M}_S)$. Each term is the preimage of a null set under a finite composition of local diffeomorphisms, hence null; a countable union of null sets is null.

Part 2, uniform version on a compact subsegment R_K. Two points need verification beyond a bare "countable dense net" argument: (i) the manifold and diffeomorphism constructions must run in neighborhoods whose size is bounded below *uniformly* along R_K; (ii) an orbit may converge to a limit $z' \in R_K$ that is not a net point, and one must say why its tail still lands in the null set attached to a *nearby* net point.

(i) *Uniform radii.* Since R_K is compact and contained in the open set $\{w : \eta(w)\,\sigma_{\max}(B) > 1\}$, there is $\kappa > 0$ with $q_{\max}(w) \ge 1 + \kappa$ on R_K, so the top conjugate eigenvalue pair of $J_S(w)$ satisfies $|\mu|^2 = 1 - q_{\max} + q_{\max}^2 \ge 1 + \kappa + \kappa^2$ uniformly; at least this pair is uniformly expanding on all of R_K. For each $w \in R_K$ choose a splitting radius $r_*(w) \in \big(1, \sqrt{1 + \kappa + \kappa^2}\big)$ avoiding the finitely many moduli in $\mathrm{spec}(D\Phi(w))$; by continuity of the spectrum the same $r_*(w)$ keeps a positive spectral gap on a relatively open piece of R_K around w, and compactness extracts a finite cover of R_K by such pieces, each with a fixed splitting radius and a uniform gap. On each piece, the graph-transform construction of the local pseudo-stable manifold for C^1 maps (Remark 11; no invertibility needed) produces, at every ray point w, a C^1 patch $W^{cs,r_*}_{loc}(w)$ of dimension $d - d_{su}(w) \le d - 2$ (where $d_{su}(w) \ge 2$ counts the eigenvalues with $|\mu| > r_*$) containing every orbit that stays in the ball $B(w, 2\rho_w)$ forever, since bounded orbits grow slower than any geometric rate $r_*^t > 1$. The admissible radius ρ_w depends only on the spectral gap at r_*, the norms of the associated spectral projections, and the modulus of continuity of $D\Phi$ on a fixed compact neighborhood of R_K (on which Lemma 23 applies with uniform constants: the complementarity margins in its proof are continuous in w and positive, hence uniformly bounded below on R_K). All of these are continuous in w, hence uniformly controlled on the compact R_K. Likewise $|\det D\Phi_S| \ge (3/4)^{\dim \mathcal{M}_S/2} > 0$ on R_K by Lemma 23(c), so by continuity there is a uniform $\rho' > 0$ such that Φ_S is a local diffeomorphism on the ρ'-neighborhood of R_K in $\mathcal{M}_S$. Set $\rho := \min(\inf_{w \in R_K} \rho_w, \rho') > 0$ and let V_K be the ρ-neighborhood of R_K in $\mathcal{M}_S$. The uniform codimension delivered here is $d_{su} \ge 2$, which may be smaller than the pointwise $d_u(w)$ of Part 1; codimension ≥ 2 is all the measure-zero conclusion needs. All uniform bounds degenerate at the noncompact ends of $R_{>1}$ (as $c_i \to 0$ or ∞ the state approaches the boundary of $\mathcal{Z}$ or the map data blow up, and as $\eta\sigma_{\max} \to 1$ the gap κ vanishes), which is why the statement is restricted to compact R_K.

(ii) *Convergence to a non-net point.* By compactness choose a finite net $w_1, \dots, w_J \in R_K$ whose $\rho/2$-balls cover R_K. Suppose a Φ_S-orbit $\{\Phi_S^t(z)\}$ remains in V_K and converges to some $z' \in R_K$, and pick j with $\|z' - w_j\| < \rho/2$. Since the orbit converges to z', there is a finite T from which on it stays in $B(z', \rho/2) \subseteq B(w_j, \rho) \subseteq B(w_j, 2\rho_{w_j})$. The inclusion of (i) at the *net point* w_j requires only

that the orbit remain in that ball forever, not that it converge to w_j; hence $\Phi_S^t(z) \in W_{loc}^{cs,r_*}(w_j) \cap \mathcal{M}_S$ for all $t \geq T$, a C^1 patch of dimension $\leq \dim \mathcal{M}_S - 2$, Lebesgue-null in $\mathcal{M}_S$. Define

$$E_K := \bigcup_{j \leq J} \bigcup_{T \geq 0} \big(\Phi_S|_{V_K}\big)^{-T} \big(W_{loc}^{cs,r_*}(w_j) \cap \mathcal{M}_S\big).$$

The tail of the orbit lies in E_K (take $T = 0$ terms), and the starting point z lies in E_K as well, since the whole orbit up to time T stays in V_K, where Φ_S has everywhere invertible derivative, so each preimage is taken under finite compositions of local diffeomorphisms and is null. E_K is a countable union of null sets, hence Lebesgue-null in $\mathcal{M}_S$, and the local basin of R_K is contained in E_K. □

Proof of Proposition 2. Parts (a) and (b) are Proposition 5 parts 1 and 2, applied at $z^\dagger$ (single-point statement) and at a compact subsegment through $z^\dagger$ (uniform statement). □

Remark 12 (Trapped tails lie in the null set by construction). Via Lemma 6, a trajectory is absorbed into $\mathcal{M}_S$ in one step near a ray, so the tail of a trajectory that converges to a step-size-violating point $z_\infty \in \mathrm{Fix}(\Phi)$ remains in the neighborhood V of Proposition 5(2), hence in the exceptional Lebesgue-null set. The measure-zero statement therefore excludes no individual orbit; its force is the genericity heuristic of Remark 14.

Remark 13 (Open gap: no full-space measure-zero statement). Near the ray, whenever some coordinate is off-support, Φ maps the S^c coordinates to exactly 0 (Lemma 23(a)), so $\det D\Phi = 0$ and Φ is *not invertible*; preimages of Lebesgue-null sets under non-invertible maps can have positive measure. The "almost every initial point avoids strict saddles" argument of Lee et al. (2019) requires $\det D\Phi \neq 0$ everywhere and therefore does *not* apply to our map on the full state space. This is why Proposition 5(2) is asserted only inside $\mathcal{M}_S$ (when both players are fully mixed, $\mathcal{M}_S$ is the whole state space and the statement is already full-space). A full-space version would need a different argument: for instance, absolute continuity properties of the push-forward of Lebesgue measure onto a neighborhood of $\mathcal{M}_S$ after one step, or a direct analysis of the preimage structure of Φ away from equilibrium. We leave this open.

Remark 14 (Open gap: fixed initialization). The algorithm's standard initialization $\tilde{r}^{(1)} = \mathbf{0}$ is a single deterministic initial condition, and the genericity statement of Proposition 5(2) has no direct force for it. Its scope covers perturbed initial conditions or mid-trajectory states in $\mathcal{M}_S$: there the exceptional set is Lebesgue-null (of codimension ≥ 2), and landing of one specific trajectory exactly on such a set is a non-generic coincidence. Any statement about the $\tilde{r}^{(1)} = \mathbf{0}$ trajectory itself remains heuristic in this sense (Remark 12).

Remark 15 (Open gap: the boundary $q = 1$ is non-hyperbolic). When $\eta\sigma_{\max}(B) = 1$, i.e. $q_{\max} = 1$, the corresponding eigenvalues are $\mu = \pm i$ (fourth roots of unity): the fixed point is non-hyperbolic, higher-order terms on the center manifold decide between weak stability, weak rotation, and weak instability, and the linearization is inconclusive. Proposition 5 does not exclude convergence to such boundary points at a sublinear rate. We have not observed this in the numerical suite (the spectral data allow checking the distribution of $q_{\max}$ against 1; see Figure 1), but it remains open.

D Proofs for the Ratio Certificate

D.1 Proofs

This subsection proves the increment identity, the two-sided increment bound and the certificate itself. It opens with the precise form of the frozen local regime, described informally in Section 4.3: the list (P1)–(P6) below fixes the constants and the coordinates that the rest of the section uses. The gap–distance sharpness of Lemma 30 rests on a separate first-order analysis of the gap map, and its proof, together with the Lipschitz upper bound it builds on, is in Appendix D.

D.1.1 Standing conventions: the frozen local regime

Throughout this subsection and Appendix D, Assumptions 1 and 2 are in force and we work in the frozen local regime constructed in the proof of Theorem 4. We collect the facts imported from there; none is reproved here.

(P1) Under Assumption 1 the equilibrium support is square, $|S_1| = |S_2| =: m$, and $m \geq 2$; the pure saddle $m = 1$ is set aside, since there $B = 0$, the transverse space W introduced in (P2) below is trivial, and convergence is exact in finitely many steps, so nothing below is needed. The value-centered support matrix $B = \tilde{A} - v^* \mathbf{1}\mathbf{1}^\top$ of Lemma 10 satisfies $B\psi = 0$, $B^\top \xi = 0$, with $\ker B = \operatorname{span}\{\psi\}$ and $\ker B^\top = \operatorname{span}\{\xi\}$; its nonzero singular values are $\sigma_{\max} = \sigma_1 \geq \cdots \geq \sigma_{m-1} = \sigma^+_{\min} > 0$, with singular vectors $Bv_k = \sigma_k u_k$, $B^\top u_k = \sigma_k v_k$.

(P2) In the support coordinates of $\mathcal{M}_S$, the state space decomposes orthogonally as $E \oplus W$, where $E = \operatorname{span}\{(\xi, 0), (0, \psi)\}$ contains the fixed-point cone E_+ (Proposition 1) and $W = \bigoplus_k \operatorname{span}\{(u_k, 0), (0, v_k)\}$. Orthogonality follows from $B^\top \xi = 0$ and $B\psi = 0$. We write P_E, P_W for the orthogonal projections and $w = (w_1, w_2) \in W$, so that $w_1 \in \operatorname{range}(B)$, $w_2 \in \operatorname{range}(B^\top)$ and $\|w\|^2 = \|w_1\|^2 + \|w_2\|^2$.

(P3) There are radii $r_1, \delta_0 > 0$, fixed in the proof of Theorem 4, such that the regime is the set of states $z = z_f + w$ with $z_f = (a\xi, b\psi) \in E_+$, $\|z_f - z^*\| \leq r_1$ (hence $a \in [c_1/2, 3c_1/2]$, $b \in [c_2/2, 3c_2/2]$), and $\|w\| \leq \delta_0$. On the regime the effective step size $\eta(a, b) = 1/\sqrt{ab}$ ranges in $[\eta_-, \eta_+]$ with $0 < \eta_- \leq \eta_+ < \infty$.

