EDGE OF STABILITY OCCURS IN MINIMAX OPTIMIZATION

Anonymous authors
Paper under double-blind review

ABSTRACT

The Edge of Stability (EoS) phenomenon is widely observed in minimization, where gradient methods continue to decrease the loss even while remaining near their stability thresholds, a behavior associated with a self-stabilization mechanism. Yet an analogous phenomenon has not been clearly identified in minimax optimization. In this paper, we formulate EoS for minimax optimization, demonstrate that it indeed occurs, and characterize the self-stabilization mechanism underlying it. We first define EoS for minimax optimization in terms of possibly complex dominant Jacobian eigenvalues reaching the boundary of the stability region. Using two-player zero-sum games, we observe minimax EoS for both Extragradient (EG) and Optimistic Gradient Descent Ascent (OGDA). However, when the dominant eigenvalues lie near the real axis, OGDA still exhibits EoS whereas EG does not. Finally, we generalize the self-stabilization mechanism to minimax methods and show how it explains these behaviors.

1 INTRODUCTION

In minimization, the stability of gradient descent (GD) is closely tied to the largest eigenvalue of loss Hessian. During neural network training, this eigenvalue often progressively increases toward the classical stability threshold $2/\eta$, where η is the learning rate, a phenomenon known as progressive sharpening (PS) (Cohen et al., 2021). Remarkably, upon reaching this threshold, training does not immediately diverge; instead, the eigenvalue remains near, and often slightly above, the threshold while the loss continues to decrease. This behavior is known as the Edge of Stability (EoS) (Cohen et al., 2021). Several works have sought to understand the dynamics underlying EoS (Arora et al., 2022; Ma et al., 2022; Ahn et al., 2022; Lyu et al., 2022; Wang et al., 2022; Cohen et al., 2025). In particular, Damian et al. (2023) identify a self-stabilization mechanism of GD, in which instability along the top Hessian eigendirection near the stability threshold induces a feedback that reduces the largest eigenvalue and thereby keeps the dynamics near the edge of stability.

In contrast, an analogous EoS phenomenon has not been clearly identified, either empirically or theoretically, in minimax optimization.[1] This absence is particularly intriguing because minimax gradient methods already have well-established convergence guarantees under certain stability conditions (Mokhtari et al., 2020; Diakonikolas et al., 2021; Böhm, 2023), including nonexpansiveness of the update operator (Gorbunov et al., 2023). Yet it remains unclear whether minimax training progressively approaches these stability boundaries and continues to make progress near them, as observed in minimization. This raises a natural question:

Does the Edge of Stability also occur in minimax optimization?

In this paper, we show that EoS indeed occurs in minimax optimization. To characterize EoS beyond minimization, we formulate it in terms of the stability region in which the update operator is nonexpansive, whose boundary reduces to the classical threshold $2/\eta$ in minimization. We use the Scaled Relative Graph (SRG) (Ryu et al., 2022; Huang et al., 2026) to represent this stability region in the complex plane, allowing us to visualize and analyze PS and EoS in minimax dynamics. We derive the corresponding stability regions for two representative minimax methods, Extragradient

[1] A few works, *e.g., (Long & Bartlett, 2024; Feng et al., 2026)*, study EoS in minimax settings, but primarily from a minimization-based perspective.

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.

(EG) (Korpelevich, 1977) and Optimistic Gradient Descent Ascent (OGDA) (Popov, 1980). We then use the dominant eigenvalues of the saddle-gradient Jacobian reaching the boundary of the stability region to formalize EoS in minimax optimization.

Since the dominant eigenvalues of the saddle-gradient Jacobian can be complex-valued, EoS in minimax optimization can occur along trajectories in the complex plane rather than only along the real axis as in minimization. We investigate this form of minimax EoS across three two-player zero-sum game settings: a two-stage Rock-Paper-Scissors game, a pursuit-evasion game (Isaacs, 1951), and a minimax formulation of neural optimal transport (Korotin et al., 2023). Across these settings, we observe that both EG and OGDA exhibit EoS. When the dominant eigenvaluse lie near the real axis, however, EG becomes unstable, whereas OGDA continues to exhibit EoS.

We then generalize the self-stabilization framework of Damian et al. (2023) to minimax methods. Unlike GD, for which self-stabilization is automatic, this property is not guaranteed for minimax methods and depends on the dynamics of the dominant eigenvalues and the update operator. We characterize conditions under which EG and OGDA self-stabilize, which also explains their different EoS behaviors. When a minimax method exhibits self-stabilization, we further show that its training dynamics are closely approximated by projected dynamics constrained to keep the eigenvalues of the saddle-gradient Jacobian within the stability region, extending Damian et al. (2023, Theorem 1) to minimax optimization.

Our main contributions are summarized as follows:

- **Minimax EoS Formulation.** In Section 3, we use the Scaled Relative Graph to describe the stability regions of minimax methods, and formulate the minimax notion of the Edge of Stability (EoS) relative to these regions, using the dominant eigenvalues of the saddle-gradient Jacobian that reach the stability boundary.
- **Empirical Evidence of Minimax EoS.** In Section 4, we provide empirical evidence across three two-player zero-sum game settings that both Extragradient (EG) and Optimistic Gradient Descent Ascent (OGDA) exhibit EoS. When the the dominant eigenvalues of the saddle-gradient Jacobian lie near the real axis, however, EG becomes unstable, whereas OGDA continues to exhibit EoS.
- **Self-Stabilization of Minimax Methods.** In Section 5, we introduce a self-stabilization property for minimax methods and study when it holds for EG and OGDA. In Section 6, we show how this property gives rise to a self-stabilization mechanism and further show that, when a minimax method self-stabilizes, its training dynamics are closely approximated by projected dynamics that keep the the eigenvalues of the saddle-gradient Jacobian within the stability region.

1.1 Notations

For $x \in \mathbb{C}\setminus\{0\}$, let $\arg(x) \in (-\pi, \pi]$ denote its principal argument. For a vector $v \in \mathbb{C}^d$, $v^\top$ denotes its transpose, v^H its conjugate transpose, and $\|v\|$ its Euclidean norm. For $v, w \in \mathbb{R}^d$, we write $v \cdot w := v^\top w$. For a square matrix A, $\mathrm{eig}(A)$ denotes the set of its eigenvalues. An order-k tensor $\mathcal{A} \in \mathbb{C}^{d_1 \times \cdots \times d_k}$ is an array with entries $\mathcal{A}_{i_1,\ldots,i_k}$. For vectors $v_j \in \mathbb{C}^{d_j}$, its tensor multiplication is defined by $\mathcal{A}[v_1, \ldots, v_k] := \sum_{i_1,\ldots,i_k} \mathcal{A}_{i_1,\ldots,i_k} (v_1)_{i_1} \cdots (v_k)_{i_k}$. For a set A, we denote its boundary by ∂A.

2 Preliminaries

Stability Threshold of Gradient Descent Consider gradient descent (GD), $z_{t+1} = z_t - \eta \nabla g(z_t)$, with step size η for the minimization problem $\min_z g(z)$. For a twice-differentiable g, let $S(z)$ be the largest eigenvalue of the Hessian $\nabla^2 g(z)$. If $S(z) \le L$ for all z, where $L > 0$, GD satisfies the descent lemma (Nesterov, 2018)

$$g(z_{t+1}) \le g(z_t) - \eta \left(\frac{2 - \eta L}{2}\right) \|\nabla g(z_t)\|^2, \qquad \forall t. \tag{1}$$

This classical guarantee motivates $2/\eta$ as the stability threshold for $S(z)$ in minimization.

Progressive Sharpening and Edge of Stability In practical neural network training, however, such a global bound on $S(z)$ is rarely available (Zhang et al., 2020). Instead, $S(z)$ progressively increases toward $2/\eta$, a phenomenon known as *Progressive Sharpening* (Cohen et al., 2021). Once $S(z)$ reaches $2/\eta$, it typically stops increasing and oscillates around this threshold. Although the descent condition in equation 1 no longer holds, training does not diverge; instead, the loss continues to decrease. This phenomenon is referred to as *Edge of Stability* (Cohen et al., 2021).

Self-Stabilization Property and Mechanism Damian et al. (2023, Lemma 2) establish the following self-stabilization property of GD:

$$\nabla S(z) = \nabla^3 g(z)[u(z), u(z)], \tag{2}$$

where $u(z)$ is the normalized eigenvector corresponding to $S(z)$. When $S(z)$ exceeds the stability threshold $2/\eta$, the GD dynamics become unstable along $u(z)$. Expanding the gradient along $u(z)$, this growing component induces a term in the GD update proportional to $-\nabla S(z)$, thereby reducing $S(z)$. This negative-feedback effect is called the self-stabilization mechanism. With this mechanism, Damian et al. (2023) show that the GD trajectory is closely approximated by the projected GD trajectory $z_{t+1}^\dagger = \mathrm{proj}_{\mathcal{M}}(z_t^\dagger - \eta \nabla g(z_t^\dagger))$ for $\mathcal{M} := \{z : S(z) \le 2/\eta,\ \nabla g(z)^\top u(z) = 0\}$. The projected trajectory corresponds to GD applied to the constrained problem $\min_{S(z)\le 2/\eta} g(z)$.

Minimax Optimization In this paper, we consider the minimax optimization problem

$$\min_x \max_y \ f(x, y), \tag{3}$$

where f is twice differentiable. Let $z = (x, y)$ and $F(z) := (\nabla_x f(z), -\nabla_y f(z))^\top$ denote the saddle-gradient operator of f. We consider the following two representative minimax methods:

- Extragradient (EG) (Korpelevich, 1977): $z_{t+1} = z_t - \eta F(z_t - \eta F(z_t))$,
- Optimistic Gradient Descent Ascent (OGDA) (Popov, 1980): $z_{t+1} = z_t - \eta\big(2F(z_t) - F(z_{t-1})\big)$.

Unlike Gradient Descent Ascent, which diverges even on simple bilinear problems (Mescheder et al., 2018), both EG and OGDA are guaranteed to converge for broad classes of minimax problems (Mokhtari et al., 2020; Diakonikolas et al., 2021; Böhm, 2023; Gorbunov et al., 2023).

3 Generalizing EoS: From Minimization to Minimax Optimization

3.1 Nonexpansiveness as a Stability Criterion

Consider an iterative method of the form

$$z_{t+1} = z_t - \eta T(z_t) = (I - \eta T)(z_t), \qquad \forall t. \tag{4}$$

A standard stability property of fixed-point iterations is the nonexpansiveness of the update operator. We consider the following stronger form, which reduces to nonexpansiveness when $\alpha = 0$:

$$\|(I - \eta T)(z) - (I - \eta T)(\tilde{z})\|^2 \le \|z - \tilde{z}\|^2 - \alpha \|T(z) - T(\tilde{z})\|^2, \quad \forall z, \tilde{z}, \tag{5}$$

for some $\alpha \ge 0$. Such inequalities are widely used in convergence analyses of minimax methods (Diakonikolas et al., 2021; Böhm, 2023; Gorbunov et al., 2023).

Recovering the classical stability threshold of GD For convex and L-smooth g, the cocoercivity of ∇g yields equation 5 for $T = \nabla g$ with $\alpha = \eta\left(\frac{2}{L} - \eta\right)$. Hence, the GD is nonexpansive when $\eta \le 2/L$, recovering the same classical stability threshold $L = 2/\eta$. The same threshold applies locally to nonconvex objectives whenever the Hessian has no negative eigenvalues. This correspondence provides a natural basis for using nonexpansiveness to extend the stability notion underlying EoS to minimax optimization.

3.2 Geometric Representation of Stability Regions via SRG

To represent this stability criterion geometrically and relate it to the eigenvalues of the Jacobian, we use the Scaled Relative Graph (SRG) (Ryu et al., 2022). The SRG of T is defined as

$$\mathcal{G}(T) := \left\{ \tfrac{\|T(z) - T(\tilde{z})\|}{\|z - \tilde{z}\|} \exp\left(\pm i \angle (T(z) - T(\tilde{z}), z - \tilde{z})\right) :\ z \ne \tilde{z} \right\},$$

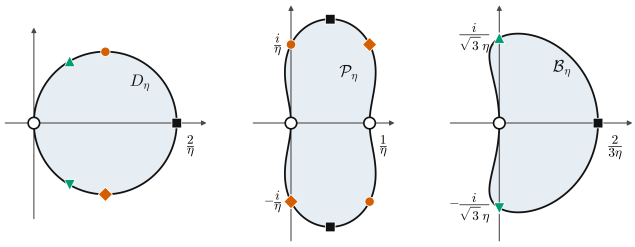

Figure 1: Stability regions in the complex plane. Left: Stability region associated with the eigenvalues of DT for general iterative methods. Middle and right: Corresponding stability regions for the eigenvalues of DF for EG and OGDA, respectively. Matching markers indicate the correspondence between the eigenvalues of DT and DF; see Appendix A.2.

where $\angle(a, b) \in [0, 2\pi)$ is the angle between a and b. It follows that, for any $\ell > 0$, the stability inequality in equation 5 holds with $\alpha = \eta(\ell - \eta)$ if and only if

$$\mathcal{G}(T) \subset D_\ell := \{\omega \in \mathbb{C} : |\omega - 1/\ell| \leq 1/\ell\},$$

see Proposition A.1. Hence, if $\mathcal{G}(T) \subset D_\ell$, the iteration is nonexpansive whenever $\eta \leq \ell$.

The SRG stability condition $\mathcal{G}(T) \subset D_\eta$ is directly related to the eigenvalues of the Jacobian DT. In particular, it implies $\mathrm{eig}(DT(z)) \subset D_\eta$ for all z. The converse also holds when $DT(z)$ is normal for all z; see Proposition A.2.

For minimax methods, such as EG and OGDA, whose updates are defined in terms of F, the Jacobian stability condition $\mathrm{eig}(DT(z)) \subset D_\eta$ can be approximately[2] translated into a corresponding condition on the eigenvalues of the saddle-gradient Jacobian DF. In particular, for EG and OGDA, the condition $\mathrm{eig}(DT(z)) \subset D_\eta$ corresponds approximately to $\mathrm{eig}(DF(z)) \subset \mathcal{P}_\eta$ and $\mathrm{eig}(DF(z)) \subset \mathcal{B}_\eta$, respectively, where

$$\mathcal{P}_\eta := \left\{\omega \in \mathbb{C} : \omega - \eta\omega^2 \in D_\eta\right\} \quad \text{and} \quad \mathcal{B}_\eta := \left\{\omega \in \mathbb{C} : \omega + \tfrac{1}{2\eta} \pm \sqrt{\omega^2 + \tfrac{1}{4\eta^2}} \in D_\eta\right\}, \quad (6)$$

see Appendix A.2. We refer to $\mathcal{P}_\eta$ and $\mathcal{B}_\eta$ as the *stability regions* of EG and OGDA, respectively.

3.3 Progressive Sharpening and Edge of Stability in Minimax Optimization

We now formulate PS and EoS in minimax optimization using the saddle-gradient Jacobian DF and the stability regions. We refer to the progressive increase in the magnitudes of some eigenvalues of $DF(z_t)$ during training as *Progressive Sharpening*. We let $S(z_t)$ denote the dominant eigenvalue, defined as one that reaches the stability boundary of the given method.[3] In minimization, $S(z_t)$ reduces to the largest Hessian eigenvalue, and in our minimax experiments (Section 4), it typically corresponds to the largest-magnitude eigenvalue. When $S(z_t)$ reaches the boundary and remains near it while training continues to make progress, we refer to this regime as the *Edge of Stability*.

4 Edge of Stability Occurs in Minimax Optimization

We investigate minimax EoS in three two-player zero-sum game settings: a two-stage Rock-Paper-Scissors game, a pursuit-evasion game (Isaacs, 1951), and a minimax formulation of neural optimal transport (Korotin et al., 2023). In each setting, we parameterize each player's strategy with a neural network. Detailed game settings are given in Appendix B.

4.1 EoS Occurs Across Two-Player Zero-Sum Game Settings

We start with the Rock-Paper-Scissors (RPS) game. However, training dynamics converge before progressive sharpening reaches the stability boundary, possibly due to its small state-action space.

[2] In general, the eigenvalues of DT cannot be determined exactly from those of DF. Assuming that the Jacobian varies only slightly over a single update, we obtain the approximate stability regions in equation 6. As shown in Section 4, their boundaries accurately predict the empirical behavior.

[3] When the eigenvalues reaching the stability boundary form a complex-conjugate pair, we choose the one with nonnegative imaginary part.