(P4) On the regime the clipping pattern is frozen (Lemma 6): for each player, on-support coordinates of both $r^{(t)}$ and $r^{(t)} + g^{(t)}$ are strictly positive (no clipping at either the prediction or the update stage), while every off-support coordinate $j \notin S_i$ satisfies $r^{(t)}_{i,j} \leq -\delta/2$ and $|g^{(t)}_{i,j}| \leq \delta/4$, hence $r^{(t)}_{i,j} + g^{(t)}_{i,j} \leq -\delta/4 < 0$ (clipped exactly to 0), where $\delta > 0$ is the strict-complementarity margin of Assumption 1. Consequently the one-round map restricted to $\mathcal{M}_S$ is a composition of analytic maps on the regime.

(P5) By analyticity and compactness of the closed regime, the maps $z \mapsto (x(z), y(z))$ (main-iterate strategies) and $z \mapsto g_i(z)$ are C^2 with uniform first- and second-derivative bounds; the second-order constants below (C_q, K_g, C_3) are finite and uniform on the regime. They are *not* exhibited in closed form. This affects only the size of the thresholds δ_2, δ_3, and not the form of any constant that enters κ.

(P6) The exact increment decomposition of Zhang et al. (2025a) reads: for each player i,

$$\dot{N}_i(t) = \left\|g_i^{(t)}\right\|^2 - D_i(t), \qquad D_i(t) = \left\|\left[r_i^{(t)}\right]_+ + g_i^{(t)}\right\|^2 - \left\|\left[r_i^{(t)} + g_i^{(t)}\right]_+\right\|^2 \geq 0.$$

Since E is a linear subspace, P_E is the (orthogonal) nearest-point map, and on the regime $P_E z$ lies in the *interior* of the cone E_+ by (P3); hence

$$\operatorname{dist}(z, E) = \operatorname{dist}(z, \mathrm{Fix}) = \|w\|. \tag{28}$$

D.1.2 Proof of Lemma 24

Lemma 24 (Frozen-regime increment identity)**.** *In the frozen local regime, for each player i and every step t,*

$$\dot{N}_i(t) \;=\; \sum_{j \in S_i} \big(g^{(t)}_{i,j}\big)^2 \;=\; \big\|g_i^{(t)}\big|_{S_i}\big\|^2,$$

where $g_i^{(t)}$ is player i's update vector; the increment equals the on-support update energy, with no error term.

Proof of Lemma 24. Fix a player i and drop the index. Start from the decomposition (P6). We compute the slack $D(t)$ componentwise under the frozen pattern (P4).

On-support coordinates $j \in S_i$. Here $r_j > 0$ (no clipping in the prediction step) and $r_j + g_j > 0$, so $([r]_+ + g)_j = r_j + g_j = ([r + g]_+)_j$: the two squares agree and the contribution to $D(t)$ cancels exactly.

Off-support coordinates $j \notin S_i$. Here $r_j \leq -\delta/2 < 0$ so $([r]_+)_j = 0$ and $([r]_+ + g)_j = g_j$; and $r_j + g_j \leq -\delta/4 < 0$ so $([r + g]_+)_j = 0$. The contribution to $D(t)$ is therefore $g_j^2 - 0 = g_j^2$.

Summing, $D(t) = \sum_{j \notin S_i} g_j^2$, and hence

$$\dot{N}_i(t) = \|g\|^2 - D(t) = \sum_{j \in S_i} g_j^2,$$

which is the claim. □

Remark 16 (Numerical verification). The identity is exact, and this exactness is a hard prediction of the frozen-pattern analysis. Our numerical suite checks only its visible consequence: the exact-zero pattern of $\tilde{r}$ freezes in finite time on 99.5% of instances (Section 5); the full clipping pattern behind the identity was not logged, so the on-support energy is used as an empirical proxy for $\dot{N}_{\mathrm{tot}}$ (Appendix E). The pattern-frozen identity is what distinguishes this lemma from the anchor decomposition (P6), whose slack term is in general strictly between 0 and $\|g\|^2$.

D.1.3 First-order expansion of the update vector

Lemma 25 (Update expansion). *In the frozen regime, for $z = z_f + w$ with $z_f = (a\xi, b\psi)$ and $w = (w_1, w_2) \in W$,*

$$g_1\big|_{S_1} = -\eta^2\, BB^\top w_1 + O\big(K_g\|w\|^2\big), \qquad g_2\big|_{S_2} = -\eta^2\, B^\top B\, w_2 + O\big(K_g\|w\|^2\big),$$

with $\eta = \eta(a,b)$ and K_g the uniform constant of (P5).

Proof. The map $z \mapsto g_1(z)$ is analytic on the regime (P4) and vanishes identically on E_+ (at fixed points $y = \tilde{y}$, so the residual and hence g are zero). Therefore $g_1(z_f + w) = Dg_1(z_f)[w] + O(K_g\|w\|^2)$ with a uniform constant by (P5), and it remains to compute the differential in the direction w. We use the cancellation identities established for Lemma 9 (all evaluated at the fixed point z_f; the symbol m_i below denotes player i's payoff vector, not the support size):

(i) The residual is $\mathrm{res}_1 = A(y - \tilde{y})$; since both y and $\tilde{y}$ are supported in S_2 on the regime, $\mathrm{res}_1|_{S_1} = \tilde{A}\,(y - \tilde{y})|_{S_2}$.

(ii) Perturbing $(\tilde{r}_1, \tilde{r}_2)$ by (w_1, w_2): the pre-iterate differential is $d\tilde{y} = P_\psi w_2/b$ (normalization at level $\|\tilde{r}_2^*\|_1 = b$, with $P_\psi = I - \psi\mathbf{1}^\top$). For the main iterate $r_2 = \tilde{r}_2 + m_2 - \gamma_2\mathbf{1}$ we have $dm_2 = -\tilde{A}^\top P_\xi\, w_1/a = -B^\top w_1/a$ (Lemma 10) and $d\gamma_2 = 0$ (cancellation, Lemma 9); hence $dy = P_\psi\big(w_2 - B^\top w_1/a\big)/b$ and

$$d(y - \tilde{y}) = -\tfrac{1}{ab}\, P_\psi B^\top w_1.$$

(iii) Therefore the first-order term of $\mathrm{res}_1|_{S_1}$ is $-\frac{1}{ab}\tilde{A}P_\psi B^\top w_1 = -\eta^2 BB^\top w_1$, using $\tilde{A}P_\psi = B$ (Lemma 10).

(iv) The centering term: to first order $\langle \mathrm{res}_1, x\rangle = \xi^\top(-\eta^2 BB^\top w_1) = -\eta^2(B^\top\xi)^\top B^\top w_1 = 0$ by (P1), and x is exactly supported in S_1 on the regime; hence $g_1|_{S_1} = \mathrm{res}_1|_{S_1} - \langle \mathrm{res}_1, x\rangle\mathbf{1}$ has the same first-order term as $\mathrm{res}_1|_{S_1}$.

Player 2 is symmetric. □

D.1.4 Proof of Lemma 26

Lemma 26 (Two-sided increment bound). *There exist $0 < n_- \le n_+$ and a threshold $\delta_3 > 0$ such that, in the frozen regime with $\|w^{(t)}\| \le \delta_3$, for every single step t,*

$$n_-\,\|w^{(t)}\|^2 \;\le\; \dot{N}_{\mathrm{tot}}(t) \;\le\; n_+\,\|w^{(t)}\|^2, \qquad n_- = \tfrac{1}{2}\,\eta_-^4\,(\sigma^+_{\min})^4, \quad n_+ = 2\,\eta_+^4\,\sigma^4_{\max}.$$

Proof of Lemma 26. By Lemma 24 and Lemma 25,

$$\dot{N}_{\mathrm{tot}}(t) = \big\|g_1|_{S_1}\big\|^2 + \big\|g_2|_{S_2}\big\|^2 = \eta^4\Big(\big\|BB^\top w_1\big\|^2 + \big\|B^\top B w_2\big\|^2\Big) + O\big(C_3\|w\|^3\big),$$

with C_3 uniform on the regime (cross terms of the expansion; (P5)). The key structural point is that $w \in W$ by definition of the decomposition (P2), so $w_1 \in \mathrm{range}(B) = \mathrm{span}\{u_k\}$ and $w_2 \in \mathrm{range}(B^\top) = \mathrm{span}\{v_k\}$. Writing $w_1 = \sum_k \alpha_k u_k$,

$$\big\|BB^\top w_1\big\|^2 = \sum_k \sigma_k^4\alpha_k^2 \;\in\; \Big[\big(\sigma^+_{\min}\big)^4\|w_1\|^2,\; \sigma^4_{\max}\|w_1\|^2\Big],$$

and symmetrically for w_2; adding and using $\|w\|^2 = \|w_1\|^2 + \|w_2\|^2$ gives

$$\eta_-^4 \big(\sigma_{\min}^+\big)^4 \|w\|^2 - C_3\|w\|^3 \;\le\; \dot N_{\mathrm{tot}}(t) \;\le\; \eta_+^4\, \sigma_{\max}^4 \|w\|^2 + C_3\|w\|^3.$$

For $\|w\| \le \delta_3 := \eta_-^4 (\sigma_{\min}^+)^4/(2C_3)$ the cubic error is absorbed: the lower bound loses at most a factor $\frac{1}{2}$, and since $\delta_3 \le \eta_+^4\sigma_{\max}^4/C_3$, the upper bound gains at most a factor 2. This yields the stated $n_- = \frac{1}{2}\eta_-^4(\sigma_{\min}^+)^4$ and $n_+ = 2\eta_+^4\sigma_{\max}^4$. □

D.1.5 Proof of Theorem 5 and Corollary 2

Proof of Theorem 5. For every $t \ge t_{\mathrm{reg}}$ the hypotheses put $z^{(t)}$ in the regime with $\|w^{(t)}\| \le \min(\delta_2, \delta_3)$, so Lemmas 30 and 26 give

$$c_-\|w^{(t)}\| \le \mathrm{gap}(t) \le L_{\mathrm{gap}}\|w^{(t)}\|, \qquad n_-\|w^{(t)}\|^2 \le \dot N_{\mathrm{tot}}(t) \le n_+\|w^{(t)}\|^2.$$

The hypothesis $\dot N_{\mathrm{tot}}(t_1) > 0$ forces $w^{(t_1)} \ne 0$ through the upper increment bound, hence $\mathrm{gap}(t_1) \ge c_-\|w^{(t_1)}\| > 0$, and every ratio below is well defined. Dividing the gap bounds at t_2 and t_1 and eliminating the amplitudes via the increment bounds,

$$\frac{\mathrm{gap}(t_2)}{\mathrm{gap}(t_1)} \;\le\; \frac{L_{\mathrm{gap}}\|w^{(t_2)}\|}{c_-\|w^{(t_1)}\|} \;\le\; \frac{L_{\mathrm{gap}}}{c_-}\sqrt{\frac{\dot N_{\mathrm{tot}}(t_2)/n_-}{\dot N_{\mathrm{tot}}(t_1)/n_+}} \;=\; \kappa\sqrt{\frac{\dot N_{\mathrm{tot}}(t_2)}{\dot N_{\mathrm{tot}}(t_1)}},$$

and the lower bound is symmetric. All four constants $L_{\mathrm{gap}}, c_-, n_\pm$ are fixed attributes of the regime ((P3)–(P5)): none depends on t_1, t_2 or on $\|w^{(t)}\|$, and no factor $1/(1-\rho^*)$ enters: they are evaluated on the fixed scale rectangle of (P3), and the amplitudes through which the trajectory contracts have been cancelled. The factorization of κ and the bound $\sqrt{n_+/n_-} \le 2(\eta_+/\eta_-)^2(\sigma_{\max}/\sigma_{\min}^+)^2$ are immediate from the formulas for $n_\pm$. □

Proof of Corollary 2. From the same sandwich, $C(t) := \log \dot N_{\mathrm{tot}}(t) - 2\log \mathrm{gap}(t)$ satisfies $\log(n_-/L_{\mathrm{gap}}^2) \le C(t) \le \log(n_+/c_-^2)$, an interval of width $\log\big((n_+/n_-)(L_{\mathrm{gap}}/c_-)^2\big) = 2\log\kappa$. For the least-squares slope over a window, write $x_t = \log \mathrm{gap}(t)$, $y_t = 2x_t + C(t)$; then $\hat\theta = \mathrm{Cov}(x,y)/\mathrm{Var}(x) = 2 + \mathrm{Cov}(x,C)/\mathrm{Var}(x)$, and Cauchy–Schwarz bounds the second term by $\mathrm{sd}(C)/\mathrm{sd}(x)$. A variable confined to an interval of width $2\log\kappa$ has population standard deviation at most half the width, $\log\kappa$ (Popoviciu's inequality):

$$\big|\hat\theta - 2\big| \le \frac{\mathrm{sd}(C)}{\mathrm{sd}(x)} \le \frac{\log\kappa}{\mathrm{sd}(\log \mathrm{gap})}.$$

As the window spans more decades of gap decay, $\mathrm{sd}(\log\mathrm{gap}) \to \infty$ and $\hat\theta \to 2$: the recorded times are distinct integers and, in the regime, $\|w^{(t+s)}\|_* \le \theta^s\|w^{(t)}\|_*$ (Lemma 16) with $\mathrm{gap} \asymp \|w\|$ (Lemma 30), so $\mathrm{gap}(t+s) \le C\theta^s \mathrm{gap}(t)$ from every t: at most a bounded number of records fall in any one decade, and a growing range of $\log\mathrm{gap}$ forces a growing standard deviation. □

The frozen local regime and the conventions (P1)–(P6) are fixed in Section D.1; this appendix uses them without change. It proves the gap–distance sharpness of Lemma 30, whose two sides rest on a Lipschitz upper bound and a first-order analysis of the gap map, and records two remarks on the per-player bound and on the finite-time behavior of the fitted slope.

D.2 A Lipschitz bound for the gap

Lemma 27 (Gap Lipschitz bound)**.** *There is $L_{\mathrm{gap}} < \infty$ such that* gap *is L_{gap}-Lipschitz on the regime, and $\mathrm{gap}(z_f) = 0$ for every $z_f \in E_+$ in the regime.*

Proof. On the regime the strategy map $z \mapsto (x(z), y(z))$ is a composition of the frozen-pattern algorithm maps (P4): an affine prediction step, the analytic solution $\gamma(\cdot)$ of the norm constraint, the clipping with fixed pattern (affine), and the normalizations $r \mapsto r/\|r\|_1$. The normalization denominators are bounded below on the regime: writing $\tilde r_1 = a\xi + w_1$ with $w_1 \perp \xi$ ((P2)), $\|\tilde r_1\|_1 \ge \|\tilde r_1\|_2 \ge a\|\xi\|_2 \ge \frac{c_1}{2}\|\xi\|_2 \ge c_1/(2\sqrt{m})$ by (P3), and the prediction preserves the ℓ_2

norm, so $\|p_1\|_1 \geq \|p_1\|_2 = \|\tilde{r}_1\|_2$ obeys the same bound; symmetrically for player 2 with ψ. The elementary estimate $\|u/\|u\|_1 - v/\|v\|_1\|_1 \leq 2\|u-v\|_1/\max(\|u\|_1, \|v\|_1)$ makes each normalization Lipschitz there. Hence $z \mapsto (x(z), y(z))$ is Lipschitz on the compact closure of the regime. The gap $\mathrm{gap}(z) = \max_i (Ay(z))_i - \min_j (x(z)^\top A)_j$ is a difference of a max and a min of finitely many affine functions of (x, y), each $\|A\|_{1\to\infty}$-Lipschitz; composing yields L_{gap}, depending only on $\|A\|$, c_1, c_2, and the dimensions. Every $z_f \in E_+$ in the regime is a fixed point whose strategies are exactly (x^*, y^*) (Proposition 1), so $\mathrm{gap}(z_f) = 0$. □

D.3 Proof of Lemma 30

The upper bound is immediate from Lemma 27 and (28): $\mathrm{gap}(z) = \mathrm{gap}(z) - \mathrm{gap}(P_E z) \leq L_{\mathrm{gap}}\|z - P_E z\| = L_{\mathrm{gap}}\|w\|$. The remainder of this subsection proves the lower bound. We need two sublemmas.

Lemma 28 (Max lemma)**.** *Let $\xi \in \operatorname{int}\Delta_m$ and $u \in \mathbb{R}^m$ with $\xi^\top u = 0$. Then $\max_i u_i \geq \xi_{\min}\|u\|_\infty \geq (\xi_{\min}/\sqrt{m})\|u\|_2$.*

Proof. Let k attain $\|u\|_\infty = |u_k|$. If $u_k > 0$ then $\max_i u_i \geq \|u\|_\infty \geq \xi_{\min}\|u\|_\infty$. If $u_k < 0$, then from $0 = \sum_i \xi_i u_i \leq \xi_k u_k + (1-\xi_k)\max_i u_i$ we get $\max_i u_i \geq \xi_k\|u\|_\infty/(1-\xi_k) \geq \xi_{\min}\|u\|_\infty$. The ℓ_2 bound is $\|u\|_\infty \geq \|u\|_2/\sqrt{m}$. □

Lemma 29 (First-order gap map and its smallest singular value)**.** *Define $T_{a,b} : W \to \mathbb{R}^{2m}$,*

$$T_{a,b}\, w := (p, q), \qquad p := \tfrac{1}{b}\, B\big(w_2 - \tfrac{1}{a} B^\top w_1\big), \qquad q := \tfrac{1}{a}\, B^\top\big(w_1 + \tfrac{1}{b} B w_2\big).$$

Then $T_{a,b}$ is block-diagonal across the orthogonal planes $W_k = \operatorname{span}\{(u_k, 0), (0, v_k)\}$: in the coordinates $w_1 = \alpha u_k$, $w_2 = \beta v_k$ one has $(p, q) = \sigma_k\big((M_k(\alpha,\beta)^\top)_1 u_k,\ (M_k(\alpha,\beta)^\top)_2 v_k\big)$ with

$$M_k = \begin{pmatrix} -\sigma_k/(ab) & 1/b \\ 1/a & \sigma_k/(ab) \end{pmatrix}, \qquad \det M_k = -\eta^2\big(1 + \eta^2\sigma_k^2\big), \quad \eta = \eta(a, b),$$

and consequently

$$\sigma_{\min}(T_{a,b}) \;=\; \min_k\, \sigma_k\, \sigma_{\min}(M_k) \;\geq\; \min_k \frac{\sigma_k\,|\det M_k|}{\|M_k\|_F} \\ \geq\; \frac{\sigma^+_{\min}\,\eta^2}{\sqrt{2\sigma^2_{\max}\eta^4 + 1/a^2 + 1/b^2}} \;=:\; c_T(a, b) \;>\; 0.$$

Proof. Block structure: by (P1), $Bv_k = \sigma_k u_k$ and $B^\top u_k = \sigma_k v_k$, so for $w_1 = \alpha u_k$, $w_2 = \beta v_k$,

$$p = \tfrac{1}{b}\big(\beta\sigma_k u_k - \tfrac{\alpha}{a}\sigma_k^2 u_k\big) = \sigma_k\Big(-\tfrac{\sigma_k}{ab}\alpha + \tfrac{1}{b}\beta\Big)u_k, \qquad q = \sigma_k\Big(\tfrac{1}{a}\alpha + \tfrac{\sigma_k}{ab}\beta\Big)v_k,$$

which is the stated 2×2 action; the images lie along the orthonormal directions u_k, v_k, so the blocks are orthogonal to one another. The determinant is $-\sigma_k^2/(ab)^2 - 1/(ab) = -\eta^2(1+\eta^2\sigma_k^2)$. For a 2×2 matrix, $\sigma_{\min} = |\det|/\sigma_{\max} \geq |\det|/\|\cdot\|_F$; using $|\det M_k| \geq \eta^2$ and $\|M_k\|_F^2 = 2\sigma_k^2\eta^4 + 1/a^2 + 1/b^2 \leq 2\sigma^2_{\max}\eta^4 + 1/a^2 + 1/b^2$ gives the bound, and $\min_k \sigma_k = \sigma^+_{\min}$. □

On the regime, using $a \geq c_1/2$, $b \geq c_2/2$ and $\eta \in [\eta_-, \eta_+]$ from (P3),

$$c_T := \inf_{(a,b)\in\text{regime}} c_T(a, b) \;\geq\; \frac{\sigma^+_{\min}\,\eta_-^2}{\sqrt{2\sigma^2_{\max}\eta_+^4 + 4/c_1^2 + 4/c_2^2}}, \tag{29}$$

which is the explicit constant announced in Lemma 30.