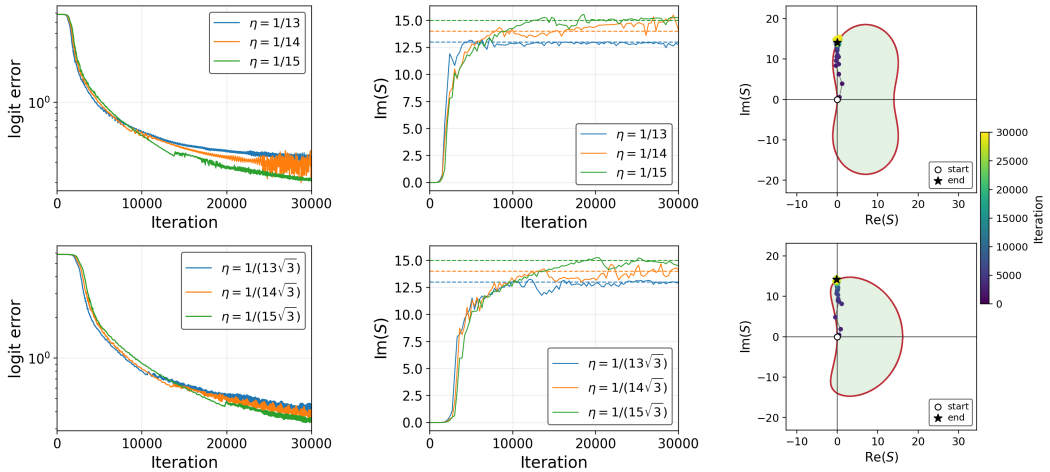

Figure 2: Experiment on the two-stage variant of the RPS game. Top and bottom rows show EG and OGDA, respectively. Right shows the trajectories of $S(z_t)$ in the complex plane.

We therefore consider a two-stage variant of the RPS game with an enlarged state-action space; see Appendix B for details. Figure 2 shows that, for both EG and OGDA, $S(z_t)$ progressively reaches the stability boundary along the imaginary direction and then oscillates near the boundary. Throughout PS and EoS, the logit error (see Definition B.3) continues to decrease.

We also study the pursuit–evasion (PE) game (Isaacs, 1951), a multi-stage Markov game that allows us to examine EoS over a longer horizon. As shown in Figure 7 in Appendix B, $S(z_t)$ progressively approaches the stability boundary away from the real axis and remains near it for both EG and OGDA, while the logit error continues to decrease.

4.2 Near the Real Axis: OGDA Exhibits EoS, While EG Is Unstable

In both settings above, $S(z_t)$ approaches the stability boundary away from the real axis, and both EG and OGDA exhibit EoS. Surprisingly, a qualitative difference between EG and OGDA emerges when $S(z_t)$ approaches the stability boundary near the real axis: OGDA exhibits EoS, whereas EG does not. We first consider the simple minimization problem of Liu et al. (2025), in which PS occurs entirely along the real axis. As shown in Figure 3, OGDA reaches the stability boundary and continues to decrease the loss, whereas EG becomes unstable after crossing its stability boundary, with the loss increasing sharply rather than exhibiting EoS. This contrast also appears in minimax settings, including another PE setting and a minimax formulation of neural optimal transport (Korotin et al., 2023), even when $S(z_t)$ lies close to, but not exactly on, the real axis; see Figures 8 and 9 in Appendix B.

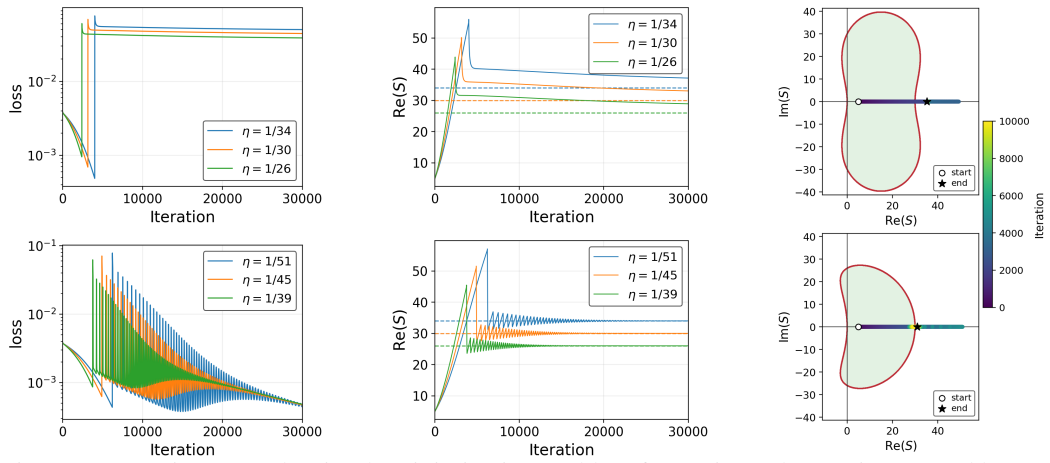

Figure 3: Experiment on the simple minimization problem from Liu et al. (2025). Top and bottom rows show EG and OGDA, respectively. Right shows the trajectories of $S(z_t)$ in the complex plane.

This difference is notable because both EG and OGDA have similar convergence guarantees under standard minimax assumptions (Mokhtari et al., 2020; Diakonikolas et al., 2021; Böhm, 2023; Gorbunov et al., 2023). In the next section, building on Damian et al. (2023), we explain the EoS behaviors observed in Sections 4.1 and 4.2 through self-stabilization.

5 Self-Stabilization Property of Minimax Methods

In this section, we generalize the self-stabilization property of GD in equation 2 to minimax methods and characterize when a minimax method has this property.

5.1 Generalizing Self-Stabilization Property

The self-stabilization property of GD can be understood as follows: when the largest eigenvalue of $D(\nabla g)$ exceeds the threshold $2/\eta$, the term $-D^2(\nabla g)$ in the GD update aligns with the direction that pushes the eigenvalue back below $2/\eta$. This suggests that self-stabilization is characterized by how the dominant eigenvalue moves relative to the stability boundary. Analogously, we generalize the self-stabilization property to minimax methods by asking whether $-D^2T$ aligns with the direction that moves an eigenvalue of DT back into the stability region D_η after it leaves D_η.

To formalize this, let $\Lambda(z)$ be the eigenvalue of $DT(z)$ that leaves D_η. It corresponds to the dominant eigenvalue $S(z)$ of $DF(z)$; for example, $\Lambda(z) \approx S(z) - \eta S(z)^2$ for EG (see equation 6). To characterize the motion of $\Lambda(z)$ relative to D_η, we define

$$\phi(z) := \frac{|1 - \eta\Lambda(z)|^2 - 1}{2\eta},$$

so that $\phi(z) > 0 \iff \Lambda(z) \notin D_\eta$. We then quantify how $-D^2T$ aligns with $-\nabla\phi$ through the self-stabilization coefficient:

$$\beta(z) := \nabla\phi(z) \cdot D^2T(z)[u(z), \bar{u}(z)],$$

where $u(z)$ is the eigenvector of $\Lambda(z)$, normalized $\|u\| = 1$ if $\Lambda(z) \in \mathbb{R}$ and $\|u\| = 2$ otherwise.

If $\beta(z) > 0$, then $-D^2T(z)[u(z), \bar{u}(z)]$ has a component in the direction that decreases $\phi(z)$, thus moving $\Lambda(z)$ toward D_η. We say that T has the *self-stabilization property* at z if $\beta(z) > 0$. This recovers the self-stabilization property of GD, for which $\beta = \|\nabla S\|^2 > 0$ whenever $\nabla S \neq 0$, following directly from equation 2. For minimax methods, however, the sign of β depends on both the dynamics of $\Lambda(z)$ (and the associated $S(z)$) and the update operator T.

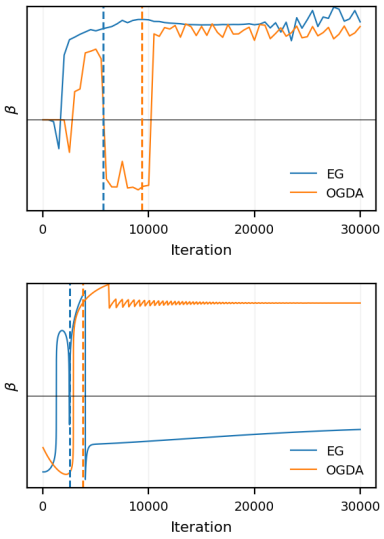

Figure 4: Self-stabilization coefficient $\beta(z_t)$ for two-stage RPS (top) and simple minimization (bottom). The dashed line indicates the first iteration at which $\Lambda(z_t) \notin D_\eta$.

Figure 4 plots the dynamics of β for the two-stage RPS game and the simple minimization problem in Figures 2 and 3, respectively, to examine the relation between self-stabilization and EoS. In the cases with EoS, $\beta(z_t)$ is either already positive when $\Lambda(z_t)$ reaches ∂D_η or becomes positive shortly afterward. In contrast, for the latter problem with EG, where EoS does not occur, $\beta(z_t)$ becomes negative after $\Lambda(z_t)$ crosses the stability boundary and remains negative thereafter. This shows that the self-stabilization property plays a key role in inducing EoS.

5.2 When Does the Self-Stabilization Property Hold?

To understand when the self-stabilization property holds, we relate the sign of β to the location of Λ relative to the other eigenvalues of DT, assuming DT is normal. In Appendix C.2, we empirically show that DT remains approximately normal throughout training.

Proposition 5.1. *Suppose $T \in C^2(\mathbb{R}^d;\mathbb{R}^d)$ with $d \geq 2$, DT is normal in a neighborhood of z, and $\Lambda(z) \in \partial D_\eta$ is simple. Then for all $\lambda \in \mathrm{eig}(DT(z)) \setminus \{\Lambda(z), \bar{\Lambda}(z)\}$,*

$$\begin{aligned} |\Lambda(z)| \geq \sqrt{2}|\lambda| \qquad &\Longrightarrow \quad \beta(z) \geq 0, \\ |\Lambda(z)| \leq |\lambda|\sqrt{\tfrac{2}{1+|1-\eta\lambda|^2}} \quad \textit{and} \quad |\operatorname{Re}\lambda - \Lambda(z)| \geq |\operatorname{Im}\lambda| \quad &\Longrightarrow \quad \beta(z) \leq 0. \end{aligned} \tag{7}$$

See Appendix C.1 for the proof.

Proposition 5.1 helps explain why $\beta(z_t) \leq 0$ for EG when $S(z_t)$ lies near the real axis. As $S(z_t)$ approaches the EG stability boundary $1/\eta$ on the real axis, the corresponding eigenvalue $\Lambda(z_t) \approx S(z_t) - \eta S(z_t)^2$ approaches 0; see Figure 1. Hence, as $|\Lambda(z_t)|$ decreases, the sufficient condition for $\beta \leq 0$ becomes easier to satisfy. In Appendix C.2, we further show that the sufficient conditions in Proposition 5.1 for either $\beta \geq 0$ or $\beta \leq 0$ are typically satisfied in our experiments. We leave a complete characterization of the sign of β in terms of $\Lambda(z)$ and T for future work.

6 Self-Stabilization Mechanism of Minimax Methods

Building on Damian et al. (2023), we study the self-stabilization mechanism of minimax methods under $\beta > 0$. We first explain how the mechanism gives rise to EoS, and then show that the minimax dynamics closely track the following projected dynamics:

$$z_{t+1}^\dagger := \mathrm{proj}_{\mathcal{M}}\big(z_t^\dagger - \eta T(z_t^\dagger)\big), \tag{8}$$

where $\mathcal{M} := \{z : \phi(z) \leq 0,\ u(z)^H T(z) = 0\}$. For the analysis, we make the following assumptions along the projected trajectory.

Assumption 6.1. Let $\alpha(z) := -\nabla\phi(z) \cdot T(z)$ and assume $\alpha(z_t^\dagger) > 0$, $\beta(z_t^\dagger) > 0$, and $DT(z_t^\dagger)$ is normal for every t.

Here, $\alpha(z_t^\dagger) > 0$ corresponds to progressive sharpening, since the update direction $-T(z_t^\dagger)$ increases ϕ; $\beta(z_t^\dagger) > 0$ gives the self-stabilization property, while normality is assumed for the analysis.

6.1 A Simplified Four-Stage Self-Stabilization Mechanism

We first fix a reference point $z^\star := z_0^\dagger$ on the stability boundary, so that $\phi(z^\star) = 0$. For notational simplicity, we set

$$\begin{gathered} T := T(z^\star), \quad DT := DT(z^\star), \quad D^2T := D^2T(z^\star), \quad u := u(z^\star), \\ \Lambda := \Lambda(z^\star), \quad \nabla\Lambda := \nabla\Lambda(z^\star), \quad \nabla\phi := \nabla\phi(z^\star), \quad \alpha := \alpha(z^\star), \quad \beta := \beta(z^\star). \end{gathered}$$

As in the simplified derivation of Damian et al. (2023, Section 4.1), we further assume $T, \nabla\Lambda \in \ker(DT)$ and $\nabla\Lambda \perp u$.[4] In Section 6.2, we remove these simplifying assumptions and develop the general analysis.

We track the deviation of z_t along the unstable direction u and the movement of Λ by defining

$$x_t := \tfrac{1}{\|u\|} u^H (z_t - z^\star) \quad \text{and} \quad y_t := (\eta\Lambda - 1)\, \nabla\Lambda^H (z_t - z^\star).$$

Here, the factor $(\eta\Lambda - 1)$ in the definition of y_t ensures that its real and imaginary parts capture the normal and tangential movements of Λ relative to ∂D_η, respectively. In particular, $\mathrm{Re}(y_t) = \nabla\phi \cdot (z_t - z^\star)$ and $\mathrm{Im}(y_t) = -1/\eta \nabla \arg(\eta\Lambda - 1) \cdot (z_t - z^\star)$. The first-order approximations of both $\Lambda(z_t)$ and $\phi(z_t)$ can be represented by y_t as

$$\Lambda(z_t) \approx \Lambda + \nabla\Lambda^\top (z_t - z^\star) = \Lambda + (\eta\Lambda - 1)\bar{y}_t, \qquad \phi(z_t) \approx \mathrm{Re}(y_t).$$

Stage 1: Progressive Sharpening While both x_t and y_t are small, we have $T(z_t) \approx T$. Then

$$\mathrm{Re}(y_{t+1} - y_t) \approx -\eta \nabla\phi^\top T = \eta\alpha > 0,$$

so $\mathrm{Re}(y_t) \approx \phi(z_t)$ increases linearly.

[4] We provide an explicit example satisfying these assumptions in Appendix D.1.

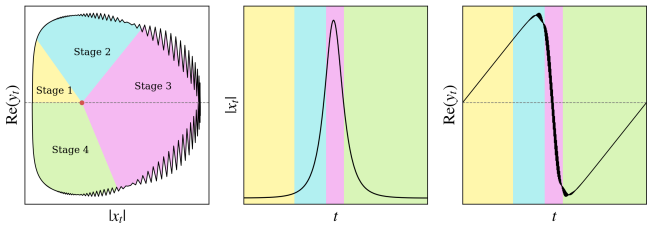

Figure 5: Self-stabilization dynamics of $|x_t|$ and $\mathrm{Re}(y_t)$ with the toy example in Appendix D.1.

Stage 2: Blowup The local dynamics of x_t satisfy

$$x_{t+1} = x_t - \frac{\eta}{\|u\|} u^H T(z_t) \approx (1 - \eta\Lambda(z_t))x_t \approx (1 - \eta\Lambda)(1 + \eta\bar{y}_t)x_t.$$

Since $|1 - \eta\Lambda| = 1$, we have $|x_{t+1}| \approx |1 + \eta y_t||x_t|$. As $|1 + \eta y_t| > 1$ approximately corresponds to $\phi(z_t) > 0$, once Stage 1 drives $\Lambda(z_t)$ outside D_η, $|x_t|$ begins to grow exponentially.

Stage 3: Self-Stabilization Once $|x_t|$ becomes sufficiently large, the approximation $T(z_t) \approx T$ is no longer sufficient, and a higher-order expansion is needed. When Λ is non-real, we use $z_t - z^\star \approx \frac{1}{\sqrt{2}}(x_t u + \overline{x_t u})$, under which the quadratic term is approximated as $\frac{1}{2}D^2T[z_t - z^\star, z_t - z^\star] \approx \frac{|x_t|^2}{2}D^2T[u, \bar{u}] + \frac{1}{2}\,\mathrm{Re}\left(x_t^2 D^2T[u, u]\right)$. Therefore, the dynamics of $\mathrm{Re}(y_t)$ is

$$\begin{aligned}\mathrm{Re}(y_{t+1} - y_t) = \eta\nabla\phi^\top(-T(z_t)) &\approx \eta\nabla\phi \cdot \left(-T - DT(z_t - z^\star) - \frac{1}{2}D^2T[z_t - z^\star, z_t - z^\star]\right) \\ &\approx \eta\left(\alpha - \frac{\beta}{2}|x_t|^2 - \frac{\gamma}{2}|x_t|^2\cos\left(2\arg(x_t) + \zeta\right)\right)\end{aligned}$$

for some constants γ and ζ independent of t. The γ-term oscillates with the rotating phase of x_t and is therefore expected to have a smaller cumulative effect than the non-oscillatory β-term. Therefore, as $|x_t|^2$ becomes sufficiently large, the positivity of β makes $\mathrm{Re}(y_t)$ decrease. For real Λ, the same dynamics hold with $\gamma = 0$.