Lemma 30 (Gap–distance sharpness)**.** *There exist $0 < c_- \leq L_{\mathrm{gap}} < \infty$ and a threshold $\delta_2 > 0$ such that, in the frozen regime with $\|w\| \leq \delta_2$,*

$$c_-\,\|w\| \;\leq\; \mathrm{gap}(z) \;\leq\; L_{\mathrm{gap}}\,\|w\|,$$

where L_{gap} is a Lipschitz constant of gap *in state coordinates and $c_- > 0$ is explicit (built from the smallest on-support equilibrium probabilities and the spectrum of B); the lower bound is where nondegeneracy enters.*

Proof of Lemma 30 (lower bound). Let $z = z_f + w$ be in the regime, $z_f = (a\xi, b\psi)$, and let $(x, y) = (x(z), y(z))$ be the main-iterate strategies; by (P4) their supports are contained in S_1, S_2. We proceed in four steps.

Step 1: exact localization to the support. By strict complementarity and continuity (shrinking the regime thresholds if necessary, absorbed into δ_2), every off-support row satisfies $(Ay)_i \le v^* - \delta/2$ for $i \notin S_1$, while on the support $\max_{i\in S_1}(\tilde{A}y)_i \ge \xi^\top \tilde{A} y = v^* \mathbf{1}^\top y = v^*$, using $\xi^\top \tilde{A} = v^* \mathbf{1}^\top$ (Lemma 10) and that y is a probability vector on S_2. Hence the global max is attained on S_1; symmetrically the global min is attained on S_2. Using $\tilde{A}\psi = v^*\mathbf{1}$ and $\tilde{A}^\top \xi = v^* \mathbf{1}$,

$$\mathrm{gap}(z) = \underbrace{\max_{i\in S_1} \big(\tilde{A}(y-\psi)\big)_i}_{=:G_1 \ge 0} + \underbrace{\Big(-\min_{j\in S_2} \big(\tilde{A}^\top (x-\xi)\big)_j\Big)}_{=:G_2 \ge 0} \qquad \text{exactly.}$$

Step 2: sharp lower bound on the max via equilibrium weighting. Let $u := \tilde{A}(y-\psi)$ and $v := \tilde{A}^\top(x-\xi)$. Then $\xi^\top u = v^* \mathbf{1}^\top (y-\psi) = 0$ and $\psi^\top v = 0$ *exactly*. Lemma 28 (and its mirror for the min) gives

$$\mathrm{gap}(z) \;\ge\; \frac{\min(\xi_{\min}, \psi_{\min})}{\sqrt{m}} \big(\|u\| + \|v\|\big) \;\ge\; \frac{\min(\xi_{\min}, \psi_{\min})}{\sqrt{m}} \, \|(u, v)\|.$$

Step 3: first-order nondegeneracy of (u, v) in w. The map $z \mapsto (u, v)$ is analytic on the regime and vanishes identically on E_+ (there $(x, y) = (\xi, \psi)$). Its differential at z_f in the direction w is computed from the cancellation identities of Lemmas 9 and 10 ($d\gamma = 0$ at fixed points; $\tilde{A} P_\psi = B$, $\tilde{A}^\top P_\xi = B^\top$): writing the differentials of the normalized iterates as in the proof of Lemma 9,

$$u = T^{(1)}_{a,b} w + O\big(C_q \|w\|^2\big), \qquad v = T^{(2)}_{a,b} w + O\big(C_q \|w\|^2\big),$$

where $(T^{(1)}_{a,b}, T^{(2)}_{a,b}) w = (p, q) = T_{a,b} w$ is exactly the map of Lemma 29 and C_q is the uniform second-order constant of (P5). Hence

$$\|(u, v)\| \;\ge\; \|T_{a,b} w\| - \sqrt{2}\, C_q \|w\|^2 \;\ge\; c_T \|w\| - \sqrt{2}\, C_q \|w\|^2,$$

by Lemma 29 and (29).

Step 4: composition. Combining Steps 2 and 3,

$$\mathrm{gap}(z) \;\ge\; \frac{\min(\xi_{\min}, \psi_{\min})}{\sqrt{m}} \Big(c_T \|w\| - \sqrt{2}\, C_q \|w\|^2\Big) \;\ge\; \frac{\min(\xi_{\min}, \psi_{\min})}{2\sqrt{m}}\, c_T \, \|w\| \;=\; c_- \|w\|$$

whenever $\|w\| \le \delta_2 := \min\big\{\delta_0,\ c_T/(2\sqrt{2}\, C_q)\big\}$. □

Remark 17 (Where the lower bound fails: degenerate equilibria). All the nontrivial content of the lower bound is the injectivity of $T_{a,b}$ on W (Lemma 29), which rests on the kernel condition of Assumption 1. If the equilibrium is not unique, every strictly complementary equilibrium is degenerate (Lemma 11): for one player, say the column player, there is a direction d with $\tilde{A} d = 0$, $\mathbf{1}^\top d = 0$ beyond $\mathrm{span}\{\psi\}$; moving along it leaves $u = \tilde{A}(y - \psi)$ identically zero while $\mathrm{dist}(z, E)$ grows linearly, so $\mathrm{gap} \gtrsim \mathrm{dist}(z, E)$ *fails*. The correct comparison object is then $\mathrm{dist}(z, \mathrm{Fix})$ with respect to the full (higher-dimensional) fixed-point set, whose normal-bundle analysis we leave open. The failure is localized to the $E \oplus W$ decomposition: E is two-dimensional and c_- is built from $\sigma^+_{\min}$ over W. B keeps a smallest positive singular value whenever $\mathrm{rank}\, B > 0$; the obstruction is that the kernel of B exceeds $\mathrm{span}\{\psi\}$, so W as defined in (P2) no longer spans the complement of E. Viewed structurally, gap is a max minus a min of finitely many affine functions, so its growth order on its zero set is a polyhedral sharpness phenomenon: uniqueness of the equilibrium is exactly linear growth, and Lemma 30 quantifies the sharpness constant through $(\xi_{\min}, \psi_{\min}, \sigma^+_{\min}, \eta)$. As $\sigma^+_{\min} \to 0$ (near-degenerate game) or $\xi_{\min} \to 0$ (near-boundary support, near-failure of strict complementarity) the constant c_- degrades continuously, consistent with the slow large-support instances observed in Section 5.

Remark 18 (Measuring $\dot{N}_{\mathrm{tot}}$). Differencing recorded N^2 fails in double precision deep in the linear regime, where rounding swamps the increment. The right-hand side of the increment identity is a sum

of small positive terms from $g^{(t)}$; the experiments accumulate it in a separate register $G_S(t)$, which stays small ($O(N^2)$, against the $O(10)$ of the full-energy cumulant), and difference G_S between recorded times under a representability guard (Appendix E.3). The guard is a diagnostic on the differenced increment, not an error bound on the accumulated sum.

Remark 19 (Both players by design). The *per-player* analogue of Lemma 26 is false, because the rotation carries $w^{(t)}$ through $\{w_1 = 0\}$ (Remark 20, which also records a windowed repair). The certificate therefore uses the *aggregate* $\dot{N}_{\mathrm{tot}}$, for which Lemma 26 holds at every step, with no windowing or genericity condition.

D.4 Why the certificate aggregates both players

Remark 20 (Failure and partial repair of the per-player bound). The per-player lower bound $\dot{N}_1(t) \gtrsim \|w^{(t)}\|^2$ is *false* pointwise in t: by Lemma 25, $\dot{N}_1(t) = \eta^4 \|BB^\top w_1^{(t)}\|^2 + O(\|w^{(t)}\|^3)$ vanishes to first order whenever $w_1^{(t)} = 0$, i.e. whenever the local rotation carries the transverse state through the opponent's axis of the modal plane W_k (the invariant plane $W_k = \mathrm{span}\{(u_k, 0), (0, v_k)\}$ of the kth singular pair, Lemma 29). Two partial repairs are available. (i) *Two-step window, single mode (rigorous).* Suppose $w^{(t)} \in W_k$ for a single k. The linearized dynamics on W_k is conjugate to $|\mu_k| R(\varphi_k)$, a rotation–dilation with $\varphi_k = \arg \mu_k \in (0, \pi)$ and $\sin \varphi_k \neq 0$ (Theorem 4). For any nonzero linear functional ℓ on the plane and rotation R by angle φ with $\sin \varphi \neq 0$, the 2×2 matrix with rows ℓ and $\ell \circ R$ has determinant $-\sin \varphi \, \|\ell\|^2 \neq 0$, whence $\max\big(|\ell(\zeta)|, |\ell(R\zeta)|\big) \geq c(\ell, \varphi)\|\zeta\|$ for all ζ. Applying this with ℓ the w_1-coordinate functional gives $\max\big(\dot{N}_1(t), \dot{N}_1(t{+}1)\big) \geq c_k \|w^{(t)}\|^2 - O(\|w^{(t)}\|^3)$. This covers the case in which the trajectory sits in one modal plane; the three-step bound (ii) below covers the general case. (ii) *Three-step window, linearized dynamics.* The modal planes W_k are invariant under the frozen-face *Jacobian* (Lemma 13), not under the nonlinear map itself, so the statement below is about the linearization $w \mapsto J_S w$; the nonlinear map adds the $O(\|w\|^2)$ remainder of Lemma 15, which the cubic error term absorbs over a window of fixed length. Writing $w_1^{(t)} = \sum_k \alpha_k(t) u_k$, $w_2^{(t)} = \sum_k \beta_k(t) v_k$, the linearized dynamics decouple per mode as $\big(\alpha_k(t{+}1), \beta_k(t{+}1)\big)^\top = J_k \big(\alpha_k(t), \beta_k(t)\big)^\top$ with $J_k = \left(\begin{smallmatrix} 1-q_k & \sigma_k/b \\ -\sigma_k/a & 1-q_k \end{smallmatrix}\right)$, $q_k := \eta^2 \sigma_k^2$ (Theorem 4). Three consecutive w_1-observations of one mode read $M_k \big(\alpha_k(t), \beta_k(t)\big)^\top$ with

$$M_k := \begin{pmatrix} 1 & 0 \\ 1 - q_k & \sigma_k/b \\ (1-q_k)^2 - q_k & 2\sigma_k(1-q_k)/b \end{pmatrix}, \qquad \det(M_k^\top M_k) = \frac{\sigma_k^2}{b^2}\big(|\mu_k|^4 + 4(1-q_k)^2 + 1\big) > 0.$$

The two columns are never collinear (their first entries are 1 and 0, and $\sigma_k/b \neq 0$), so each mode is observable from three consecutive w_1 observations with *no condition on* q_k; with $\lambda_k := \lambda_{\min}(M_k^\top M_k) > 0$, $\sum_{s=0}^2 \alpha_k(t{+}s)^2 \geq \lambda_k(\alpha_k(t)^2 + \beta_k(t)^2)$. By Lemma 25, $\dot{N}_1(t) = \eta^4 \sum_k \sigma_k^4 \alpha_k(t)^2 + O(\|w^{(t)}\|^3)$, and $\sum_k (\alpha_k^2 + \beta_k^2) = \|w\|^2$ (orthonormal singular vectors); summing over $s = 0, 1, 2$ and absorbing the cubic error as in Lemma 26 gives, below a threshold,

$$\sum_{s=0}^{2} \dot{N}_1(t{+}s) \;\geq\; \tfrac{1}{2}\, \eta_-^4 \, c_O \, \|w^{(t)}\|^2, \qquad c_O := \inf \min_k \sigma_k^4 \lambda_k > 0,$$

with the infimum over the compact regime in (η, a, b); positivity follows from the displayed determinant and compactness. Three steps suffice, independent of $\dim W$, because each block is a rotation–dilation, and with the infimum the bound is uniform in t. Repeated singular values need no separate treatment: the singular vectors can still be chosen orthonormal, the blocks J_k coincide for equal σ_k, and the modal energies α_k^2 add across the merged block, so the same determinant computation applies. The certificate itself needs neither repair: it uses $\dot{N}_{\mathrm{tot}}$, for which Lemma 26 holds at every step.

D.5 Finite-time correction to the fitted slope

Remark 21 (Open: finite-time correction to θ). Corollary 2 controls the fitted slope through the *width* of the offset band only, and the offset $C(t)$ need not converge. Each transverse mode is a rotation–dilation, so $\dot{N}_{\mathrm{tot}}(t)$ is proportional to the squared modal radius while $\mathrm{gap}(t)$ carries a

phase-dependent factor (the maximum of finitely many linear forms on the rotating state). Their log-ratio keeps oscillating with the phase as the radius contracts. A sharper finite-time statement would therefore expand $C(t) = C_\infty(\varphi_t) + O(\|w^{(t)}\|)$ with an explicit phase-dependent leading term and an explicit constant in the remainder; the latter is inherited from the non-closed-form second-order constants of (P5). Likewise, the practical choice of the measurement window (discarding the pre-regime segment) requires an operational regime-detection criterion. Candidate rules combine relative stagnation of the norm increments with detection of support freezing, and we treat them as part of the experimental protocol (Section 5); the theorem does not supply one. Both items are open.

E Experimental Details

This appendix documents the protocol of Section 5: the instance family, the algorithm implementation, every measurement window and numerical guard, and the excluded instances accounted for one by one. All experiments are pure NumPy/SciPy, single-threaded per instance, with fixed seeds throughout; no GPU is required.

E.1 Instance family

Each instance is a square zero-sum matrix game $A \in \mathbb{R}^{n\times n}$ (row player maximizes) drawn from a grid of

$$\begin{aligned} & n \in \{5, 20, 50\} \\ & \times \text{ distribution} \in \{\textsf{gaussian}, \textsf{uniform}, \textsf{pm1_sparse}\} \\ & \times \text{ generator} \in \{\textsf{plain}, \textsf{antisym}\} \end{aligned}$$

with 12 seeds per cell, i.e. $3 \times 3 \times 2 \times 12 = 216$ instances. The distributions are: i.i.d. standard Gaussian entries; i.i.d. uniform entries on $[-1, 1]$; and sparse-sign entries, where each entry is ± 1 with equal probability and is then kept with probability 0.3 (zero otherwise). The plain generator uses the raw draw G; the antisym generator uses the antisymmetrization $(G - G^\top)/\sqrt{2}$, which produces symmetric (value-zero) games with qualitatively larger equilibrium supports. Seeds are deterministic ($10\,000 + 997\,u$ for instance index u).

For every game we solve both players' linear programs (LP; SciPy/HiGHS), obtaining an equilibrium (x^*, y^*), the value v^*, the supports (threshold 10^{-8}; HiGHS default primal and dual feasibility tolerances 10^{-7}), the strict-complementarity margin, and a duality-gap sanity check, where the margin is

$$\delta := \min\big(\min_{i\notin S_1}(v^* - (Ay^*)_i),\ \min_{j\notin S_2}((x^{*\top}A)_j - v^*)\big).$$

The LP equilibrium is used only for reference statistics (support-match rates, $|\gamma - v^*|$); all spectral quantities in the protocols for verifications (3) and (5) of this appendix are rebuilt from the *actually converged* state, because on $34/216$ instances (15.7%) the realized exact-zero pattern differs from the thresholded LP support; multiple equilibria are one possible explanation, not an established classification.

E.2 Algorithm implementation

We implement Algorithm 2 of Zhang et al. (2025a) (IREG-PRM$^+$) in the extra-gradient setup, with the prediction $m^{(t)}$ given by the payoff against the opponent's pre-iterate. Two lines deserve emphasis because a plausible-looking variant breaks the theory's sharpest prediction:

- *Prediction step (Algorithm 2, line 9).* $r^{(t)} = \tilde{r}^{(t)} + m^{(t)} - \gamma^{(t)}\mathbf{1}$ is *unclipped*; the strategy is $x^{(t)} = \big[r^{(t)}\big]_+ / \|\big[r^{(t)}\big]_+\|_1$. The shift $\gamma^{(t)}$ solves $\|\big[\tilde{r}^{(t)} + m^{(t)} - \gamma\mathbf{1}\big]_+\|_2 = N(t)$ by the sorted piecewise closed form followed by a Newton refinement on $f(\gamma) = \|[v-\gamma]_+\|_2^2 - N(t)^2$, which avoids the catastrophic cancellation of the pure closed form when $N(t) \ll \|v\|$.
- *Update step (Algorithm 2, line 13).* $\tilde{r}^{(t+1)} = \big[r^{(t)} + g^{(t)}\big]_+$ with the *unclipped* $r^{(t)}$, where $g^{(t)}$ is the instantaneous regret of the correction step.

Keeping the prediction step unclipped is what makes the analysis sharp. Clipping it as well ($\tilde{r}^{(t+1)} = \left[\left[r^{(t)}\right]_+ + g^{(t)}\right]_+$) still gives a valid regret minimizer (it retains norm monotonicity and scale invariance), but near a strict-complementarity fixed point it emits $O(\|g\|)$ off-support "dust," so its one-step map is not smooth at the fixed point and exact support freezing fails. All results in this paper use the strict Algorithm 2; with it, the exact-zero pattern freezes in finite time as predicted by Lemma 6, and this freeze is used as a regression check of the implementation.

Horizons are $T = 2 \times 10^5$ for $n \in \{5, 20\}$ and $T = 10^6$ for $n = 50$. Time series (gap, norms, cumulative $\sum\|g\|^2$, cumulative truncation loss, $|\gamma - v^*|$) are recorded every $\max(10, T/4000)$ iterations; the gap is additionally checked every 10 iterations for the snapshot trigger below.

E.3 Measurement protocols

Support freezing (verification 1). We track two patterns after every iteration: the relative pattern $\{i : \tilde{r}_i > 10^{-6}\, N(t)\}$ and the exact-zero pattern $\{i : \tilde{r}_i > 0\}$ (both players jointly). t_{freeze} is the iteration of the *last* pattern change over the whole run. The headline numbers (99.5% frozen before $T/2$; median $t_{\text{freeze}}/T = 2.7 \times 10^{-4}$, p90 1.5×10^{-3}, max 0.52) refer to the exact-zero pattern, which is the hard prediction of Lemma 6; the relative pattern gives identical medians.

Snapshot. A state snapshot $z = (\tilde{r}_1, \tilde{r}_2)$ is stored at the first $t \geq 200$ (checked every 10 iterations) at which the gap falls below 10^{-4} of its value at $t = 1$ *and is still strictly positive*. On this suite the second condition is the only one that ever blocks the trigger: exactly 32 instances have no snapshot, because they jump past the threshold to a floating-point zero gap before a positive-gap check lands, and report a gap of exactly 0.0 from then on (29 at $n = 5$, 3 at $n = 20$). Those 32 are excluded from the spectral and θ measurements (counted below), since no positive-gap segment remains to measure. Every instance with a snapshot (184 of 216) also converged (final gap below 10^{-10} of its initial value), so all 216 instances converge and the *measurable* pool, on which the spectral, θ and step-size statistics are computed, is the 184 with a snapshot.

The 32 finite-step convergers. A gap of exactly 0.0 in double precision could be rounding. Two checks say it is not. First, 31 of the 32 games have a pure saddle point (an entry that is both a column maximum and a row minimum), the $m = 1$ case in which Section 4.2 predicts convergence in one round after absorption; the remaining game is a ± 1 sparse instance of value 0 with a non-unique equilibrium and no pure saddle. Second, re-running all 32 in 60-digit arithmetic (bisection for the γ-shift, 200 rounds) gives a gap that is *exactly* zero, at the same round as the double-precision run in 30 instances and one round later in the other two; the latest such round is 29, for the game without a pure saddle. Finite-step convergence on these instances is a property of the iteration, and the excluded group is exactly the set on which no exponential segment exists.

Finite-difference (FD) Jacobian (verification 2). The one-step map Φ acts on the state $z = (\tilde{r}_1, \tilde{r}_2)$; its Jacobian at the snapshot is formed column by column with step $\varepsilon = 10^{-7} \max_j z_j$: central differences for coordinates with $z_j > 10\varepsilon$, forward differences otherwise (negative $\tilde{r}$ is not a reachable state). A consistency probe recomputes up to 6 columns with step 10ε and reports the median relative column difference; across instances this consistency indicator has median 6.0×10^{-7} (p90 1.7×10^{-6}), which is an accuracy diagnostic for the FD probe and is exactly the scale of the maximum deviation seen in verification (3). Eigenvalues with $|\lambda| > 1 - 10^{-5}$ are classified as the *unit band*, the numerical stand-in for the two exact unit eigenvalues of Theorem 4(i) (median count 2, min 2, max 10), and $\rho_2(J)$ is the largest modulus below the band, the numerical stand-in for ρ^*. A slowly contracting mode with $\eta\sigma < 3 \times 10^{-3}$ would also fall in the band, so a band count above 2 suggests a degenerate equilibrium without proving one; the 8 instances so flagged are confirmed or refuted structurally by the kernel dimension of B on the converged support (verification 6). The predicted per-iteration slope is $\log \rho_2(J)$; this verification tests the mechanism (rate = frozen-face spectrum) with a numerical Jacobian, and verification 3 tests the closed form against the same Jacobian.