Stage 4: Return to Stability Region The negative drift from Stage 3 eventually gives $|1+\eta y_t| < 1$. Hence $|x_t|$ begins to shrink, and the dynamics eventually return to Stage 1.

Together, the four stages are summarized by

$$\begin{aligned}x_{t+1} &\approx (1 - \eta\Lambda)(1 + \eta\bar{y}_t)x_t, \\ \mathrm{Re}(y_{t+1}) &\approx \mathrm{Re}(y_t) + \eta\left(\alpha - \frac{\beta}{2}|x_t|^2 - \frac{\gamma}{2}|x_t|^2\cos(2\arg(x_t) + \zeta)\right).\end{aligned} \tag{9}$$

We refer to these four stages as the self-stabilization mechanism of minimax methods; see Figure 5 for the dynamics. In the minimization setting, where $\Lambda = 2/\eta$ and $\gamma = 0$, equation 9 reduces to Damian et al. (2023, Equation 3). The oscillatory γ-term therefore captures a new rotational effect that does not arise in minimization.

While the updates along u, $\bar{u}$ and $\nabla\phi$ are self-stabilized, progress continues in the remaining directions. In Appendix D.4, Lemma D.10 shows that this remaining update coincides with the projected update $z_t^\dagger$ up to higher-order terms. The next section extends this result without the simplifying assumptions introduced at the beginning of this section.

6.2 Projected Dynamics Closely Approximate Minimax Dynamics During EoS

We now analyze the self-stabilization mechanism along the projected trajectory $\{z_t^\dagger\}$ without fixing a reference point.

Definition 6.2. At $z_t^\dagger$, write

$$\begin{aligned}&T_t := T(z_t^\dagger), \quad &&DT_t := DT(z_t^\dagger), \quad &&D^2T_t := D^2T(z_t^\dagger), \quad &&\Lambda_t := \Lambda(z_t^\dagger), \\ &u_t := u(z_t^\dagger), \quad &&\nabla\phi_t := \nabla\phi(z_t^\dagger), \quad &&\nabla\Lambda_t := \nabla\Lambda(z_t^\dagger), \quad &&\alpha_t := \alpha(z_t^\dagger).\end{aligned}$$

Let P_t be the orthogonal projector onto $U_t := \mathrm{span}_{\mathbb{R}}\{\mathrm{Re}\, u_t, \mathrm{Im}\, u_t\}$, and write $a^\perp := (I - P_t)a$. Since U_t and $\nabla\Lambda_t$ are not orthogonal in general, we define $\beta_t^\perp := \nabla\phi_t^\perp \cdot D^2T_t[u_t, \bar{u}_t]$. We also define the quantities used in our main result as $\delta_t := \sqrt{\frac{2\alpha_t}{\beta_t^\perp}}$, $\delta := \sup_t \delta_t$, $\epsilon_t := \eta\sqrt{\alpha_t}$, and $\epsilon := \sup_t \epsilon_t$.

As in Section 6.1, we track deviations along u_t and $\nabla\Lambda_t$. We introduce the predicted dynamics $v_t^\star$ (see Definition D.2) intended to approximate $z_t - z_t^\dagger$, and define $x_t^\star := \frac{1}{\|u_t\|}u_t^H v_t^\star$ and $y_t^\star := (\eta\Lambda_t - 1)(\nabla\Lambda_t^\perp)^H v_t^\star$. They exhibit an analogous self-stabilization mechanism; see Lemma D.5.

We impose the following assumptions for the proof, which play the same roles as their counterparts in Damian et al. (2023).

Assumption 6.3. For every t, we assume that

1. For some fixed $\Delta > 0$, $\min_{\lambda\in\{\Lambda_t,\bar{\Lambda}_t\}} \mathrm{dist}(\lambda, \mathrm{eig}(DT_t) \setminus \{\lambda\}) > \frac{\Delta}{\eta}$.
2. Let ρ_2, ρ_3 be the minimum constants such that $\|D^2T(z)\|_{op} \le \rho_2$ for all z and D^2T is ρ_3-Lipschitz with respect to $\|\cdot\|_{op}$. Assume $\rho_3 = O(\eta\rho_2^2)$.
3. $\frac{\alpha_t}{\|T_t\|\|\nabla\phi_t^\perp\|}, \frac{\beta_t^\perp}{\|\nabla\phi_t^\perp\|^2}, \frac{\|\nabla\phi_t^\perp\|}{\rho_2} = \Theta(1)$.
4. $\frac{\|D^2T_t[v,w]\|}{\rho_2\|v\|\|w\|}, \frac{|\langle v_t^{\star,\perp}, DT_t v_t^{\star,\perp}\rangle|}{\|DT_t\|\|v_t^{\star,\perp}\|^2}, \frac{\mathrm{dist}(\lambda, D_\eta)}{\|DT_t\|} \le O(\epsilon) \quad \forall v, w \in U_t^\perp, \lambda \in \mathrm{eig}(DT_t) \setminus \{\Lambda_t, \bar{\Lambda}_t\}$.

In the minimax setting, additional assumptions are required to obtain the same type of approximation result. A detailed discussion of their roles is provided in Appendix D.2.

Assumption 6.4. For every t, we assume that

1. $\left\|D\left(DT(DT)^\top - (DT)^\top DT\right)(z_t^\dagger)\right\|_{op} = O\left(\frac{\epsilon\rho_2}{\eta}\right)$.
2. $\|P_{U_t,\nabla\phi_t}^\perp D^2T_t[u_t, \bar{u}_t]\| = O(\epsilon\rho_2)$, where $P_{U_t,\nabla\phi_t}^\perp$ is the projector onto $\mathrm{span}_{\mathbb{R}}\{U_t, \nabla\phi_t\}^\perp$.

Theorem 6.5. *Let $\mathcal{T} = O(\epsilon^{-1})$. Assume Assumptions 6.1–6.4, $P_0^\perp(z_0 - z_0^\dagger) = 0$, $\|z_0 - z_0^\dagger\| = O(\delta)$, and ϵ is sufficiently small*[5]*. For $t \le \mathcal{T}$, we have*

$$z_t = z_t^\dagger + v_t^\star + O(\epsilon\delta), \tag{10}$$

$$\Lambda(z_t) = \Lambda_t - (1 - \eta\Lambda_t)\bar{y}_t^\star + \frac{\nabla\Lambda_t^\top u_t}{\|u_t\|}x_t^\star + O(\epsilon^2/\eta). \tag{11}$$

This provides a natural extension of Damian et al. (2023, Theorem 1) to minimax optimization. It shows that the minimax trajectory closely follows the projected dynamics, while the dominant eigenvalue remains near the stability boundary through self-stabilization.

7 Conclusion

In this paper, we generalized the notion of Edge of Stability (EoS) to minimax optimization and empirically showed that both Extragradient (EG) and Optimistic Gradient Descent Ascent (OGDA) exhibit EoS in three two-player zero-sum games. A notable distinction appears when the dominant eigenvalue of the saddle-gradient Jacobian lies near the real axis: OGDA continues to exhibit EoS, whereas EG becomes unstable. We extended the self-stabilization property to minimax methods and derived a self-stabilization mechanism that explains both the emergence of EoS and its absence for EG in the real-axis regime.

Limitation and Future Work Our analysis relies on normality assumptions on DT, and some conditions required by Theorem 6.5 are only approximately satisfied in our experiments. Relaxing these assumptions and validating the self-stabilization mechanism under more realistic conditions are important directions for future work. A more complete characterization of self-stabilization property also remains open. Another direction is to study EoS in a broader range of minimax problems.

[5]Although the small-ϵ assumption is not stated explicitly in Damian et al. (2023), it is also required in their proof.

AI USE STATEMENT

In this work, we used OpenAI Astra to assist with the derivation of Lemma C.1 and Lemma D.13. We independently verified the resulting proofs and take responsibility for their correctness. We also used Astra to assist with translation, clarity editing, and LaTeX formatting. We reviewed all AI-assisted content and take responsibility for the final content of this work.

REPRODUCIBILITY STATEMENT

Detailed experimental settings, including problem formulations, model architectures, optimization settings, evaluation metrics, and hyperparameters, are provided in Appendix B. The toy example illustrating the self-stabilization mechanism is described in Appendix D.1. The assumptions underlying our theoretical results are stated in the main text, with proofs and supporting derivations provided in Appendices A, C, and D.

REFERENCES

Kwangjun Ahn, Jingzhao Zhang, and Suvrit Sra. Understanding the unstable convergence of gradient descent. In *International conference on machine learning*, pp. 247–257. PMLR, 2022.

Ahmet Alacaoglu, Luca Viano, Niao He, and Volkan Cevher. A natural actor-critic framework for zero-sum markov games. In *International Conference on Machine Learning*, pp. 307–366. PMLR, 2022.

Sanjeev Arora, Zhiyuan Li, and Abhishek Panigrahi. Understanding gradient descent on the edge of stability in deep learning. In *International Conference on Machine Learning*, pp. 948–1024. PMLR, 2022.

Axel Böhm. Solving nonconvex-nonconcave min-max problems exhibiting weak minty solutions. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=Gp0pHyUyrb.

Jiseok Chae, Kyuwon Kim, and Donghwan Kim. Two-timescale extragradient for finding local minimax points. In *International Conference on Learning Representations*, volume 2024, pp. 35905–35944, 2024.

Jeremy Cohen, Simran Kaur, Yuanzhi Li, J Zico Kolter, and Ameet Talwalkar. Gradient descent on neural networks typically occurs at the edge of stability. In *International Conference on Learning Representations*, 2021. URL https://openreview.net/forum?id=jh-rTtvkGeM.

Jeremy Cohen, Alex Damian, Ameet Talwalkar, J Zico Kolter, and Jason D. Lee. Understanding optimization in deep learning with central flows. In *The Thirteenth International Conference on Learning Representations*, 2025. URL https://openreview.net/forum?id=sIE2rI3ZPs.

Alex Damian, Eshaan Nichani, and Jason D. Lee. Self-stabilization: The implicit bias of gradient descent at the edge of stability. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=nhKHA59gXz.

Jelena Diakonikolas, Constantinos Daskalakis, and Michael I Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2746–2754. PMLR, 2021.

Jiaojiao Fan, Shu Liu, Shaojun Ma, Hao-Min Zhou, and Yongxin Chen. Neural Monge map estimation and its applications. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL https://openreview.net/forum?id=2mZSlQscj3. Featured Certification.

Yi Feng, Andrea Paudice, and Stratis Skoulakis. Adversarial training with large step sizes: Implicit bias and evolution of sharpness. In *New Frontiers in Game-Theoretic Learning - NExT-Game*, 2026. URL https://openreview.net/forum?id=HLq4IwkUJv.

Burak M Gonultas and Volkan Isler. Pursuit-evasion for car-like robots with sensor constraints. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3800–3807. IEEE, 2025.

Eduard Gorbunov, Adrien Taylor, Samuel Horváth, and Gauthier Gidel. Convergence of proximal point and extragradient-based methods beyond monotonicity: the case of negative comonotonicity. In *International Conference on Machine Learning*, pp. 11614–11641. PMLR, 2023.

Xinmeng Huang, Ernest K Ryu, and Wotao Yin. Scaled relative graph of normal matrices. *Journal of Convex Analysis*, 33(1&2):145–154, 2026.

Rufus Isaacs. *Games of Pursuit*. RAND Corporation, Santa Monica, CA, 1951.

Alexander Korotin, Daniil Selikhanovych, and Evgeny Burnaev. Neural optimal transport. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.net/forum?id=d8CBRlWNkqH.

Galina M Korpelevich. Extragradient method for finding saddle points and other problems. *Matekon*, 13(4):35–49, 1977.

Michael L Littman. Markov games as a framework for multi-agent reinforcement learning. In *Machine learning proceedings 1994*, pp. 157–163. Elsevier, 1994.

Liming Liu, Zixuan Zhang, Simon Shaolei Du, and Tuo Zhao. A minimalist example of edge-of-stability and progressive sharpening. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=yst8MHfcgP.

Philip M Long and Peter L Bartlett. Sharpness-aware minimization and the edge of stability. *Journal of Machine Learning Research*, 25(179):1–20, 2024.

Kaifeng Lyu, Zhiyuan Li, and Sanjeev Arora. Understanding the generalization benefit of normalization layers: Sharpness reduction. *Advances in Neural Information Processing Systems*, 35: 34689–34708, 2022.

Chao Ma, Lei Wu, and Lexing Ying. The multiscale structure of neural network loss functions: The effect on optimization and origin. *arXiv preprint arXiv:2204.11326*, 14, 2022.

Lars Mescheder, Sebastian Nowozin, and Andreas Geiger. Which training methods for GANs do actually converge? In *International Conference on Machine Learning (ICML)*, 2018.

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

Yurii Nesterov. *Lectures on Convex Optimization*, volume 137 of *Springer Optimization and Its Applications*. Springer, Cham, 2 edition, 2018. ISBN 978-3-319-91578-4. doi: 10.1007/978-3-319-91578-4.

Leonid Denisovich Popov. A modification of the Arrow-Hurwicz method for search of saddle points. *Mathematical notes of the Academy of Sciences of the USSR*, 28(5):845–848, 1980.

Litu Rout, Alexander Korotin, and Evgeny Burnaev. Generative modeling with optimal transport maps. In *International Conference on Learning Representations*, 2022. URL https://openreview.net/forum?id=5JdLZg346Lw.

Ernest K Ryu, Robert Hannah, and Wotao Yin. Scaled relative graphs: Nonexpansive operators via 2d euclidean geometry. *Mathematical Programming*, 194(1):569–619, 2022.

Lloyd S Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10): 1095–1100, 1953.

Zixuan Wang, Zhouzi Li, and Jian Li. Analyzing sharpness along GD trajectory: Progressive sharpening and edge of stability. *Advances in Neural Information Processing Systems*, 35:9983–9994, 2022.

Sihan Xu and Zhaohui Dang. Emergent behaviors in multiagent pursuit evasion games within a bounded 2d grid world. *Scientific Reports*, 15(1):29376, 2025.

Geonhui Yoo, Minhak Song, and Chulhee Yun. Understanding sharpness dynamics in NN training with a minimalist example: The effects of dataset difficulty, depth, stochasticity, and more. In *Forty-second International Conference on Machine Learning*, 2025. URL https://openreview.net/forum?id=XfjrLEPOQV.

Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *International Conference on Learning Representations*, 2020. URL https://openreview.net/forum?id=BJgnXpVYwS.

Yulai Zhao, Yuandong Tian, Jason Lee, and Simon Du. Provably efficient policy optimization for two-player zero-sum markov games. In *International Conference on Artificial Intelligence and Statistics*, pp. 2736–2761. PMLR, 2022.

CONTENTS

1 Introduction **1**

1.1 Notations 2

2 Preliminaries **2**

3 Generalizing EoS: From Minimization to Minimax Optimization **3**

3.1 Nonexpansiveness as a Stability Criterion 3

3.2 Geometric Representation of Stability Regions via SRG 3

3.3 Progressive Sharpening and Edge of Stability in Minimax Optimization 4

4 Edge of Stability Occurs in Minimax Optimization **4**

4.1 EoS Occurs Across Two-Player Zero-Sum Game Settings 4

4.2 Near the Real Axis: OGDA Exhibits EoS, While EG Is Unstable 5

5 Self-Stabilization Property of Minimax Methods **6**

5.1 Generalizing Self-Stabilization Property 6

5.2 When Does the Self-Stabilization Property Hold? 6

6 Self-Stabilization Mechanism of Minimax Methods **7**

6.1 A Simplified Four-Stage Self-Stabilization Mechanism 7

6.2 Projected Dynamics Closely Approximate Minimax Dynamics During EoS 8

7 Conclusion **9**

A Proofs and Missing Details in Section 3 **15**

A.1 Stability Region: Propositions and Proofs 15

A.2 Derivation of Stability Region for EG and OGDA 17

B Experiments Details of Section 4 **19**

B.1 Two-player Zero-Sum Markov Games 19

B.1.1 Two-stage Rock Paper Scissors Game 20

B.1.2 Pursuit–Evasion Games 20

B.2 Neural Optimal Transport 24

B.3 Simple Minimization Problem 24

C Proofs and Missing Details in Section 5 **26**

C.1 Proof of Proposition 5.1 26

C.2 Verification of the conditions in Proposition 5.1. 28

D Proofs and Missing Details in Section 6 **30**

D.1 Toy-Example for Self-Stabilization Mechanism 30

D.2 Discussion of the Assumptions 30
D.3 Predicted Dynamics 32
D.4 Properties of the constrained trajectory 34
D.5 Proofs of Theorem 41

A Proofs and Missing Details in Section 3

A.1 Stability Region: Propositions and Proofs

We first prove the following Nonexpansive Lemma of given update operator $I - \eta T$. The proof follows almost directly from the definition of Scaled Relative Graph (SGR).

Proposition A.1 (Nonexpansive Lemma)**.** *For some $\ell > 0$,*

$$\mathcal{G}(T) \subset D_\ell := \left\{ \omega \in \mathbb{C} : \left| \omega - \frac{1}{\ell} \right| \leq \frac{1}{\ell} \right\}$$

is equivalent to

$$\|(I - \eta T)(z) - (I - \eta T)(w)\|^2 \leq \|z - w\|^2 - \eta(\ell - \eta)\|T(z) - T(w)\|^2, \quad \forall z, w. \tag{12}$$

for any $\eta > 0$.

Proof. Let us fix z, w, and let

$$\Delta z := z - w, \qquad \Delta T := T(z) - T(w).$$

The claim is immediate if $\Delta z = 0$. Assume that $\Delta z \neq 0$. Let $\omega \in \mathcal{G}(T)$ be an SRG point associated with the pair (z, w). By the definition of the SRG,

$$|\omega| = \frac{\|\Delta T\|}{\|\Delta z\|}, \qquad \operatorname{Re} \omega = \frac{\langle \Delta T, \Delta z \rangle}{\|\Delta z\|^2}. \tag{13}$$

The assumed disk containment gives

$$\left| \omega - \frac{1}{\ell} \right|^2 \leq \frac{1}{\ell^2}.$$

Expanding this inequality yields

$$|\omega|^2 - \frac{2}{\ell} \operatorname{Re} \omega \leq 0,$$

or equivalently,

$$2 \operatorname{Re} \omega \geq \ell |\omega|^2.$$

Multiplying by $\|\Delta z\|^2$ and using the SRG identities equation 13, we obtain

$$2\langle \Delta T, \Delta z \rangle \geq \ell \|\Delta T\|^2.$$

Therefore,

$$\begin{aligned} \|(I - \eta T)(z) - (I - \eta T)(w)\|^2 &= \|\Delta z - \eta \Delta T\|^2 \\ &= \|\Delta z\|^2 - 2\eta \langle \Delta z, \Delta T \rangle + \eta^2 \|\Delta T\|^2 \\ &\leq \|\Delta z\|^2 - \eta \ell \|\Delta T\|^2 + \eta^2 \|\Delta T\|^2 \\ &= \|\Delta z\|^2 - \eta(\ell - \eta)\|\Delta T\|^2. \end{aligned}$$

The converse follows by reversing the preceding argument; the details are left to the reader. □

This implies that if $\mathcal{G}(T) \subset D_\ell$ with $\eta \leq \ell$, then the update operator $I - \eta T$ is nonexpansive. Since $D_\ell \subset D_\eta$ for $\eta \leq \ell$, the condition $\mathcal{G}(T) \subset D_\eta$ characterizes the stability region for nonexpansiveness of $I - \eta T$.

The SRG of T is closely related to the Jacobian matrix DT of the update operator T. While previous works (Ryu et al., 2022, Theorem 1), (Huang et al., 2026, Theorem 4.3) established the relationship for linear operators, we extend their results to non-linear operators.

Proposition A.2. *Let $D_\eta = \{\omega : |1 - \eta\omega| \leq 1\}$ be a disk. If T is differentiable, then*

$$\mathcal{G}(T) \subset D_\eta \implies \operatorname{eig}(DF(z)) \subset D_\eta, \quad \forall z. \tag{14}$$

Moreover, if $DF(z)$ is a normal matrix for all z, then

$$\mathcal{G}(T) \subset D_\eta \iff \operatorname{eig}(DF(z)) \subset D_\eta, \quad \forall z. \tag{15}$$

Proof. From the proof of Proposition 5, we have $\omega \in D_\ell \iff 2\operatorname{Re}\omega \geq \ell|\omega|^2$. By the SRG identities equation 13, for a pair z, w with $z \neq w$, the corresponding SRG point ω satisfies

$$|\omega| = \frac{\|T(z) - T(w)\|}{\|z - w\|}, \qquad \operatorname{Re}\omega = \frac{\langle T(z) - T(w), z - w\rangle}{\|z - w\|^2}.$$

Consequently,

$$\mathcal{G}(T) \subset D_\ell \iff 2\langle T(z) - T(w), z - w\rangle \geq \ell\|T(z) - T(w)\|^2, \qquad \forall z, w. \tag{16}$$

We first prove the first statement equation 14 that does not require normality. Suppose that $\mathcal{G}(T) \subset D_\ell$. Fix z and $u \in \mathbb{R}^d$. Applying equation 16 to $z + tu$ and z, dividing by t^2, and letting $t \downarrow 0$ gives

$$2\langle DT(z)u, u\rangle \geq \ell\|DT(z)u\|^2.$$

Writing $A := DT(z)$, this is equivalent to

$$A + A^\top - \ell A^\top A \succeq 0. \tag{17}$$

Let $\lambda \in \operatorname{eig}(A)$, and let $q \in \mathbb{C}^d \setminus \{0\}$ be a corresponding eigenvector, so that $Aq = \lambda q$. Since the matrix in equation 17 is real, symmetric, and positive semidefinite, its quadratic form is nonnegative also on complex vectors. Hence

$$0 \leq q^H \left(A + A^\top - \ell A^\top A\right) q = \left(2\operatorname{Re}\lambda - \ell|\lambda|^2\right) \|q\|^2.$$

It follows that

$$2\operatorname{Re}\lambda \geq \ell|\lambda|^2,$$

and therefore $\lambda \in D_\ell$. Since z was arbitrary,

$$\operatorname{eig}\big(DT(z)\big) \subset D_\ell, \qquad \forall z.$$

We now prove the converse under the normality assumption. Suppose that $A(z) := DT(z)$ is normal and

$$\operatorname{eig}(A(z)) \subset D_\ell, \qquad \forall z.$$

Since $A(z)$ is normal, it is unitarily diagonalizable over $\mathbb{C}$. Thus, for each z, there exist a unitary matrix U and a diagonal matrix

$$\Lambda = \operatorname{diag}(\lambda_1, \ldots, \lambda_d)$$

such that

$$A(z) = U\Lambda U^H.$$

Therefore,

$$\begin{aligned} A(z) + A(z)^\top - \ell A(z)^\top A(z) &= U\left(\Lambda + \Lambda^H - \ell\Lambda^H\Lambda\right)U^H \\ &= U \operatorname{diag}\left(2\operatorname{Re}\lambda_i - \ell|\lambda_i|^2\right)_{i=1}^d U^H. \end{aligned}$$

Since $\lambda_i \in D_\ell$ for every i,

$$2\operatorname{Re}\lambda_i - \ell|\lambda_i|^2 \geq 0.$$

Hence

$$A(z) + A(z)^\top - \ell A(z)^\top A(z) \succeq 0, \qquad \forall z.$$

Equivalently,

$$2\langle A(z)u, u\rangle \geq \ell\|A(z)u\|^2, \qquad \forall z, \quad \forall u \in \mathbb{R}^d. \tag{18}$$

Fix z, w, let $d := z - w$, and define

$$z_t := w + td, \qquad t \in [0, 1].$$

By the fundamental theorem of calculus,

$$T(z) - T(w) = \int_0^1 DT(z_t)d\, dt.$$

Using equation 18, we obtain

$$\begin{aligned} 2\langle T(z) - T(w), d\rangle &= 2\int_0^1 \langle DT(z_t)d, d\rangle \, dt \\ &\geq \ell \int_0^1 \|DT(z_t)d\|^2 \, dt \\ &\geq \ell \left\| \int_0^1 DT(z_t)d \, dt \right\|^2 \\ &= \ell \|T(z) - T(w)\|^2, \end{aligned}$$

where the second inequality follows from Jensen's inequality. Hence equation 16 holds, and therefore

$$\mathcal{G}(T) \subset D_\ell.$$

□

A.2 Derivation of Stability Region for EG and OGDA

Extragradient The EG iteration and its update operator are

$$z_{t+1} = z_t - \eta\, T_{\mathrm{eg}}(z), \qquad\qquad T_{\mathrm{eg}}(z) := F(z - \eta F(z))\,.$$

The chain rule gives the exact Jacobian

$$DT_{\mathrm{eg}}(z) = DF(z - \eta F(z))\,(I - \eta DF(z))\,.$$

In general, the eigenvalues of $DT_{\mathrm{eg}}(z)$ cannot be determined solely from those of $DF(z)$. However, assuming that the Jacobian changes only slightly over one update, we assume that $DF(z-\eta F(z)) = DF(z)$. It follows that,

$$\begin{aligned} \mathrm{eig}(DT_{\mathrm{eg}}(z)) \subset D_\eta &\iff \mathrm{eig}\big(DF(z) - \eta(DF(z))^2\big) \subset D_\eta \\ &\iff \mathrm{eig}(DF(z)) \subset \mathcal{P}_\eta, \end{aligned}$$

where

$$\mathcal{P}_\eta := \left\{\omega \in \mathbb{C} : \omega - \eta\omega^2 \in D_\eta\right\}.$$

It has a peanut-shaped stability region; see Figure 1. This region was first introduced by Chae et al. (2024), who analyze the convergence of a variant of EG in the nonconvex–nonconcave setting.

If an eigenvalue of DT_{eg} lies on ∂D_η, the corresponding eigenvalues of DF also lie on $\partial\mathcal{P}_\eta$. We highlight several key correspondences between the spectra of DT_{eg} and DF. The left-hand side denotes an eigenvalue of DT_{eg}, and the right-hand side the corresponding eigenvalue of DF:

$$0 \longleftrightarrow \left\{0, \frac{1}{\eta}\right\}, \qquad \frac{2}{\eta} \longleftrightarrow \left\{\frac{1+i\sqrt{7}}{2\eta}, \frac{1-i\sqrt{7}}{2\eta}\right\},$$
$$\frac{1+i}{\eta} \longleftrightarrow \left\{\frac{i}{\eta}, \frac{1-i}{\eta}\right\}, \qquad \frac{1-i}{\eta} \longleftrightarrow \left\{-\frac{i}{\eta}, \frac{1+i}{\eta}\right\}.$$

Note that the eigenvalue 0 of DT_{eg} corresponds to two eigenvalues of DF, namely 0 and $1/\eta$. In particular, the latter correspondence underlies the unstable EG dynamics observed in Section 4.2.

Optimistic gradient descent ascent Let $Z_t = (z_t, z_{t-1})^\top$. Then OGDA can be written exactly as

$$Z_{t+1} = Z_t - \eta T_{\mathrm{ogda}}(Z_t), \qquad\qquad T_{\mathrm{ogda}}(z, \tilde{z}) := \begin{pmatrix} 2F(z) - F(\tilde{z}) \\ \eta^{-1}(\tilde{z} - z) \end{pmatrix}.$$

At $(z, \tilde{z})$, the exact Jacobian is

$$DT_{\mathrm{ogda}}(z, \tilde{z}) = \begin{pmatrix} 2DF(z) & -DF(\tilde{z}) \\ -\eta^{-1} I & \eta^{-1} I \end{pmatrix}.$$

In general, the eigenvalues of $DT_{\mathrm{ogda}}(z, \tilde{z})$ cannot be exactly determined from those of $DF(z)$. Under the similar assumption $DF(z_{t+1}) = DF(z_t)$ to the EG case, let $(u^\top, v^\top)^\top$ be an eigenvector of the block Jacobian with eigenvalue μ. The eigenvalue equation is given by

$$\begin{pmatrix} 2DF(z_t) & -DF(z_t) \\ -\eta^{-1}I & \eta^{-1}I \end{pmatrix} \begin{pmatrix} u \\ v \end{pmatrix} = \mu \begin{pmatrix} u \\ v \end{pmatrix}.$$

This yields a system of two equations:

$$2DF(z_t)u - DF(z_t)v = \mu u, \tag{19}$$

$$-\eta^{-1}u + \eta^{-1}v = \mu v. \tag{20}$$

From equation 20, we directly obtain the relation $u = (1 - \eta\mu)v$. Substituting this into equation 19 gives

$$DF(z_t)\big[2(1 - \eta\mu) - 1\big]v = \mu(1 - \eta\mu)v \quad \Longrightarrow \quad DF(z_t)v = \frac{\mu(1 - \eta\mu)}{1 - 2\eta\mu}v.$$

This implies that v must be an eigenvector of $DF(z_t)$. Letting $\lambda \in \mathrm{eig}(DF(z_t))$ be the corresponding eigenvalue, we have $\lambda = \frac{\mu(1-\eta\mu)}{1-2\eta\mu}$. Rearranging this yields the characteristic quadratic equation for μ:

$$\mu^2 - \left(2\lambda + \frac{1}{\eta}\right)\mu + \frac{\lambda}{\eta} = 0,$$

which gives the two eigenvalue branches generated by λ:

$$\mu_\pm(\lambda) = \lambda + \frac{1}{2\eta} \pm \sqrt{\lambda^2 + \frac{1}{4\eta^2}}.$$

Therefore, the spectrum of the OGDA update operator is obtained by applying the two branches $\mu_\pm$ to the eigenvalues of $DF(z_t)$. Hence,

$$\begin{aligned} \mathrm{eig}(DT_{\mathrm{ogda}}(Z_t)) \subset D_\eta &\iff \left\{\lambda + \frac{1}{2\eta} \pm \sqrt{\lambda^2 + \frac{1}{4\eta^2}} : \lambda \in \mathrm{eig}(DF(z_t))\right\} \subset D_\eta \\ &\iff \mathrm{eig}(DF(z_t)) \subset \mathcal{B}_\eta, \end{aligned}$$

where

$$\mathcal{B}_\eta := \left\{\omega \in \mathbb{C} : \omega + \frac{1}{2\eta} \pm \sqrt{\omega^2 + \frac{1}{4\eta^2}} \in D_\eta\right\}.$$

It has a bean-shaped stability region; see Figure 1.

If an eigenvalue of DT_{ogda} lies on ∂D_η, the corresponding eigenvalues of DF also lie on $\partial\mathcal{B}_\eta$. We highlight several key correspondences between the spectra of DT_{ogda} and DF. The left-hand side denotes an eigenvalue of DT_{ogda}, and the right-hand side the corresponding eigenvalue of DF:

$$0 \longleftrightarrow 0, \qquad \frac{2}{\eta} \longleftrightarrow \frac{2}{3\eta},$$

$$\frac{1 + i\sqrt{3}}{2\eta} \longleftrightarrow \frac{i}{\sqrt{3}\eta}, \qquad \frac{1 - i\sqrt{3}}{2\eta} \longleftrightarrow -\frac{i}{\sqrt{3}\eta}.$$

Unlike EG, the eigenvalue 0 of DT_{ogda} corresponds only to the eigenvalue 0 of DF.

B EXPERIMENTS DETAILS OF SECTION 4

B.1 TWO-PLAYER ZERO-SUM MARKOV GAMES

We consider a two-player zero-sum Markov game (Shapley, 1953; Littman, 1994), which consists of two competing players x and y, a state space $\mathcal{S}$, action spaces $\mathcal{A}_x$ and $\mathcal{A}_y$, a transition kernel $P(s'|s, a_x, a_y)$, and a reward function $R(s, a_x, a_y)$. The two players use policies $\pi_x(\cdot|s)$ and $\pi_y(\cdot|s)$, where player x minimizes and player y maximizes the same objective. Given an initial distribution ρ, discount factor $\gamma \in (0, 1)$, and entropy regularization parameter $\tau \geq 0$, the objective is

$$\min_{\pi_x} \max_{\pi_y} J_\tau(\pi_x, \pi_y) \quad \text{where}$$
$$J_\tau(\pi_x, \pi_y) := \mathbb{E}\left[\sum_{t=0}^{\zeta-1} \gamma^t \left\{R(s_t, a_{x,t}, a_{y,t}) - \tau H(\pi_x(\cdot|s_t)) + \tau H(\pi_y(\cdot|s_t))\right\}\right], \tag{21}$$

where $s_0 \sim \rho$, $a_{x,t} \sim \pi_x(\cdot|s_t)$, $a_{y,t} \sim \pi_y(\cdot|s_t)$, $s_{t+1} \sim P(\cdot|s_t, a_{x,t}, a_{y,t})$, ζ is the termination time, and $H(p) := -\sum_a p_a \log p_a$. The case $\tau = 0$ recovers the unregularized game.

A pair of policies $(\pi_x^\star, \pi_y^\star)$ is a Nash equilibrium if

$$J_\tau(\pi_x^\star, \pi_y) \leq J_\tau(\pi_x^\star, \pi_y^\star) \leq J_\tau(\pi_x, \pi_y^\star)$$

for every π_x and π_y. When the state and action spaces are finite, a Nash equilibrium exists for the two-player zero-sum Markov game (Shapley, 1953).