Measured rate (verification 2, continued). The measured slope is an ordinary least-squares fit of $\log \text{gap}$ against t on the exponential segment. The recorded points are those with gap in $[10^3 \cdot \text{floor},\ 0.5 \cdot \text{gap}_{\text{snap}}]$, where floor is the smallest positive gap attained in the run (if the run never reaches its floor, the lower cut is $1.5 \times$ floor). An instance is *usable* if it has a snapshot, the fit uses at

least 5 points, $R^2 \geq 0.9$, and the slope is negative. Accounting over all 216 instances: 155 usable, 32 without a snapshot (the finite-step convergers of the snapshot paragraph, all $n \in \{5, 20\}$ with final gap exactly 0.0), 29 with no fittable segment (runs that hit the double-precision floor almost immediately after the snapshot). Both excluded groups are visible in the raw traces: the no-snapshot group reaches a zero gap, in double and in 60-digit arithmetic alike, before the snapshot trigger fires. Each of the 29 non-fittable instances has fewer than 5 usable points below half the snapshot gap (one has fewer than 10 positive gaps overall), so the R^2 filter never triggers. The $216 = 155 + 32 + 29$ accounting is produced by one pass over the traces and reproduced independently by a second implementation of the same filter; the two agree on all three counts. The freeze-time split of the 155 reported in Section 5 comes from the same accounting.

Closed form versus FD (verification 3). From the eligible pool (the 184 instances with a snapshot, all of which have final gap below 10^{-10} of its initial value) we sample 32 instances uniformly without replacement at a fixed, recorded seed, so the draw is deterministic and auditable. Each is re-run for $4\,t_{\mathrm{snap}} + 2000$ iterations to a fully converged state z. The support blocks are taken as the exact-zero pattern of $\tilde{r}$, and the normalized strategies as $\xi = \tilde{r}_1|_{S_1}/c_1$, $\psi = \tilde{r}_2|_{S_2}/c_2$ with $c_i = \|\tilde{r}_i\|_1$. The closed-form spectrum is $\mu = 1 + \lambda + \lambda^2$, $\lambda = \pm i\sqrt{q}$, over the nonzero eigenvalues q of $Q = \tilde{A} P_\psi \, \tilde{A}^\top P_\xi / (c_1 c_2)$ (kernel modes, $|q| \leq 10^{-9}$ relative, are the unit-band directions and are excluded from the rate). The comparison FD Jacobian is evaluated at the *same* state z. Relative errors on ρ_2: median 2.4×10^{-10}, maximum 4.1×10^{-7}, i.e. at the FD consistency limit quoted above, in all $32/32$ instances, including one degenerate instance (5×8 support blocks, 9 unit eigenvalues). Building B from the thresholded LP support gives relative errors up to 2.66%, with the discrepancy concentrated on instances whose realized support differs from the thresholded LP support (the 84.3% support-match rate of Section E.1). Rebuilding it from the converged support removes this discrepancy.

Certificate exponent θ (verification 4). The window is the linear regime $t \geq t_{\mathrm{snap}}$, clipped to start at $t \geq \max\{t_{\mathrm{snap}}, t_{\mathrm{freeze}}\}$, where t_{freeze} is the time of the last exact-zero support-pattern change. The window is therefore a *candidate* frozen window: the exact-zero pattern of $\tilde{r}$ is fixed on it, but the full clipping pattern behind Lemma 24 (positivity of r and $r + g$ on the support at both the prediction and the update stage) was not separately checked, so the on-support energy is an empirical proxy for $\dot{N}_{\mathrm{tot}}$ on this window. The primary estimator $\hat{\theta}_{\mathrm{supp}}$ is built from the theorem's quantity: at each step we add the *on-support* update energy $\sum_i \sum_{j \in S_i} (g_{i,j}^{(t)})^2$ into a running cumulative $G_S(t)$ stored at the recorded times; once the pattern is frozen this equals $\dot{N}_{\mathrm{tot}}(t)$ by Lemma 24. Between consecutive recorded times $t_k < t_{k+1}$ we form the interval average $\overline{\dot{N}}_{\mathrm{tot}}^{[k]} := (G_S(t_{k+1}) - G_S(t_k))/(t_{k+1} - t_k)$. What is measured is thus a per-unit-time average of the on-support energy over the interval; it equals the average of the theorem's pointwise $\dot{N}_{\mathrm{tot}}$ only if the full clipping pattern is frozen on the window. The pointwise sandwich of Theorem 5 transfers to such averages with the same κ when the interval gap is read as the root-mean-square gap over the interval (average Corollary 2's two-sided bound over $t \in [t_k, t_{k+1})$). We pair $\overline{\dot{N}}_{\mathrm{tot}}^{[k]}$ with the geometric mean $\sqrt{\mathrm{gap}(t_k)\mathrm{gap}(t_{k+1})}$ of the endpoint gaps as an empirical proxy for that RMS; for a geometrically decaying gap with oscillation the two differ by a bounded factor that the theory does not pin down. The interval statistics below therefore test the ratio law empirically; the pointwise inequality itself is not verified. The cumulative G_S is small ($O(N^2)$, against the $O(10)$ of the full-energy cumulant), which is what keeps its differences representable deep in the linear regime (Remark 18); a difference is kept only if it passes the representability guard ($> 64\,\varepsilon_{\mathrm{mach}}$ relative to $G_S(t_{k+1})$) and the interval's geometric-mean gap exceeds the absolute floor 10^{-13} (payoffs are $O(1)$). We require at least 6 surviving increments and fit the OLS slope of $\log \overline{\dot{N}}_{\mathrm{tot}}^{[k]}$ against the log geometric-mean gap. The cross-check estimator $\hat{\theta}_{xy}$ applies the same protocol to the *full* update energy ΔG^2 of the cumulative $\sum_t \|g^{(t)}\|^2$ (which adds the off-support energy), and $\hat{\theta}_{\mathrm{ratio}}$ is the ratio of the OLS time-slopes of $\log(\Delta G^2/\Delta t)$ and of $\log \mathrm{gap}$, which averages the complex-spectrum oscillation out of numerator and denominator before dividing. This yields $129/216$ usable instances under all three estimators; the exclusions are the 32 no-snapshot instances plus instances with fewer than 6 guard-surviving increments (fast floor-hitters). Usable windows span a median of 2.8 decades of gap decay (p10 2.4). Results: $\hat{\theta}_{\mathrm{supp}}$ median 1.986 (p10 1.911, p90 2.038; 96.1% between 1.8 and 2.2), $\hat{\theta}_{xy}$ median 1.995 (98.4%), and $\hat{\theta}_{\mathrm{ratio}}$ median 2.004 (98.4%); the three estimators agree with each other and with the

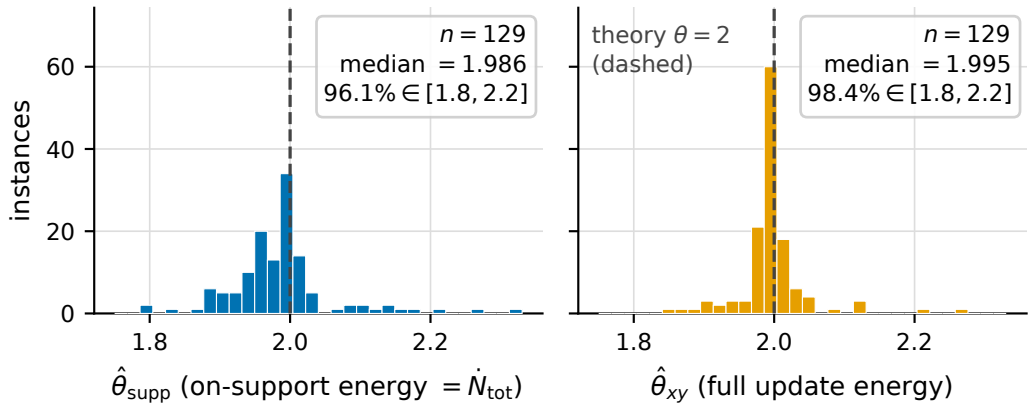


Figure 2: Certificate exponent θ on 129 usable instances: on-support energy $\hat{\theta}_{\mathrm{supp}}$ and full-update-energy cross-check $\hat{\theta}_{xy}$, both concentrated on the theory line $\theta = 2$ (dashed), medians 1.986 and 1.995, 96.1% and 98.4% between 1.8 and 2.2.

predicted 2 to better than 1%. A fourth, deprecated estimator differences the recorded cumulants (cumulative $\sum\|g\|^2$ minus cumulative truncation loss); it gives a consistent but degraded median of 1.948 with only 92/128 (71.9%) between 1.8 and 2.2; one instance lacks enough guarded increments for this estimator. Its failure mode is instructive and is *not* statistical noise: once the cumulants reach order 10^1, their accumulated floating-point rounding ($\sim 10^{-13}$–10^{-12} absolute) exceeds the true late-regime increments, producing a spurious constant plateau. The instances it loses are recovered exactly by the on-support estimator, whose cumulant stays small (Remark 18).

The naive measurement that the protocol above replaces, differencing N^2 across log-spaced checkpoints over the whole run, yields a median exponent of 0.375 on the same data. Its two failure modes are instructive: (i) it mixes the transient and linear regimes in one fit, and (ii) once gap $\sim 10^{-12}$ the true increment of N^2 is below the representable resolution of double precision, so the computed differences are dominated by rounding noise. Both are measurement artifacts; the guarded protocol above eliminates them without touching the dynamics.

Self-stabilization (verification 5). This verification runs on the *entire* measurable pool, all 184 instances with a snapshot, with no sampling. Each instance is re-run to $4\,t_{\mathrm{snap}} + 2000$ iterations, B is rebuilt from the converged support and strategies as above, and $\eta_\infty \sigma_{\max}(B)$ is evaluated with $\eta_\infty = 1/\sqrt{c_1 c_2}$ at the realized saturation levels. Reported statistics: 0 of 184 instances violate $\eta_\infty \sigma_{\max}(B) \le 1$; the distribution has median 0.790, p90 0.946, and maximum 0.9985, with 18 instances above 0.95; the minimum distance of $q_{\max} = (\eta_\infty \sigma_{\max}(B))^2$ to the boundary 1 is 0.0029. So the constraint holds with room to spare: the median instance sits 0.21 below the boundary, and the 18 instances above 0.95 form the tail on which the heuristic is really tested. What this tests is the self-stabilization heuristic (Remark 1). It does not test Proposition 2: the proposition permits convergence to a violating ray along its center-stable set. A converged instance above 1 would therefore refute the heuristic and exhibit such an exceptional orbit, without contradicting the theorem.

Ratio certificate (verification 6). On the 129 θ-usable instances we re-run each trajectory to $\min(T,\, 4\,t_{\mathrm{snap}} + 2000)$ iterations and work inside the same linear-regime window as the θ measurement. The empirical deviation is

$$\kappa_{\mathrm{emp}} \;=\; \max_{t_1 < t_2 \in \mathrm{window}} \max\!\left\{ R(t_1,t_2),\, R(t_1,t_2)^{-1} \right\}, \qquad R(t_1,t_2) = \frac{\mathrm{gap}(t_2)/\mathrm{gap}(t_1)}{\sqrt{\dot{N}_{\mathrm{tot}}(t_2)/\dot{N}_{\mathrm{tot}}(t_1)}},$$

where t_1, t_2 range over recorded times and $\dot{N}_{\mathrm{tot}}$ and gap stand for the interval quantities of the θ protocol ($\dot{N}_{\mathrm{tot}}$ and the geometric-mean gap), computed under three numerators. The primary numerator is the on-support energy of Lemma 24 accumulated between recorded times (Remark 18; 124 instances qualify). The full-update-energy increments $\Delta \sum_t \|g^{(t)}\|^2$ serve as a cross-check

(124 instances). The third numerator is the telescoped increments $N(t+1)^2 - N(t)^2$ with the same representability guard as the θ protocol (at least 6 guarded increments required; $123/124$ instances qualify), which is the naive monitor one would track by differencing the recorded norm. The theoretical constant is assembled per instance from the definitions in Theorem 5: L_{gap} and c_- from the converged support, strategies, and value (rebuilt as in verification 3), and $\sqrt{n_+/n_-}$ from the extreme nonzero singular values of B and the realized range of η over the window. The formulas are those of the proofs, evaluated on the window's realized range of the scales $a_t = N_1(t)/\|\xi\|_2$, $b_t = N_2(t)/\|\psi\|_2$ (with $\eta_\pm$ the extreme values of $1/\sqrt{a_t b_t}$; these equal the projection coordinates a, b of the decomposition $\tilde{r}_1 = a\xi + w_1$ only on the cone, and differ from them by $O(\|w\|^2)$ inside the regime): $n_- = \frac{1}{2}\eta_-^4(\sigma_{\min}^+)^4$, $n_+ = 2\eta_+^4\sigma_{\max}^4$, $c_- = \frac{\min(\xi_{\min},\psi_{\min})}{2\sqrt{m}}\, c_T$ with $c_T = \sigma_{\min}^+\eta_-^2 \big/ \sqrt{2\sigma_{\max}^2\eta_+^4 + a_-^{-2} + b_-^{-2}}$, and $L_{\mathrm{gap}} = \sqrt{2}\max_k \sigma_k\, \sigma_{\max}\left(\begin{smallmatrix} -\sigma_k/(a_-b_-) & 1/b_- \\ 1/a_- & \sigma_k/(a_-b_-) \end{smallmatrix}\right)$, a first-order estimate of the Lipschitz constant of the gap map on W (the proof's L_{gap} is existential; this explicit value is what the experiment uses, and its ordering relative to the theorem's constant is not established, so the assembled κ is an uncertified benchmark). The realized window range of η replaces the regime constant in the same spirit. Regime membership ($\|w\| \le \min(\delta_2, \delta_3)$) is not checked, since those radii are not explicit; the window is defined operationally by the frozen support and the linear-regime start. A failure of the inequality against this benchmark would therefore not refute the theorem, and success does not verify it; the comparison calibrates the empirical deviation against a first-order scale. Five instances are excluded before assembly: four with degenerate kernels ($\dim\ker B > 1$ on the converged frozen support, all in the $n = 20$ sparse antisymmetric family) and one with a non-square support, for which c_- and the decomposition $E \oplus W$ are not defined by the theory. Reported statistics: with the on-support numerator (the on-support energy increment) the inequality $\kappa_{\mathrm{emp}} \le \kappa$ holds on $124/124$ instances, with κ_{emp} median 1.63 / p90 2.35 / max 4.1; the full-energy cross-check agrees ($124/124$; median 1.63 / p90 2.18 / max 4.1); with the telescoped monitor it holds on $119/123$ instances, the 4 exceptions being accumulated-rounding artifacts (Remark 18) that inflate the apparent deviation in the conservative direction. All four hold under the on-support numerator. The assembled κ has median 1.7×10^8, so the median slack is about 10^8.

Calibrated point estimate (verification 7). On the same 129 θ-usable instances and the same linear-regime window, fix a calibration time t_1 at the first guarded increment and, for a sensitivity check, at the one-third and one-half points of the window. At every later guarded increment compute the estimate $\widehat{\mathrm{gap}}(t_2) = \mathrm{gap}(t_1)\sqrt{\dot{N}_{\mathrm{tot}}(t_2)/\dot{N}_{\mathrm{tot}}(t_1)}$ and the multiplicative deviation $D(t_2) = \widehat{\mathrm{gap}}(t_2)/\mathrm{gap}(t_2)$. D is the empirical deviation of the interval proxy: Theorem 5 bounds the pointwise ratio, so no theorem interval applies to D directly, and we compare it against the first-order benchmark κ of verification 6. Record per instance $D_{\max} = \max_{t_2} \max\{D(t_2), D(t_2)^{-1}\}$, the calibration sensitivity (ratio of $D_{\max}$ under the two alternate calibration points to that under the first), and the internal consistency identity $D_{\max} \le \kappa_{\mathrm{emp}}$ against the κ_{emp} of verification 6 (guaranteed by construction; any violation flags an implementation error). Statistics use the on-support numerator ($129/129$ instances): $D_{\max}$ median 1.50 / p90 1.90 / max 3.9; $D_{\max} < 2$ on 91.5% and < 5 on all instances; calibration sensitivity median 1.13; the identity check passes on all instances. The forward gap estimate is therefore accurate to a factor of 2 on more than nine instances in ten, and the calibration point matters at the 13% level. The telescoped monitor version has $D_{\max}$ median 1.52 but diverges (exceeding 10^2, the rounding-floor threshold) on $19/129$ instances, again the accumulated-rounding artifact of Remark 18.

Degenerate-instance sandwich. The 129 θ-usable instances split into three classes by two criteria fixed before assembly: *degenerate* (8) if the snapshot FD Jacobian has more than 2 unit-band eigenvalues; *near-degenerate* (13, all $n = 50$) if the LP equilibrium margin is below 1.2×10^{-3}; *regular* (108) otherwise. The unit-band count is a numerical flag (see the FD paragraph). The structural test is the kernel dimension of B on the converged frozen support, and on it the 8 split as 5 with $\dim\ker B > 1$ or a non-square support (the verification 6 exclusions, where c_- and $E \oplus W$ are undefined) and 3 with $\dim\ker B = 1$ (a third unit-band eigenvalue at the snapshot that is not a kernel mode of the converged B). For those 3, κ is assembled and the inequality holds, so they stay in the count: $124 = 129 - 5$. The *empirical* sandwich is defined on all 8 because it only reads the trajectory: over the θ-usable window κ_{emp} has median 1.80 on the 8 (vs. 1.63 on the 108 regular

and 1.61 on the 13 near-degenerate instances), and the estimate deviation stays below 2 on 7 of the 8 (below 5 on all). The ratio law itself does not break at degeneracy; this paper's constants do. That supports using the certificate as a monitor.

EFG monitoring (verification 8). The extensive-form solver runs on the two small poker games Kuhn (3 cards, 12 information sets) and Leduc (3×2 deck, 1116 information sets), one IR-PRM$^+$ instance per information set (a player's decision point), each fed the counterfactual utilities of CFR (Zinkevich et al., 2007) in the extra-gradient setup and implemented vectorized in-house following the scale-invariant algorithm of Zhang et al. (2025a). Traces are recorded at logarithmic times plus a linear block over the first 2000 iterations. The monitored quantity is the cumulative update energy $G(t) = \sum_{s \le t} \sum_j \|g_j^{(s)}\|^2$ over information sets j (the extensive-form analogue of the full-energy cumulant of the matrix protocol), already computed by the solver each iteration; as in the matrix protocol we use its per-unit-time increment $(G(t_{k+1}) - G(t_k))/(t_{k+1} - t_k)$ between recorded times. The aggregate squared regret norm $Q(t) = \sum_j \|\tilde{r}_j^{(t)}\|^2$, a level quantity, is recorded alongside; its per-unit-time increment is not used for the headline numbers: it couples to the gap at 0.92 on Kuhn and 0.998 on Leduc, against 0.998 and 0.997 for the update energy. The gap is computed from a best response over the full tree (two passes, one per player) at the same recorded times. The window and guards mirror the matrix protocol (floor 10^{-13}, representability guard, at least 6 guarded points); Kuhn runs to $T = 10^5$ and Leduc to $T = 5 \times 10^4$, all on CPU. The best-response cost is measured by wall clock. The ratio of one $\mathrm{NashConv}$ pass (one tree traversal per player, computing the extensive-form gap by summing the two players' best-response gains) to one solver iteration is 0.076 on Kuhn and 0.087 on Leduc. Replacing an every-iteration NashConv schedule by the sparse record schedule uses $126\times$ (Kuhn) and $64\times$ (Leduc) fewer best-response passes, the ratio $T/\#\{\text{recorded times}\}$; the saving is in the schedule. The coupling is reported as $\mathrm{corr}\big(\log \frac{G(t_{k+1})-G(t_k)}{t_{k+1}-t_k},\ \log \mathrm{gap}(t_{k+1})\big)$ on the guarded window, and the estimate deviation as in verification 7 with G in place of G_S.