Neural network parameterization When the action spaces are finite, the policy of each player can be represented by action probability vectors $p_x(s) \in \Delta_{|\mathcal{A}_x|}$ and $p_y(s) \in \Delta_{|\mathcal{A}_y|}$ at state s. We parameterize these vectors by neural networks with parameters x and y using softmax activation:

$$p_x(s) := \mathrm{softmax}(N_x(s)), \qquad p_y(s) := \mathrm{softmax}(N_y(s)), \tag{22}$$

where $\mathrm{softmax}(v) := \frac{1}{\sum_j e^{v_j}}(e^{v_1}, e^{v_2}, \ldots, e^{v_d}) \in \Delta_d$. This softmax activation is widely used for training two-player zero-sum Markov games (Zhao et al., 2022; Alacaoglu et al., 2022).

Evaluation To evaluate policy learning, a standard metric is the *duality gap*,

$$\mathrm{Gap}(\pi_x, \pi_y) := \max_{\widetilde{\pi}_y} J_\tau(\pi_x, \widetilde{\pi}_y) - \min_{\widetilde{\pi}_x} J_\tau(\widetilde{\pi}_x, \pi_y).$$

The duality gap is nonnegative and equals zero at a Nash equilibrium. However, it does not directly measure the distance to the Nash equilibrium, which is more closely related to the Proposition A.1 studied in this paper. To measure the distance between policies themselves, we introduce the following distance between two probability vectors.

Definition B.1 (Logit Distance)**.** For $p, q \in \mathrm{int}(\Delta_d)$[6], define

$$d_{\mathrm{logit}}(p, q) := \|r(p) - r(q)\|, \qquad r(p) := \log p - \frac{1}{d}\mathbf{1}\mathbf{1}^\top \log p.$$

Proposition B.2. *The logit distance d_{logit} is a metric on $\mathrm{int}(\Delta_d)$.*

Proof. Nonnegativity and symmetry follow directly from the Euclidean norm, and the triangle inequality follows from

$$\|r(p) - r(q)\| \leq \|r(p) - r(s)\| + \|r(s) - r(q)\|.$$

It remains to verify identity. If $d_{\mathrm{logit}}(p, q) = 0$, then $r(p) = r(q)$, and hence

$$\left(I - \frac{1}{d}\mathbf{1}\mathbf{1}^\top\right)(\log p - \log q) = 0.$$

Since $\ker\left(I - \frac{1}{d}\mathbf{1}\mathbf{1}^\top\right) = \mathrm{span}\{\mathbf{1}\}$, there exists $c \in \mathbb{R}$ such that $\log p - \log q = c\mathbf{1}$. Hence $p = e^c q$. Since $\mathbf{1}^\top p = \mathbf{1}^\top q = 1$, we have $e^c = 1$, and therefore $p = q$. □

[6]The logit distance is not finite on the boundary of the simplex. For numerical evaluation, if a policy lies on or numerically approaches the boundary, we clip its probabilities from below by a small constant $\varepsilon_{\mathrm{clip}} > 0$.

Using this distance, we measure the error of the learned policies from a reference Nash equilibrium.

Definition B.3 (Logit Error)**.** Let $(\pi_x^\star, \pi_y^\star)$ be a Nash equilibrium. For any policy pair (π_x, π_y), define the *logit error* as

$$\mathcal{E}_{\text{logit}} := \frac{1}{2} \sum_{i \in \{x,y\}} \mathbb{E}_{s\sim\rho} \left[d_{\text{logit}}\big(\pi_i(\cdot|s), \pi_i^\star(\cdot|s)\big) \right].$$

B.1.1 Two-stage Rock Paper Scissors Game

Rule. We consider a two-stage Rock-Paper-Scissors (RPS) game, which is a variant of the standard RPS game. In the first stage, each player independently draws one of R, P, and S for each of their two hands. In the second stage, after observing all four hands, the players simultaneously select either their left or right hand for the final RPS match. As in standard RPS, rock beats scissors, scissors beats paper, and paper beats rock; a win, draw, or loss gives reward $+1$, 0, or -1, respectively.

Environment. The four hands drawn in the first stage define the state

$$s = (h_{x,L}, h_{x,R}, h_{y,L}, h_{y,R}) \in \mathcal{S} := \{\mathrm{R}, \mathrm{P}, \mathrm{S}\}^4,$$

and the first-stage sampling induces the state distribution ρ. The action spaces are

$$\mathcal{A}_x = \mathcal{A}_y = \{L, R\},$$

corresponding to the choice of the left or right hand. For $a_x, a_y \in \{L, R\}$, the reward $R(s, a_x, a_y) \in \{-1, 0, 1\}$ is determined by the RPS outcome between the two selected hands. The game terminates immediately after this action, so the terminate time $\zeta = 1$. Consequently, the discount factor γ and transition kernel P has no effect on the objective. We use no entropy regularization, i.e., $\tau = 0$.

Policy networks We encode the four hands by concatenating their one-hot representations, giving a 12-dimensional input. Each player uses a separate network as in equation 22, with three ReLU hidden layers of width 12 and a two-dimensional linear output layer. The network architectures are summarized in Table 1.

Nash equilibrium Since $\zeta = 1$, the Nash equilibrium policy $\pi^\star(\cdot \mid s)$ can be computed independently for each state $s = (h_{x,L}, h_{x,R}, h_{y,L}, h_{y,R})$ as the equilibrium of a 2×2 matrix game with payoff matrix $[R(s, a_x, a_y)]_{a_x, a_y}$. If there is no pure-strategy equilibrium, i.e., $\pi_i^\star(\cdot \mid s) \notin \{0, 1\}$ for $i \in \{x, y\}$, then $\pi_i^\star(\cdot \mid s)$ can be computed as

$$\pi_x^\star(L \mid s) = \frac{R(s,R,R) - R(s,R,L)}{D(s)}, \qquad \pi_y^\star(L \mid s) = \frac{R(s,R,R) - R(s,L,R)}{D(s)},$$

where

$$D(s) = R(s,L,L) - R(s,L,R) - R(s,R,L) + R(s,R,R).$$

The right-hand probabilities are their complements.

Why standard RPS is insufficient Standard RPS is too simple to exhibit EoS: training typically converges to the Nash equilibrium before the dominant eigenvalue reaches the stability boundary. In minimization, PS has been observed more clearly in deeper neural networks (Yoo et al., 2025), motivating us to use multiple hidden layers. However, in standard RPS, the policies π_x and π_y are independent of input state s, so the network can approach the Nash equilibrium largely through its last layer without substantially changing the hidden layers. As a result, training converges before exhibiting EoS. This motivates the two-stage RPS game as our first simple testbed for observing EoS.

B.1.2 Pursuit–Evasion Games

To study the learning dynamics of neural policies under multi-stage decisions and state transitions, we consider another two-player zero-sum Markov game called *pursuit–evasion* (PE) (Isaacs, 1951). Unlike two-stage RPS, where we train only second-stage hand-selection policies, both PE players act at every step, and their joint actions affect subsequent states and future rewards. Recent works have also considered training neural policies in PE games (Xu & Dang, 2025; Gonultas & Isler, 2025).

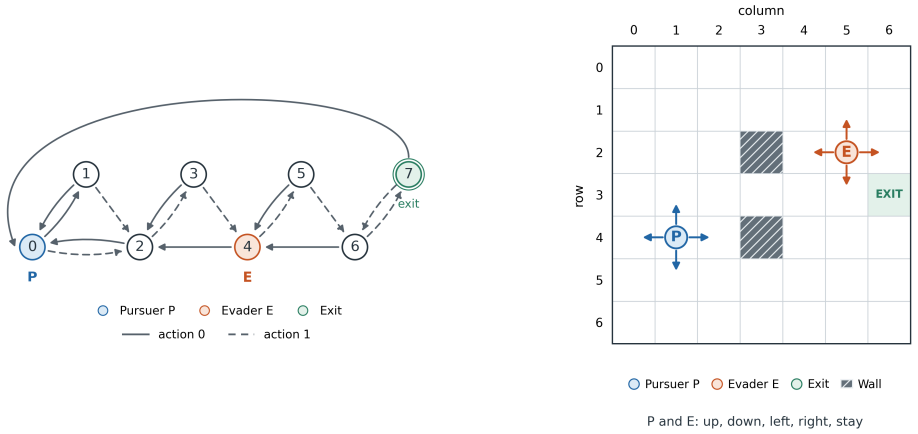

Figure 6: Pursuit–Evasion games: 8-node graph (left) and 7×7 grid (right). Graph actions 0 and 1 follow solid and dashed arrows, respectively. Both grid players can move up, down, left, right, or stay. Displayed positions are examples; initial states are uniform over all nonterminal joint positions.

Rule A pursuer P tries to catch an evader E, while the evader tries to avoid capture and reach an exit. Both players observe both positions and choose their actions simultaneously. Capture occurs when they move to the same location or exchange locations. If capture does not occur, the evader escapes upon reaching the exit. Either event ends the game. The exit is otherwise an ordinary location for the pursuer, who can occupy and leave it.

Each player's intended move independently becomes a stay with probability 0.05. The pursuer's reward at step t is

$$r_t = -0.01 + \mathbf{1}\{\text{capture at } t\} - \mathbf{1}\{\text{escape at } t\},$$

and the evader receives $-r_t$. Thus the step cost also applies to the final move.

Environment The state $s = (x_P, x_E)$ records the positions of the pursuer and the evader. Let $\mathcal{S}$ denote the set of nonterminal states, where the players occupy different locations and the evader is not at the exit. The initial distribution ρ is uniform over $\mathcal{S}$. We consider two environments of increasing state-space size, illustrated in Figure 6:

(i) 8-node graph. Both players move between the eight numbered nodes along directed arrows. Their action spaces are $\mathcal{A}_P = \mathcal{A}_E = \{0, 1\}$, where action 0 follows the solid outgoing arrow and action 1 follows the dashed outgoing arrow. For example, from node 0, either player can move to node 1 or node 2. Each arrow represents one move, regardless of its drawn length. Node 7 is the exit. The evader has seven possible non-exit positions and, for each one, the pursuer may occupy any of the other seven nodes. The environment therefore has $7 \times 7 = 49$ nonterminal states.

(ii) 7×7 *grid.* Each player can move one cell up, down, left, or right, or remain in place. A move into a wall or outside the board leaves the player in place. Rows and columns are indexed from 0 to 6; the walls are located at $(2, 3)$ and $(4, 3)$, and the exit is located at $(3, 6)$. Removing the two walls leaves 47 free cells. The evader may occupy any of the 46 non-exit cells and, for each evader position, the pursuer may occupy any of the other 46 free cells. The environment therefore has $46 \times 46 = 2116$ nonterminal states.

Objective We identify the evader with player x and the pursuer with player y in equation 21. We set the discount factor $\gamma = 0.95$ and entropy regularization parameter $\tau = 0.05$. For stationary policies π_E and π_P, let P_π denote the transition matrix among nonterminal states induced by $\pi = (\pi_E, \pi_P)$, including the action-slip probabilities, and let $r_{\pi,\tau}$ denote the corresponding expected one-step regularized reward. Then the objective in equation 21 can be written as

$$J_\tau(\pi_E, \pi_P) = \rho^\top (I - \gamma P_\pi)^{-1} r_{\pi,\tau}.$$

Policy networks Each player uses a separate network to parameterize a stationary policy as in equation 22. Both networks have three ReLU hidden layers and a linear output layer followed by

Setting	Two-stage RPS	8-node graph PE	7×7 grid PE
States	81	49	2116
Network input dimension	12	49	4
Hidden layers $\times$ width	3×12	3×12	3×32
Total optimizer updates	30k	60k	60k
Policy evaluation interval	10	200	100
EG step sizes η	$\left\{\frac{1}{13}, \frac{1}{14}, \frac{1}{15}\right\}$	$\left\{\frac{1}{6}, \frac{1}{8}, \frac{1}{12}\right\}$	$\left\{2, 1, \frac{1}{2}\right\}$
OGDA step sizes η	$\left\{\frac{1}{13\sqrt{3}}, \frac{1}{14\sqrt{3}}, \frac{1}{15\sqrt{3}}\right\}$	$\left\{\frac{1}{6\sqrt{3}}, \frac{1}{8\sqrt{3}}, \frac{1}{12\sqrt{3}}\right\}$	$\left\{\frac{4}{3}, \frac{2}{3}, \frac{1}{3}\right\}$

Table 1: Training and evaluation settings for the reported experiments.

softmax. For the 8-node graph, the input is $7e_s \in \mathbb{R}^{49}$, where e_s is the one-hot vector identifying the joint state. For the 7×7 grid, the input consists of the two players' normalized coordinates, $\frac{1}{3}(r_P, c_P, r_E, c_E)^\top - \mathbf{1} \in [-1, 1]^4$. The architectures are summarized together with the two-stage RPS network in Table 1.

Nash equilibrium We compute the reference Nash policies using an entropy-regularized variant of Shapley's value iteration (Shapley, 1953). At each iteration, we update each state's value by solving a regularized matrix game whose payoffs combine immediate rewards and discounted continuation values. Starting from $V_0 = 0$, we update

$$Q_k(s, a_E, a_P) = \bar{r}(s, a_E, a_P) + \gamma \sum_{s' \in \mathcal{S}} P(s' \mid s, a_E, a_P) V_k(s'),$$

$$V_{k+1}(s) = \min_{p_E \in \Delta_{|\mathcal{A}_E|}} \max_{p_P \in \Delta_{|\mathcal{A}_P|}} \left\{ p_E^\top Q_k(s) p_P + \tau H(p_P) - \tau H(p_E) \right\},$$

where $\bar{r}$ is the expected unregularized pursuer reward, and both $\bar{r}$ and P incorporate the action-slip probabilities. The statewise saddle-point policies at convergence define the reference equilibrium.

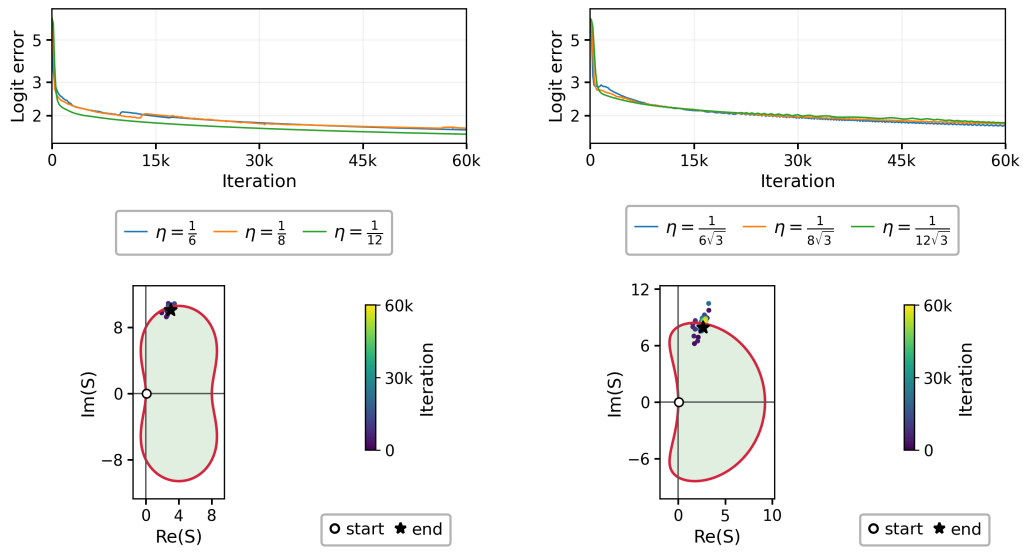

Figure 7: PE learning and eigenvalue dynamics on the 8-node graph ($\gamma = 0.95$, $\tau = 0.05$). The top row shows logit error and the bottom row shows trajectories of S, with EG on the left and OGDA on the right. Both algorithms exhibit EoS together with stable learning in this setting.

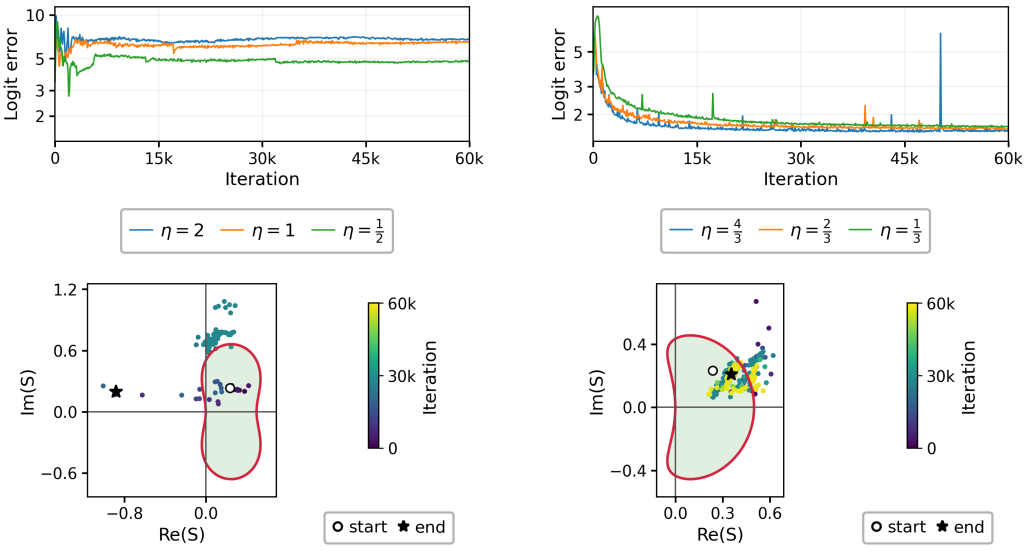

Figure 8: PE learning and eigenvalue dynamics on the 7×7 grid ($\gamma = 0.95$, $\tau = 0.05$). The top row shows logit error and the bottom row shows trajectories of S, with EG on the left and OGDA on the right. In this larger environment, EG becomes unstable and fails to reduce the logit error, whereas OGDA maintains EoS dynamics together with stable learning.