E.4 Excluded and degenerate instances: summary

We collect all exclusions in one place. (i) Spectral comparison: $155/216$ with a fitted slope; 32 no-snapshot (numerical-zero convergers, final gap exactly 0.0, all $n \in \{5, 20\}$, for which the positive-gap snapshot trigger never fires; see the snapshot paragraph), 29 no fittable exponential segment (fast convergence to the float floor). (ii) θ measurement: $129/216$ slope-measurable; the same 32 numerical-zero convergers plus fast floor-hitters with fewer than 6 representable increments. (iii) Support match: the frozen pattern differs from the LP support in 15.7% of instances; multiple equilibria are one possible explanation, not an established classification; this is why the protocols for verifications (3) and (5) in this appendix rebuild all spectral quantities from the converged state. (iv) Degenerate equilibria: instances with more than 2 unit-band eigenvalues (up to 10) are flagged as candidates for a positive-dimensional equilibrium manifold and confirmed for five, refuted for three, by the kernel dimension of the converged B; the closed-form ρ_2 remains exact there, the kernel modes landing in the unit band. (v) Ratio-certificate assembly (verification 6): 5 of 129 excluded (four with $\dim \ker B > 1$, one non-square support), as itemized in the protocol paragraph above; every exclusion is recorded per instance alongside the trace it came from, and the list is complete.

E.5 Summary of verifications

Table 3 collects the full statistics behind Table 1 across all 216 instances. Reported ranges span the 10th to the 90th percentile across instances.

F Discussion

Error bounds, active-set identification, and metric subregularity (extended related work). The error-bound route gives global linear convergence over polyhedra under a projection-residual error bound (Tseng, 1995). Our rate comes from the local spectrum of the linearized one-round map on the frozen face (Theorem 4), and the operator of IREG-PRM$^+$ satisfies a Minty condition (Cai et al., 2025). Finite-time identification followed by local linear convergence has precedents outside games (Burke & Moré, 1988; Hare & Lewis, 2004; Liang et al., 2017). The state-to-strategy map

Table 3: Full statistics behind Table 1, on 216 random matrix games (three dimensions $\times$ three distributions $\times$ two generators $\times$ 12 seeds). †Self-stabilization heuristic, empirically consistent with Proposition 2.

Prediction	Source	n_{inst}	Statistic
Exact support freezing before $T/2$	Lem. 6	216	frozen on 99.5%; median $t_{\text{freeze}}/T = 2.7 \times 10^{-4}$
$\gamma \to v^*$ at the fixed point	Prop. 1	216	median $\lvert\gamma - v^*\rvert = 4.47 \times 10^{-17}$
Rate equals FD $\rho_2(J_{\text{FD}})$	Thm. 4	155	ratio median 1.000 (p10–p90 0.984–1.016); corr. 0.9982
Closed form $\rho_2 =$ FD ρ_2	Lem. 9, 10	32	relative error median 2.4×10^{-10}, max 4.1×10^{-7}
Certificate exponent $\theta = 2$	Cor. 2	129	on-support median 1.986; 96.1% between 1.8 and 2.2; cross-check 1.995, 98.4%
$\eta_\infty \sigma_{\max}(B) \le 1$ at limits	Rem. 1†	184	0 violations; median 0.790, max 0.9985; 18 above 0.95; closest approach of $(\eta_\infty \sigma_{\max}(B))^2$ to 1 is 0.0029
Ratio certificate (empirical proxy)	Thm. 5	124	holds on 124/124; κ_{emp} median 1.63, max 4.1
Point-estimate forecast (empirical proxy)	Thm. 5	129	max forward deviation median 1.50, p90 1.90, max 3.9; below 2 on 91.5%
Progress monitoring (Kuhn, Leduc)	empirical	2	max forward deviation 1.80, 3.22; Nash-Conv pass 0.076/0.087 of an iteration
Degenerate-instance sandwich	empirical	8	κ_{emp} median 1.80 vs. 1.63; below 2 on 7/8

of IREG-PRM$^+$ is not a proximal operator, so the identification argument has to pass through that Minty condition (Anagnostides et al., 2025a; Zhang et al., 2025b; Pethick et al., 2025).

Computable progress metrics (extended related work). The closest work, the normalized duality gap of Applegate et al. (2023), is compared in Section 2. Its constants are existential or require vertex enumeration (Wei et al., 2021, Appendix G), and it measures distance to the whole solution set, so it stays finite at degenerate equilibria. Ours rests on a two-dimensional fixed-point cone and degenerates with it; on that subset the deviation is bracketed empirically, with κ_{emp} at median 1.80 against 1.63 on regular instances, and below 2 on 7 of 8 instances (Table 3). Utility-increment certificates (Burch et al., 2019) and the gap-to-regret identity (Zinkevich et al., 2007) are one-sided bounds on the averaged point, whose guarantee is $O(1/T)$. Ours is two-sided on the last-iterate and is an increment ratio, like the step ratio in the certificate of Applegate et al. (2023), measured on quantities the algorithm already computes. No per-iteration best-response pass is added, in normal or extensive form (Section 5).

From saturation to convergence, and the constant. The bound of Theorem 1 is explicit; on the testbed it sits three to six orders of magnitude above the realized $N_{i,\infty}$. It leaves the *value* of the saturation level undetermined, and that value selects the effective step size η_∞ and hence whether Assumption 2 holds. Between the global theorem (gap $\to 0$ on every game) and the local one (linear rate under the strict spectral condition) sits a quantitative theory of where the norm stops. The boundary $\eta_\infty \sigma_{\max}(B) = 1$ and the exceptional set of Proposition 2 are open (Appendix G).

Quantitative neighborhood radii. The absorbing radius around the frozen face and the constant C_ε in the rate envelope are not closed form; they depend on the equilibrium margin and other instance quantities through non-explicit constants. A quantitative version, with finite-time control of the $\theta = 2$ law of Corollary 2, would upgrade the ratio certificate from a two-sided online bound into an a priori complexity bound. Concretely, the constants $C(t)$ in Corollary 2 bracket the ratio by an existential κ, while the first-order benchmark assembled in Section 5 exceeds the measured deviation by eight orders of magnitude. Making L_{gap}, c_-, and the spectral range explicit is the technical route recorded in Appendix G.

Degenerate and non-unique equilibria. The closed-form ρ_2 stays exact at degenerate and large-support instances, but *point* convergence is not established there. The fixed-point set is the positive

cone over the NE polytope (Proposition 1), positive-dimensional when the equilibrium is not unique; Theorem 2 places the ω-limit set inside it. On the testbed no drift along that set is visible. On the 8 instances whose B has a kernel of dimension 2 or more on the converged support (Appendix E, verification 6), the state $z = (\tilde{r}_1, \tilde{r}_2)$ moves by at most 5×10^{-10} and the strategy pair by at most 4×10^{-12}, from the first checkpoint at which the gap is below 10^{-15} to the end of the run, 10^5 to 10^6 rounds later. Two regimes remain open: drift along the equilibrium set, and the non-hyperbolic boundary $\eta_\infty \sigma_{\max}(B) = 1$, where 18 of the 184 measurable instances have $\eta_\infty \sigma_{\max}(B) > 0.95$. The natural extension is a normally hyperbolic invariant manifold argument along a smooth stratum of that set, transporting the frozen-face spectral computation onto the normal bundle (Appendix G); the boundary case itself needs a center-manifold analysis of the $q = 1$ mode (Appendix G).

Extensive-form games and partial feedback. The certificate's payoff is largest in extensive-form games, where $N(t)$ aggregates per-infoset regret norms from quantities the update already produces, while the gap still requires a best-response pass. Extending the frozen-face analysis to counterfactual-value couplings is open, as is the certificate under sampled feedback, where the norm increment is dominated by sampling noise. Both have concrete starting points: a per-information-set B_j coupled through counterfactual weights (Appendix G), and variance reduction through correlated chance sampling (Appendix G).

G Future Work

The open quantitative link is the saturation *level*, which decides the effective step size and thus the spectral condition of Theorem 4. A bound on $N_{i,\infty}$ tight enough to decide $\eta_\infty \sigma_{\max}(B) < 1$ would turn the measured middle link of Remark 1 into a theorem. Convergence results for regret matching in potential games (Anagnostides et al., 2025b) and rate refinements for zero-sum problems (Anagnostides & Sandholm, 2024) suggest that the mechanism carries beyond the matrix-game testbed. Making the absorbing radius and the constant C explicit would upgrade the ratio certificate of Theorem 5 from a two-sided online bound into an a priori complexity bound, the most directly useful extension of the local theory.

The natural route to extend point convergence beyond the uniqueness hypothesis replaces the isolated fixed point by a normally hyperbolic invariant manifold and transports the frozen-face spectral computation of Theorem 4 onto its normal bundle. At the boundary $\eta_\infty \sigma_{\max}(B) = 1$ the relevant eigenvalues are $\mu = \pm i$, the fixed point is non-hyperbolic, and a center-manifold analysis of the $q = 1$ mode would sharpen the generic inequality into a dichotomy (Remark 15). Support collapse makes the one-round map non-invertible, so a full-space statement needs absolute continuity of the push-forward of Lebesgue measure (Remark 13); the genericity questions for optimistic dynamics in general-sum games (Anagnostides et al., 2022a;b) are the closest analogues.

A finite-time version of the slope-two law of Corollary 2 would make the phase-dependent offset $C(t)$ and its $O(\|w^{(t)}\|)$ remainder explicit and supply an operational regime-detection rule (Remark 21). The per-player analogue admits a three-step-window repair for the linearized dynamics (Remark 20). Under sampled feedback the norm increment is dominated by sampling noise, and variance reduction through correlated chance sampling (Li et al., 2026) acts on exactly that quantity; restricting responses under a confidence schedule for safe opponent exploitation (Li & Huang, 2026a) is the natural analogue of our monitor in that setting. The extensive-form object to be built is a per-information-set B_j coupled through counterfactual weights; the efficient, parallel, and abstracted CFR solvers of Li & Huang (2026c; 2025; 2026b); Li et al. (2024) and the black-box equilibrium search of Zhang & Sandholm (2021) are natural starting points, together with discounted regret minimization (Brown & Sandholm, 2019) and predictive Blackwell approachability (Farina et al., 2021).

H Conclusion

The linear last-iterate convergence of IREG-PRM$^+$ has a mechanical explanation. The algorithm's own normalization makes the regret norm an inverse step size; the norm stops growing on every matrix game; and from then on every round is a fixed-step extra-gradient iteration on a frozen support face whose spectrum is explicit. Two consequences reach beyond the theorems. The rate is decided by one scale-invariant number: the saturated step size times the largest singular value of the value-centered

support matrix. What a scale-invariant algorithm must control to converge fast is therefore where its norm stops, and on the testbed the realized norms stop on the contracting side of that boundary on every measurable instance. The quantity that certifies progress, the squared-norm increment, is computed by the algorithm in every round. In extensive-form games the best-response pass that defines the gap can therefore be scheduled sparsely, with the increment tracking the gap between passes. Two questions stay open: the saturation level itself, and drift when the equilibrium is not unique (Appendices F–G).