B.2 Neural Optimal Transport

Let μ and ν be probability distributions on $\mathbb{R}^d$. We consider the following minimax formulation of Neural Optimal Transport (Korotin et al., 2023):

$$\max_y \min_x \ \mathbb{E}_{\xi\sim\mu}\left[\frac{1}{2}\|\xi - T_x(\xi)\|^2 - V_y(T_x(\xi))\right] + \mathbb{E}_{\zeta\sim\nu}\left[V_y(\zeta)\right],$$

where $T_x : \mathbb{R}^d \to \mathbb{R}^d$ is the transport map and $V_y : \mathbb{R}^d \to \mathbb{R}$ is the potential, both parameterized by neural networks. Here, $(T_x)_\sharp\mu$ denotes the push-forward of μ through T_x. At the optimum, this formulation recovers the quadratic-cost optimal transport problem (Fan et al., 2023; Rout et al., 2022)

$$\min_{(T_x)_\sharp\mu=\nu} \mathbb{E}_{\xi\sim\mu}\left[\frac{1}{2}\|\xi - T_x(\xi)\|^2\right].$$

In our experiment, μ is a two-dimensional Gaussian distribution and ν is a two-dimensional Swiss-roll distribution, with 100 samples drawn from each distribution. The transport network T_x has two hidden layers of width 100 with ReLU activations, while the potential network V_y uses a Softplus activation. We evaluate training every 5000 iterations using the Wasserstein-2 distance between the transported distribution $(T_x)_\sharp\mu$ and the target distribution ν. Training is run for up to 10^6 iterations and terminated early if the iteration diverges or the loss falls below 10^{-1}.

B.3 Simple Minimization Problem

We consider a two-dimensional linear regression problem with $\xi = (\xi_1, \xi_2)^\top \sim \mathcal{N}(0, \mathrm{diag}(\lambda_1, \lambda_2))$ and target $f^\star(\xi) = \xi_2$. We use a one-hidden-layer, width-one linear network $h(\xi) = \alpha\beta_1\xi_1 + \alpha\beta_2\xi_2$ with squared loss, whose population objective is

$$L(\alpha, \beta_1, \beta_2) = \frac{\lambda_1}{2}(\alpha\beta_1)^2 + \frac{\lambda_2}{2}(\alpha\beta_2 - 1)^2.$$

Let $x = (\alpha, \beta_1, \beta_2)$ and introduce an auxiliary variable y that does not affect the objective, i.e., $f(x, y) := L(x)$. We then write the problem in minimax form as

$$\min_x \max_y f(x, y).$$

In our experiment, we set $\lambda_1 = 100$ and $\lambda_2 = 0.01$.

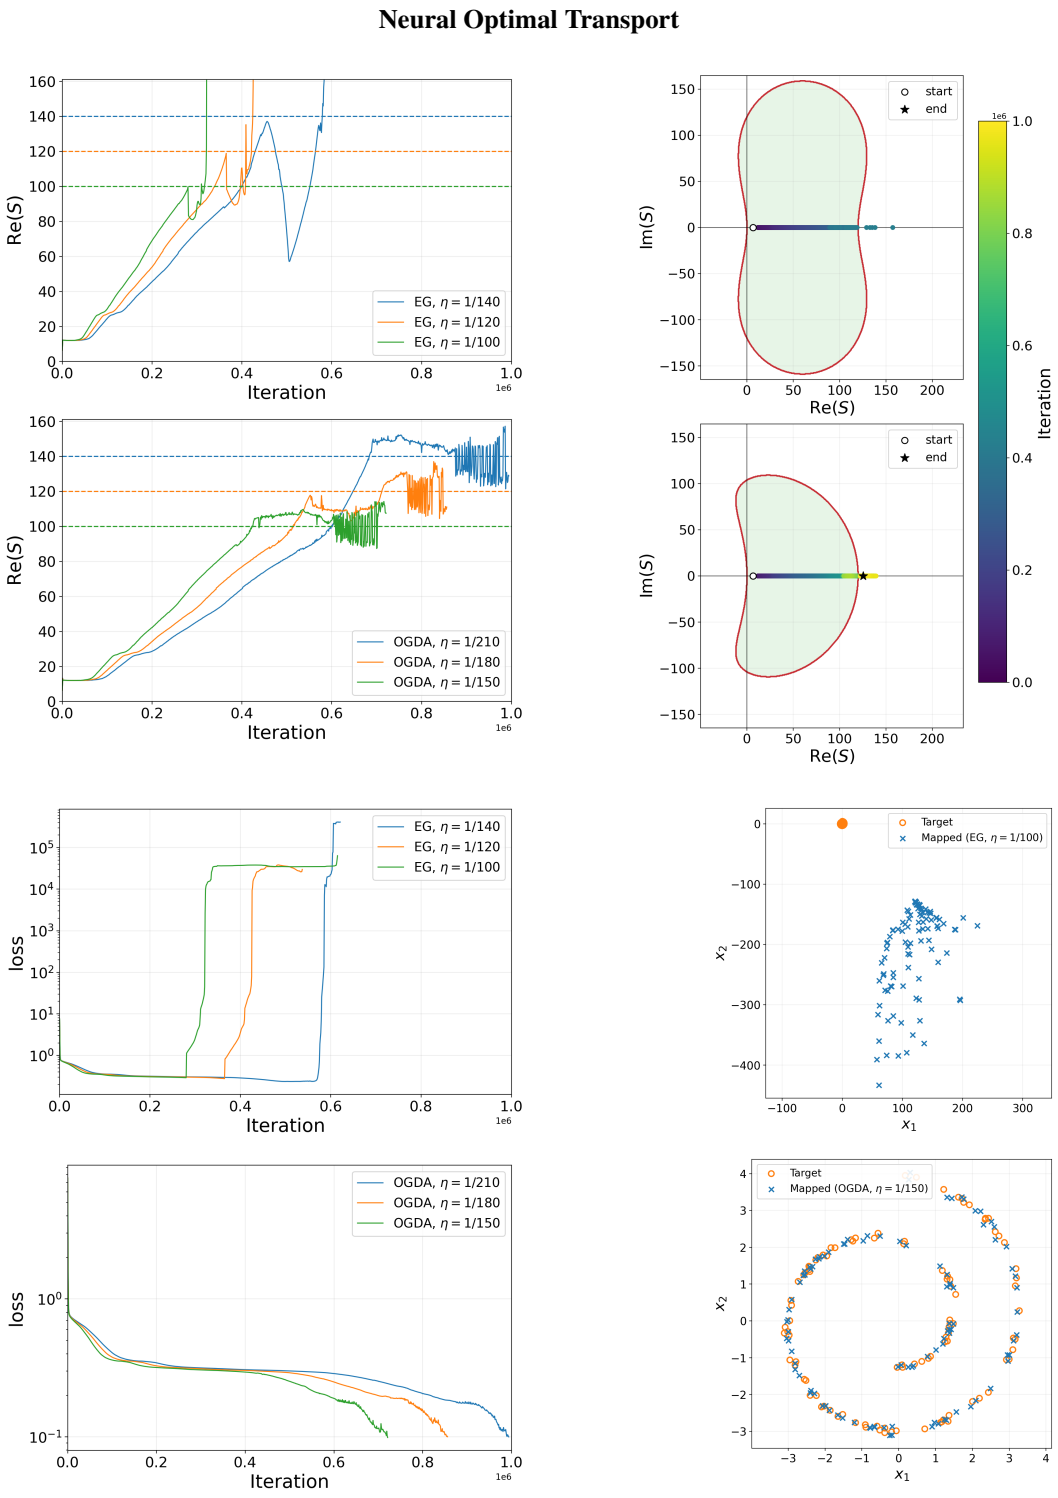

Figure 9: Neural Optimal Transport experiment. The top block shows the dominant eigenvalues trajectories for EG and OGDA together with the corresponding stability region. The bottom block shows the training losses and the transport mappings immediately before training termination for EG and OGDA. The loss denotes $W_2^2(T_\sharp\mu, \nu)$, where W_2 is the Wasserstein-2 distance.

C Proofs and Missing Details in Section 5

C.1 Proof of Proposition 5.1

We first establish the following lemma.

Lemma C.1. *Suppose T is C^2 and DT is normal in a neighborhood of z. Let Λ and u be the corresponding local eigenvalue and eigenvector branches, with $\|u\|^2 = 1$ if $\Lambda \in \mathbb{R}$ and $\|u\|^2 = 2$ otherwise. Then*

$$(DT^\top - \Lambda I) D^2T[u, \bar{u}] = \|u\|^2 (DT - \bar{\Lambda} I) \nabla \Lambda. \tag{23}$$

Consequently, if $\Lambda \notin \mathbb{R}$, then

$$D^2T[u, \bar{u}] = -2 \operatorname{Re} \nabla \Lambda - 2 \frac{DT - (\operatorname{Re} \Lambda) I}{\operatorname{Im} \Lambda} \operatorname{Im} \nabla \Lambda. \tag{24}$$

If $\Lambda \in \mathbb{R}$, then

$$D^2T[u, u] = \left[I + (DT^\top - \Lambda I)^\dagger (DT - DT^\top) \right] \nabla \Lambda, \tag{25}$$

where $\dagger$ denotes the Moore–Penrose inverse.

Proof. Normality implies $DT^\top u = \bar{\Lambda} u$ and $u^H DT = \Lambda u^H$. Differentiating $DTu = \Lambda u$ gives

$$D^2T[h, u] + DT\, Du[h] = (h^\top \nabla \Lambda) u + \Lambda Du[h],$$

so multiplication by u^H yields

$$\|u\|^2 h^\top \nabla \Lambda = u^H D^2T[h, u]. \tag{26}$$

Writing $\partial_i = \partial / \partial z_i$, then

$$\begin{aligned} \big(D(DT^\top)[\bar{u}, u]\big)_i &= \sum_{j,k=1}^{d} \bar{u}_j u_k\, \partial_j (DT^\top)_{ik} = \overline{\sum_{j,k=1}^{d} \bar{u}_j u_k\, \partial_j \partial_i T_k}, \\ &= \sum_{j,k=1}^{d} u_j \bar{u}_k\, \partial_i \partial_j T_k \\ &= \|u\|^2 \overline{\partial_i \Lambda}. \end{aligned}$$

Here we used equation 26 with $h = e_i$. Therefore,

$$D(DT^\top)[\bar{u}, u] = \|u\|^2 \overline{\nabla \Lambda}, \qquad D(DT^\top)[u, \bar{u}] = \|u\|^2 \nabla \Lambda. \tag{27}$$

Differentiating $DT^\top DT = DTDT^\top$ in direction $\bar{u}$ and applying the result to u gives

$$\begin{gathered} D(DT^\top)[\bar{u}, DTu] + DT^\top D^2T[\bar{u}, u] = D^2T[\bar{u}, DT^\top u] + DT\, D(DT^\top)[\bar{u}, u] \\ \implies (DT^\top - \bar{\Lambda} I) D^2T[u, \bar{u}] = \|u\|^2 (DT - \Lambda I) \overline{\nabla \Lambda}. \end{gathered}$$

Taking complex conjugates proves equation 23. If $\operatorname{Im} \Lambda \neq 0$, then $\|u\|^2 = 2$, and its imaginary part is

$$-\operatorname{Im} \Lambda\, D^2T[u, \bar{u}] = 2(DT - (\operatorname{Re} \Lambda) I) \operatorname{Im} \nabla \Lambda + 2 \operatorname{Im} \Lambda\, \operatorname{Re} \nabla \Lambda.$$

Dividing by $-\operatorname{Im} \Lambda$ proves equation 24.

For real Λ, u, $\|u\|^2 = 1$, so equations equation 23 and equation 26 give

$$\begin{gathered} (DT^\top - \Lambda I)\big(D^2T[u, u] - \nabla \Lambda\big) = (DT - DT^\top) \nabla \Lambda, \\ u^\top \big(D^2T[u, u] - \nabla \Lambda\big) = 0. \end{gathered}$$

Since normality implies $\ker(DT^\top - \Lambda I) = \operatorname{span}\{u\}$, applying $(DT^\top - \Lambda I)^\dagger$ proves equation 25. □

Proof of Proposition 5.1. All quantities below are evaluated at z. The boundary condition and the definition of ϕ give

$$\eta|\Lambda|^2 = 2\,\mathrm{Re}\,\Lambda \geq 0, \qquad \nabla\phi = \mathrm{Re}\big((\eta\bar{\Lambda} - 1)\nabla\Lambda\big). \tag{28}$$

Choose an orthonormal eigenbasis $u_1, \ldots, u_d$ of DT, with eigenvalues $\lambda_1, \ldots, \lambda_d$, taking conjugate pairs for nonreal eigenvalues and real vectors for real eigenvalues. By equation 23,

$$(\bar{\lambda}_j - \Lambda)u_j^H D^2T[u, \bar{u}] = \|u\|^2(\lambda_j - \bar{\Lambda})u_j^H \nabla\Lambda. \tag{29}$$

Non-real case. With the normalization $\|u\|^2 = 2$, take $u_1 = u/\sqrt{2}$ and $u_2 = \bar{u}/\sqrt{2}$. Equation equation 26 with $h = \bar{u}$ gives $2u^H\nabla\Lambda = u^H D^2T[u, \bar{u}]$. Thus the $j = 1$ instance of equation 29 yields

$$(\bar{\Lambda} - \Lambda)u^H D^2T[u, \bar{u}] = (\Lambda - \bar{\Lambda})u^H D^2T[u, \bar{u}], \qquad u^H D^2T[u, \bar{u}] = \bar{u}^H D^2T[u, \bar{u}] = 0,$$

where the second equality of projections uses the reality of $D^2T[u, \bar{u}]$. For $j \geq 3$,

$$u_j^H \nabla\Lambda = \frac{\bar{\lambda}_j - \Lambda}{2(\lambda_j - \bar{\Lambda})} u_j^H D^2T[u, \bar{u}].$$

Since the squared projection magnitudes of the real vector $D^2T[u, \bar{u}]$ agree on conjugate eigenvectors, pairing their contributions gives

$$\begin{aligned}
\beta &= \mathrm{Re}\big((\eta\bar{\Lambda} - 1)D^2T[u, \bar{u}]^H\nabla\Lambda\big) \\
&= \frac{1}{4}\sum_{j=3}^{d} \mathrm{Re}\left[(\eta\bar{\Lambda} - 1)\left(\frac{\bar{\lambda}_j - \Lambda}{\lambda_j - \bar{\Lambda}} + \frac{\lambda_j - \Lambda}{\bar{\lambda}_j - \bar{\Lambda}}\right)\right] |u_j^H D^2T[u, \bar{u}]|^2 \\
&= \frac{1}{2}\sum_{j=3}^{d} \frac{|\,\mathrm{Re}\,\lambda_j - \Lambda|^2 - (\mathrm{Im}\,\lambda_j)^2}{|\lambda_j - \Lambda|^2|\lambda_j - \bar{\Lambda}|^2} \\
&\qquad \times \left(\frac{|\Lambda|^2}{2}\big(1 + |1 - \eta\lambda_j|^2\big) - |\lambda_j|^2\right) |u_j^H D^2T[u, \bar{u}]|^2.
\end{aligned} \tag{30}$$

Here the last equality follows from $\eta|\Lambda|^2 = 2\,\mathrm{Re}\,\Lambda$. Under the first condition in equation 7, for each noncritical eigenvalue λ,

$$\begin{aligned}
|\,\mathrm{Re}\,\lambda - \Lambda|^2 - (\mathrm{Im}\,\lambda)^2 &= 2\left(\mathrm{Re}\,\lambda - \frac{\mathrm{Re}\,\Lambda}{2}\right)^2 + \frac{(\mathrm{Re}\,\Lambda)^2}{2} + (\mathrm{Im}\,\Lambda)^2 - |\lambda|^2 \\
&\geq \frac{|\Lambda|^2}{2} - |\lambda|^2 \geq 0, \\
\frac{|\Lambda|^2}{2}\big(1 + |1 - \eta\lambda|^2\big) - |\lambda|^2 &= (\eta\,\mathrm{Re}\,\Lambda - 1)|\lambda|^2 - 2\,\mathrm{Re}\,\Lambda\,\mathrm{Re}\,\lambda + |\Lambda|^2 \\
&\geq 2\left(\frac{\mathrm{Re}\,\Lambda}{|\Lambda|}|\lambda| - \frac{|\Lambda|}{2}\right)^2 + \frac{|\Lambda|^2}{2} - |\lambda|^2 \geq 0.
\end{aligned}$$

Thus equation 30 gives $\beta \geq 0$. Under the second condition,

$$|\,\mathrm{Re}\,\lambda - \Lambda|^2 - (\mathrm{Im}\,\lambda)^2 \geq 0, \qquad \frac{|\Lambda|^2}{2}\big(1 + |1 - \eta\lambda|^2\big) - |\lambda|^2 \leq 0,$$

so equation 30 gives $\beta \leq 0$.

Real case. Here $\Lambda \in \{0, 2/\eta\}$, u is real with $\|u\| = 1$, and $\nabla\phi = (\eta\Lambda - 1)\nabla\Lambda$. Taking $u_1 = u$, equation equation 25 and conjugate pairing yield

$$\beta = (\eta\Lambda - 1)\left[|u^\top\nabla\Lambda|^2 + \sum_{j=2}^{d} \frac{(\mathrm{Re}\,\lambda_j - \Lambda)^2 - (\mathrm{Im}\,\lambda_j)^2}{|\lambda_j - \Lambda|^2}|u_j^H\nabla\Lambda|^2\right]. \tag{31}$$

Under the first condition in equation 7, $\Lambda = 0$ would force every eigenvalue to vanish, contradicting simplicity since $d \geq 2$. Hence $\Lambda = 2/\eta$, and

$$(\operatorname{Re}\lambda_j - \Lambda)^2 - (\operatorname{Im}\lambda_j)^2 = 2\left(\operatorname{Re}\lambda_j - \frac{\Lambda}{2}\right)^2 + \frac{\Lambda^2}{2} - |\lambda_j|^2 \geq 0,$$

which gives $\beta \geq 0$ in equation 31. Under the second condition, $\Lambda = 2/\eta$ would imply

$$\frac{4}{\eta^2} \leq \frac{2|\lambda_j|^2}{1 + |1 - \eta\lambda_j|^2} \quad \Longleftrightarrow \quad \left|\lambda_j - \frac{2}{\eta}\right|^2 \leq 0 \qquad (j \geq 2),$$

again contradicting simplicity since $d \geq 2$. Thus $\Lambda = 0$, and the remaining inequality gives

$$\beta = -|u^\top \nabla\Lambda|^2 - \sum_{j=2}^{d} \frac{(\operatorname{Re}\lambda_j)^2 - (\operatorname{Im}\lambda_j)^2}{|\lambda_j|^2} |u_j^H \nabla\Lambda|^2 \leq 0.$$

This proves both implications. □

C.2 Verification of the conditions in Proposition 5.1.

Proposition 5.1 provides sufficient conditions for β to be positive or negative by assuming that DT is normal in a neighborhood of z. We examine these conditions in the two-stage RPS game and the simple regression problem of Liu et al. (2025) considered in Section 4.

For the two-stage RPS game, the dominant eigenvalue $S(z_t)$ lies along the imaginary axis. Hence $\Lambda(z_t)$, the eigenvalue of $DT(z_t)$, reaches the boundary of D_η near $(1 + i)/\eta$ for EG and $(1 + i\sqrt{3})/(2\eta)$ for OGDA; see Figure 1. We compute the remaining eigenvalues of DT at the final iterate. For EG, all noncritical eigenvalues satisfy $|\lambda| \leq |\Lambda|/\sqrt{2}$, the sufficient spectral condition for $\beta \geq 0$ in equation 7. For OGDA, this bound is not satisfied by all eigenvalues.

For the simple regression problem of Liu et al. (2025), the dominant eigenvalue $S(z_t)$ lies along the real axis. For EG, $S(z_t)$ approaches $1/\eta$ and the corresponding eigenvalue $\Lambda(z_t)$ approaches 0. At $\Lambda = 0$, the magnitude inequality in the sufficient condition for $\beta \leq 0$ is automatic. The remaining condition is $|\operatorname{Re}\lambda - \Lambda| = |\operatorname{Re}\lambda| \geq |\operatorname{Im}\lambda|$ for every noncritical eigenvalue λ. In Figure 10, this condition is also satisfied. For OGDA, $S(z_t)$ approaches $2/(3\eta)$ and the corresponding eigenvalue $\Lambda(z_t)$ approaches $2/\eta$; in this case, the noncritical eigenvalues satisfy the sufficient spectral condition for $\beta \geq 0$ in equation 7.

To quantify how close DT is to normal along training, we compute $\|DT^\top DT - DTDT^\top\|_F / \|DT\|_F^2$, where $\|\cdot\|_F$ denotes the Frobenius norm. This quantity remains small throughout training in all of the above experiments, indicating that DT is approximately normal along the observed trajectories. However, even though this quantity is near zero, we cannot apply the proposition 5.1 as it need DT to be normal in neighborhood of z. Tracking the sign of β without this normality assumption remains an interesting direction for future work.

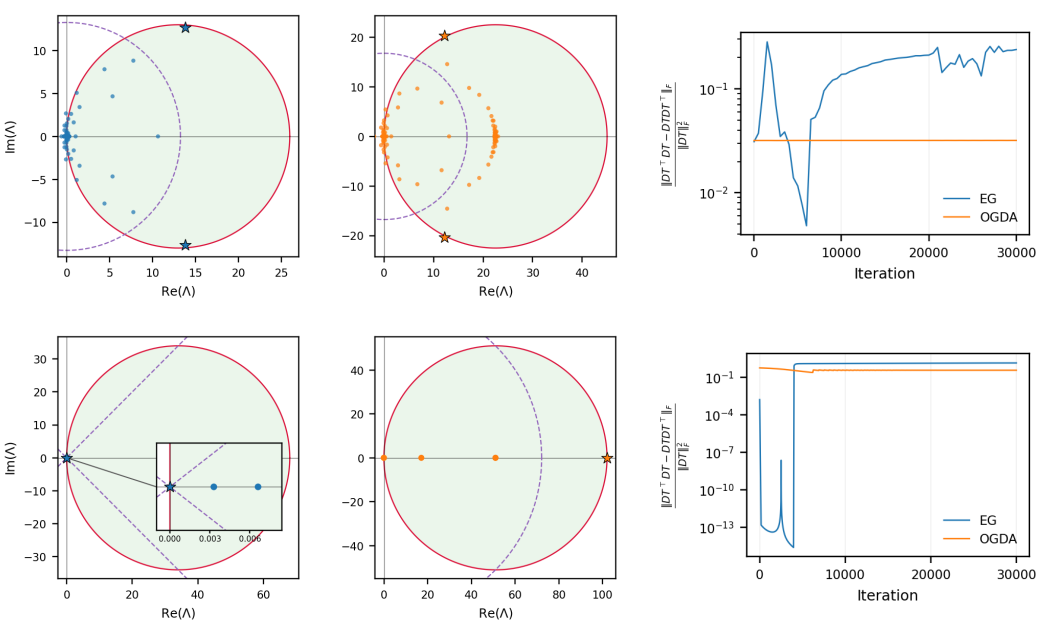

Figure 10: Eigenvalues of DT (left), where stars indicate Λ and the dashed line and circles indicate the spectral conditions in Proposition 5.1, and the normalized non-normality measure $\|DT^\top DT - DTDT^\top\|_F / \|DT\|_F^2$ (right). The top row shows the two-stage RPS game, and the bottom row shows the simple minimization problem of Liu et al. (2025). Blue and orange denote EG and OGDA, respectively.

D Proofs and Missing Details in Section 6

D.1 Toy-Example for Self-Stabilization Mechanism

For $\alpha, \beta, \gamma > 0$ and sufficiently small $\eta > 0$, consider

$$\min_{x_1,x_2} \max_{y_1,y_2} f(x_1, x_2, y_1, y_2) := \left(\frac{1}{\eta} + \frac{\sqrt{\gamma}}{2}\, x_2\right) x_1 y_1 + \frac{\beta x_2}{4\sqrt{\gamma}}(x_1^2 + y_1^2) - \frac{2\alpha}{\sqrt{\gamma}}\, x_2 + y_2.$$

Write $z = (x_1, y_1, x_2, y_2)$ and $F = (\partial_{x_1} f, -\partial_{y_1} f, \partial_{x_2} f, -\partial_{y_2} f)^\top$, so the EG operator is $T(z) = F(z - \eta F(z))$. Initialize at $z_0 = (r, -r, s_\star, 0)$. Along the reference trajectory $z_t^\dagger = (0, 0, s_\star, \eta t)$, all quantities below are constant. To leading order in η,

$$T = \begin{pmatrix} 0 \\ 0 \\ -2\alpha/\sqrt{\gamma} \\ -1 \end{pmatrix}, \qquad DF \simeq \begin{pmatrix} 0 & 1/\eta & 0 & 0 \\ -1/\eta & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad DT \simeq \begin{pmatrix} 1/\eta & 1/\eta & 0 & 0 \\ -1/\eta & 1/\eta & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$u \simeq \begin{pmatrix} 1 \\ i \\ 0 \\ 0 \end{pmatrix}, \qquad \Lambda \simeq \frac{1+i}{\eta}, \qquad \nabla\Lambda \simeq \frac{\sqrt{\gamma}}{2} \begin{pmatrix} 0 \\ 0 \\ 2+i \\ 0 \end{pmatrix}, \qquad \nabla\phi \simeq \frac{\sqrt{\gamma}}{2} \begin{pmatrix} 0 \\ 0 \\ 1 \\ 0 \end{pmatrix}.$$

The conditions $u^H T = 0$ and $T, \nabla\Lambda \in \ker(DT)$ hold exactly, while normality holds to leading order. Moreover,

$$-\nabla\phi^\top T = \alpha + O(\eta^2), \qquad \nabla\phi^\top D^2 T[u, \bar{u}] = \beta + O(\eta^2),$$
$$\left|\nabla\phi^\top D^2 T[u, u]\right| = \gamma + O(\eta^2).$$

D.2 Discussion of the Assumptions

We provide further discussion of the assumptions used in Section 6. The first four assumptions play the same roles as the corresponding assumptions in Damian et al. (2023), adapted to the update operator T. The last two assumptions are additionally required for our minimax analysis.

Assumption (Spectral separation)**.** For every t, for some fixed $\Delta > 0$,

$$\min_{\lambda \in \{\Lambda_t, \bar{\Lambda}_t\}} \operatorname{dist}(\lambda, \operatorname{eig}(DT_t) \setminus \{\lambda\}) > \tfrac{\Delta}{\eta}.$$

This assumption separates the eigenspace associated with Λ_t from the remaining eigenspaces of DT_t. It prevents the critical eigenspace from mixing strongly with the other eigenspaces as the projected trajectory evolves, and plays the same role as the eigengap condition in Assumption 2 of Damian et al. (2023).

Assumption (Regularity)**.** Let ρ_2, ρ_3 be the minimum constants such that $\|D^2 T(z)\|_{op} \le \rho_2$ for all z and $D^2 T$ is ρ_3-Lipschitz with respect to $\|\cdot\|_{op}$. Assume

$$\rho_3 = O(\eta \rho_2^2).$$

This assumption controls the Taylor expansion of the update operator T. The bound on $D^2 T$ controls the quadratic term, while the Lipschitz condition controls the higher-order Taylor error. It plays the same role as Assumption 3 of Damian et al. (2023).

Assumption (Scale conditions)**.** For every t,

$$\frac{\alpha_t}{\|T_t\|\|\nabla\phi_t^\perp\|}, \frac{\beta_t^\perp}{\|\nabla\phi_t^\perp\|^2}, \frac{\|\nabla\phi_t^\perp\|}{\rho_2} = \Theta(1).$$

This assumption requires the scales of progressive sharpening, self-stabilization, and the quadratic variation of T to remain comparable. In particular, it prevents any of the quantities governing the local self-stabilization dynamics from becoming degenerate relative to the others. It plays the same role as Assumption 4 of Damian et al. (2023).

Our assumption	Damian et al. (2023)
Progressive sharpening ($\alpha_t > 0$)	Assumption 1
Spectral separation	Assumption 2
Regularity	Assumption 3
Scale conditions	Assumption 4
Non-worst-case orthogonal dynamics	Assumption 5
Approximate normality	–
Quadratic alignment	–

Table 2: Correspondence between our assumptions and those of Damian et al. (2023).

Assumption (Non-worst-case orthogonal dynamics)**.** For every t,

$$\frac{\|D^2T_t[v,w]\|}{\rho_2\|v\|\|w\|}, \frac{|\langle v_t^{\star,\perp}, DT_t v_t^{\star,\perp}\rangle|}{\|DT_t\|\|v_t^{\star,\perp}\|^2}, \frac{\mathrm{dist}(\lambda, D_\eta)}{\|DT_t\|} \le O(\epsilon)$$

for every $v, w \in U_t^\perp$, and $\lambda \in \mathrm{eig}(DT_t) \setminus \{\Lambda_t, \bar{\Lambda}_t\}$.

This assumption excludes worst-case behavior in directions orthogonal to the eigenspace U_t. The three conditions respectively control the quadratic interaction within $U_t^\perp$, the action of DT_t on the orthogonal deviation, and the location of the non-dominant eigenvalues relative to the stability region. It plays the same role as Assumption 5 of Damian et al. (2023).

The preceding assumptions directly parallel those used in the analysis of Damian et al. (2023). In addition, our minimax analysis requires the following two assumptions to control the nonsymmetric and complex-valued dynamics.

Assumption (Approximate normality)**.** For every t,

$$\left\| D\big(DT(DT)^\top - (DT)^\top DT\big)\,(z_t^\dagger) \right\|_{op} = O\left(\frac{\epsilon\rho_2}{\eta}\right).$$

Since $DT(z_t^\dagger)$ is normal along the projected trajectory, this assumption controls how rapidly normality can be violated in a neighborhood of $z_t^\dagger$. This allows the local spectral structure around the projected trajectory to remain sufficiently close to the normal case throughout the coupling argument. This condition has no counterpart in Damian et al. (2023).

Assumption (Quadratic alignment)**.** For every t,

$$\|P^\perp_{U_t,\nabla\phi_t} D^2T_t[u_t, \bar{u}_t]\| = O(\epsilon\rho_2)$$

where $P^\perp_{U_t,\nabla\phi_t}$ denote the orthogonal projector onto $\mathrm{span}\{U_t, \nabla\phi_t\}^\perp$.

This assumption requires the quadratic correction $D^2T_t[u_t, \bar{u}_t]$ to lie approximately in the subspace spanned by U_t and $\nabla\phi_t$. Thus, the quadratic term responsible for self-stabilization does not produce a significant component in unrelated orthogonal directions. This condition is specific to our minimax analysis.

For completeness, Table 2 summarizes the correspondence with the assumptions of Damian et al. (2023). The progressive sharpening condition $\alpha_t > 0$, introduced separately in the main text, corresponds to their Assumption 1.

D.3 Predicted Dynamics

We first define the deviation of the original trajectory from the projected trajectory.

Definition D.1. Let

$$v_t := z_t - z_t^\dagger, \qquad x_t := \frac{u_t^H v_t}{\|u_t\|}, \qquad y_t := (\eta\Lambda_t - 1)(\nabla\Lambda_t^\perp)^H v_t.$$

We approximate v_t by the following predicted deviation $v_t^\star$.

Definition D.2 (Predicted Dynamics)**.** Set $v_0^\star := v_0$. Define $v_{t+1}^\star$ by

$$\begin{aligned} v_{t+1}^\star = P_{t+1}^\perp (I - \eta D T_t) P_t^\perp v_t^\star \\ + \frac{\eta}{2} \frac{\beta_t^\perp}{\|\nabla\phi_t^\perp\|^2} P_{t+1}^\perp \nabla\phi_t^\perp \left(\delta_t^2 - |x_t^\star|^2\right) \\ - \frac{\eta}{2} P_{t+1}^\perp \operatorname{Re}\left((x_t^\star)^2 D^2 T_t[u_t, u_t]\right) \mathbf{1}_{\{\Lambda_t \notin \mathbb{R}\}} \\ + \|u_t\| \operatorname{Re}\left((1 - \eta\Lambda_t)(1 + \eta\overline{y_t^\star}) x_t^\star u_{t+1}\right). \end{aligned} \tag{32}$$

Let

$$x_t^\star := \frac{u_t^H v_t^\star}{\|u_t\|}, \qquad y_t^\star := (\eta\Lambda_t - 1)(\nabla\Lambda_t^\perp)^H v_t^\star.$$

For simplicity, we define the following prediction map.

Definition D.3. For $v \in \mathbb{R}^d$, let $\operatorname{step}_t(v)$ be the vector defined by equation 32 with $v_t^\star = v$, $x = u_t^H v / \|u_t\|$, and $y = (\eta\Lambda_t - 1)(\nabla\Lambda_t^\perp)^H v$, so that we have

$$v_{t+1}^\star = \operatorname{step}_t(v_t^\star).$$

We need generalized coefficients for β and γ.

Definition D.4. Let $R_k := P_{k+1}^\perp (I - \eta D T_k) P_k^\perp$ and define

$$\beta_{s\to t} := (\nabla\phi_{t+1}^\perp)^\top R_t \cdots R_{s+1} P_{s+1}^\perp \frac{\beta_s^\perp}{\|\nabla\phi_s^\perp\|^2} \nabla\phi_s^\perp, \tag{33}$$

$$\gamma_{s\to t}(\vartheta) := (\nabla\phi_{t+1}^\perp)^\top R_t \cdots R_{s+1} P_{s+1}^\perp \operatorname{Re}\left(e^{2i\vartheta} D^2 T_s[u_s, u_s]\right) \mathbf{1}_{\{\Lambda_s \notin \mathbb{R}\}}. \tag{34}$$

An empty product is the identity. Here, both coefficients are real.

Lemma D.5 (Predicted dynamics Lemma)**.** *If $P_0^\perp v_0 = 0$, then*

$$x_{t+1}^\star = (1 - \eta\Lambda_t)(1 + \eta\overline{y_t^\star}) x_t^\star, \tag{35}$$

$$\operatorname{Re}(y_{t+1}^\star) = \frac{\eta}{2} \sum_{s=0}^{t} \left[\beta_{s\to t}\left(\delta_s^2 - |x_s^\star|^2\right) - \gamma_{s\to t}(\arg x_s^\star) |x_s^\star|^2\right]. \tag{36}$$

The phase contribution is zero when $x_s^\star = 0$. If $\phi(z_t^\dagger) = 0$, then

$$|x_{t+1}^\star| = |1 + \eta y_t^\star| \, |x_t^\star|. \tag{37}$$

Note that when the coefficients $\beta_{s\to t}$ are constant, our update reduces to the simple case discussed in Section 6.1. When $|x_t^\star|$ is large, equation 36 demonstrates that a self-stabilization force emerges, acting to decrease $\operatorname{Re}(y_{t+1}^\star)$. Unlike in Section 6.1, the strength of this force varies with t.

Proof. Since $P_0^\perp v_0 = 0$, we can inductively unfold equation 32 to obtain

$$\begin{aligned} v_{t+1}^\star = \frac{\eta}{2} \sum_{s=0}^{t} R_t \cdots R_{s+1} P_{s+1}^\perp \Bigg[\frac{\beta_s^\perp}{\|\nabla\phi_s^\perp\|^2} \nabla\phi_s^\perp \left(\delta_s^2 - |x_s^\star|^2\right) \\ - \operatorname{Re}\left((x_s^\star)^2 D^2 T_s[u_s, u_s]\right) \mathbf{1}_{\{\Lambda_s \notin \mathbb{R}\}} \Bigg] \\ + \begin{cases} (1 - \eta\Lambda_t)(1 + \eta\overline{y_t^\star}) x_t^\star u_{t+1}, & \Lambda_t \in \mathbb{R}, \\ \sqrt{2} \operatorname{Re}\left((1 - \eta\Lambda_t)(1 + \eta\overline{y_t^\star}) x_t^\star u_{t+1}\right), & \Lambda_t \notin \mathbb{R}. \end{cases} \end{aligned} \tag{38}$$

Recall that

$$x^\star_{t+1} = \frac{u^H_{t+1} v^\star_{t+1}}{\|u_{t+1}\|}, \qquad \mathrm{Re}(y^\star_{t+1}) = (\nabla \phi^\perp_{t+1})^\top v^\star_{t+1}.$$

The first coordinate extracts the critical component of equation 38, while the second annihilates it since $\nabla \phi^\perp_{t+1} \in U^\perp_{t+1}$. Thus, using $u^H_{t+1} \bar{u}_{t+1} = 0$ and $\|u_{t+1}\|^2 = 2$ in the non-real case,

$$x^\star_{t+1} = (1 - \eta \Lambda_t)(1 + \eta \overline{y^\star_t}) x^\star_t.$$

Moreover, from the definitions equation 33–equation 34 and

$$(x^\star_s)^2 = |x^\star_s|^2 e^{2i \arg x^\star_s} \qquad (x^\star_s \neq 0),$$

with the phase contribution defined to be zero when $x^\star_s = 0$, we obtain

$$\mathrm{Re}(y^\star_{t+1}) = \frac{\eta}{2} \sum_{s=0}^{t} \left[\beta_{s \to t} \left(\delta_s^2 - |x^\star_s|^2 \right) - \gamma_{s \to t} (\arg x^\star_s) |x^\star_s|^2 \right].$$

Finally, if $\phi(z^\dagger_t) = 0$, then $|1 - \eta \Lambda_t| = 1$, and hence

$$|x^\star_{t+1}| = |1 + \eta \overline{y^\star_t}| \, |x^\star_t| = |1 + \eta y^\star_t| \, |x^\star_t|.$$

□

In particular, when $\Lambda_t = 2/\eta$, $y^\star_t$ is real and $x^\star_{t+1} = -(1 + \eta y^\star_t) x^\star_t$.

D.4 Properties of the constrained trajectory

All statements below use the assumptions and conventions stated in Section 6.

Definition D.6. Let $\mathcal{S}_t := B(z_t^\dagger, \Delta/(4\eta\rho_2))$, where Δ is the fixed constant in Assumption 6.3(1) and $B(x, r)$ is the ball of radius r centered at x.

Lemma D.7 (Lipschitz Lemma)**.** *Fix an index t for which $\phi(z_t^\dagger) = 0$. With a local $C^{1,1}$ choice of phase for u, the following bounds hold on $\mathcal{S}_t$.*

1. *DT is ρ_2-Lipschitz with respect to the operator norm.*
2. *T is $O(\eta^{-1})$-Lipschitz.*
3. *Λ is $O(\rho_2)$-Lipschitz.*
4. *u and P are $O(\eta\rho_2)$-Lipschitz.*
5. *$\nabla\Lambda$ and $\nabla\phi$ are $O(\eta\rho_2^2)$-Lipschitz.*
6. *$\|D^2 u\|_{\mathrm{op}} = O(\eta^2\rho_2^2)$ almost everywhere. The bound holds everywhere if T is C^3.*

We prove the following two auxiliary lemmas before proving Lemma D.7.

Lemma D.8. *For every t with $\phi(z_t^\dagger) = 0$ and every $z \in \mathcal{S}_t$, $\Lambda(z)$ remains simple and*

$$|\Lambda(z) - \Lambda_t| \le \frac{\Delta}{4\eta}, \qquad |\lambda - \Lambda(z)| > \frac{3\Delta}{4\eta} \quad \big(\lambda \in \mathrm{eig}(DT_t) \setminus \{\Lambda_t\}\big). \tag{39}$$

Every unit right or left eigenvector w for this eigenvalue satisfies

$$\| \left(I - \frac{u_t u_t^H}{\|u_t\|^2}\right) w\| \le \frac{1}{3}, \qquad \frac{|u_t^H w|^2}{\|u_t\|^2} \ge \frac{8}{9}. \tag{40}$$

Thus we can choose u locally of class $C^{1,1}$, with $\|u(z)\| = \|u_t\|$ and $u_t^H u(z) > 0$ real, and normalize the left eigenvector ℓ by $\ell(z)^H u(z) = 1$ with $\|\ell(z)\| \le 9/(7\|u_t\|)$.

Proof. Fix t. The bound $\|D^2 T\|_{\mathrm{op}} \le \rho_2$ and the definition of $\mathcal{S}_t$ give

$$\|DT(z) - DT_t\| \le \rho_2 \|z - z_t^\dagger\| \le \frac{\Delta}{4\eta}, \qquad z \in \mathcal{S}_t. \tag{41}$$

For an eigenvalue λ of $DT(z)$ and a corresponding unit eigenvector w, normality of DT_t gives

$$\mathrm{dist}(\lambda, \mathrm{eig}(DT_t)) \le \|(DT_t - \lambda I)w\| = \|(DT_t - DT(z))w\| \le \frac{\Delta}{4\eta}.$$

Along the segment from $z_t^\dagger$ to z, no eigenvalue can therefore cross the circle of radius $\Delta/(2\eta)$ about Λ_t. This circle initially contains one eigenvalue counted with algebraic multiplicity, so it continues to contain exactly one. This proves simplicity and equation 39.

For a unit right eigenvector, expanding in an orthonormal eigenbasis of DT_t now yields

$$\frac{3\Delta}{4\eta} \| \left(I - \frac{u_t u_t^H}{\|u_t\|^2}\right) w\| \le \|(DT_t - \Lambda(z)I)w\| \le \frac{\Delta}{4\eta}.$$

Applying the same argument to $DT(z)^H$ proves the left-eigenvector bound. In particular, any unit left eigenvector has inner product of magnitude at least $\|u_t\|(8/9 - 1/9) = 7\|u_t\|/9$ with $u(z)$. The phase choice and the stated normalization follow. The simple eigenvalue and its eigenvectors are smooth functions of the matrix entries locally; since DT is $C^{1,1}$ by Assumption 6.3(2), their local branches are $C^{1,1}$. □

Lemma D.9. *Fix t with $\phi(z_t^\dagger) = 0$. For every $z \in \mathcal{S}_t$ and $v \in \mathbb{C}^d$ satisfying $u_t^H v \in \mathbb{R}$, we have*

$$\|v\| \le O(\eta)\|(DT(z) - \Lambda(z)I)v\| + \frac{|\mathrm{Re}(u(z)^H v)|}{\|u(z)\|}. \tag{42}$$

The same estimate holds for $DT(z)^H$ and $\overline{\Lambda(z)}$, using its eigenvector of norm $\|u_t\|$ with positive real inner product with u_t.

Proof. Fix $z \in \mathcal{S}_t$ and such a vector v. Set $w := v - \mathrm{Re}(u(z)^H v)u(z)/\|u(z)\|^2$. Then $u_t^H w \in \mathbb{R}$ and $\mathrm{Re}(u(z)^H w) = 0$. By equation 40,

$$\frac{|u_t^H w|}{\|u_t\|} \le \frac{\| \left(I - \frac{u_t u_t^H}{\|u_t\|^2}\right) u(z)\|}{u_t^H u(z)/\|u_t\|}\|w\| \le \frac{1}{\sqrt{8}}\|w\|, \qquad \| \left(I - \frac{u_t u_t^H}{\|u_t\|^2}\right) w\| \ge \sqrt{\frac{7}{8}}\|w\|.$$

Normality, equation 39, and equation 41 give

$$\|(DT(z) - \Lambda(z)I)w\| \ge \frac{\Delta}{4\eta}\left(3\sqrt{\frac{7}{8}} - 1\right)\|w\| = \Omega(\eta^{-1})\|w\|.$$

Since $(DT(z) - \Lambda(z)I)u(z) = 0$, this proves equation 42. The adjoint case uses the same argument and Lemma D.8. □

Proof of Lemma D.7. Fix t and take $z, z' \in \mathcal{S}_t$.

1. Integrating D^2T along the segment from z' to z gives
$$\|DT(z) - DT(z')\| \le \rho_2\|z - z'\|.$$

2. Normality at $z_t^\dagger$ and $\phi(z_t^\dagger) = 0$ give $|\Lambda_t| \le 2/\eta$ and $\|DT_t\| = \max_{\lambda \in \mathrm{eig}(DT_t)} |\lambda|$. For every noncritical eigenvalue, Assumption 6.3(4) gives
$$|\lambda| \le \frac{2}{\eta} + \mathrm{dist}(\lambda, D_\eta) \le \frac{2}{\eta} + O(\epsilon)\|DT_t\|.$$
Taking the maximum and absorbing the last term for sufficiently small ϵ yields $\|DT_t\| \le 3/\eta$. Part (1) therefore yields
$$\|DT_t\| \le \frac{3}{\eta}, \qquad \sup_{z \in \mathcal{S}_t}\|DT(z)\| \le \frac{3 + \Delta/4}{\eta} = O(\eta^{-1}). \tag{43}$$
Integrating DT proves (2). For the critical eigenvalues, the boundary condition gives $|1 - \eta\Lambda_t| = |1 - \eta\bar{\Lambda}_t| = 1$. For every other eigenvalue, Assumption 6.3(4) gives $|1 - \eta\lambda| \le 1 + O(\epsilon)\eta\|DT_t\|$. Hence
$$\|I - \eta DT_t\| = \max_{\lambda \in \mathrm{eig}(DT_t)} |1 - \eta\lambda| \le 1 + O(\epsilon), \tag{44}$$
which will be used in the deviation estimates below.

3. Differentiate $DT(z)u(z) = \Lambda(z)u(z)$ in a real direction h and multiply by $\ell(z)^H$. The terms containing $Du(z)h$ cancel, giving
$$(\nabla\Lambda(z))^\top h = \ell(z)^H D^2T(z)[h, u(z)]. \tag{45}$$
Lemma D.8 and $\|D^2T\|_{\mathrm{op}} \le \rho_2$ give $\|\nabla\Lambda(z)\| = O(\rho_2)$. Integration along a segment proves (3).

4. The differentiated eigenvector equation gives
$$(DT(z) - \Lambda(z)I)Du(z)h = -D^2T(z)[h, u(z)] + ((\nabla\Lambda(z))^\top h)u(z).$$
By (3), the right-hand side has norm $O(\rho_2)\|h\|$. The phase and fixed-norm conditions imply $u_t^H Du(z)h \in \mathbb{R}$ and $\mathrm{Re}(u(z)^H Du(z)h) = 0$. Lemma D.9 therefore yields $\|Du(z)h\| = O(\eta\rho_2)\|h\|$.

   For real Λ, use $P = uu^\top$. For non-real Λ, put $V = [\mathrm{Re}\, u, \mathrm{Im}\, u]$, so $P = V(V^\top V)^{-1}V^\top$. At the normal base point the two columns are orthogonal with norm 1, while Lemma D.8 gives $\|u(z) - u_t\| \le 2\sqrt{1 - \sqrt{8/9}} < 1$. Thus $\|(V^\top V)^{-1}\| = O(1)$. Differentiating the projector formula gives, in both cases, $\sup_{\|v\| \le 1}\|DP(z)[h, v]\| = O(\|Du(z)h\|) = O(\eta\rho_2)\|h\|$. This proves (4).

5. Apply Lemma D.9 to $DT(z)^H$ as in (4). Its eigenvector with fixed norm has derivative $O(\eta\rho_2)$; the normalization $\ell^H u = 1$ preserves this order by Lemma D.8. Hence $\|D\ell\|_{\mathrm{op}} = O(\eta\rho_2)$. Using (4) and the ρ_3-Lipschitz bound for D^2T in equation 45 gives

$$\mathrm{Lip}(\nabla\Lambda) = O(\rho_3 + \eta\rho_2^2) = O(\eta\rho_2^2).$$

Finally,

$$\nabla\phi(z) = -\operatorname{Re}\Big(\overline{1-\eta\Lambda(z)}\,\nabla\Lambda(z)\Big). \tag{46}$$

Since $|1-\eta\Lambda(z)| = O(1)$ on $\mathcal{S}_t$, part (3) and the bound just proved give $\mathrm{Lip}(\nabla\phi) = O(\eta\rho_2^2)$.

6. Differentiate the eigenvector equation twice, in real directions h, k, at points where the second derivatives exist. The $C^{1,1}$ regularity established in Lemma D.8 ensures that this holds almost everywhere. Suppressing the argument z gives

$$\begin{aligned}(DT-\Lambda I)D^2u[h,k] = &-D^3T[h,k,u] - D^2T[h, Du\,k] - D^2T[k, Du\,h]\\ &+ (h^\top\nabla^2\Lambda\,k)u + ((\nabla\Lambda)^\top h)Du\,k + ((\nabla\Lambda)^\top k)Du\,h.\end{aligned}$$

By (3)–(5) and $\|D^3T\|_{\mathrm{op}} \le \rho_3$, its right-hand side has norm $O(\rho_3+\eta\rho_2^2)\|h\|\|k\|$. Differentiating the normalization conditions twice also gives

$$u_t^H D^2u(z)[h,k] \in \mathbb{R}, \qquad \operatorname{Re}\big(u(z)^H D^2u(z)[h,k]\big) = -\operatorname{Re}\big((Du(z)h)^H Du(z)k\big).$$

We may therefore apply Lemma D.9 directly with $v = D^2u(z)[h,k]$ to obtain

$$\|D^2u(z)[h,k]\| \le O(\eta\rho_3+\eta^2\rho_2^2)\|h\|\|k\| = O(\eta^2\rho_2^2)\|h\|\|k\|.$$

This proves (6).

□

Lemma D.10 (Constrained step)**.** *For every $t \le \mathcal{T}$,*

$$\phi(z_t^\dagger) = 0, \qquad z_{t+1}^\dagger = z_t^\dagger - \eta P^{\perp}_{U_t,\nabla\phi_t} T_t + O(\epsilon^2\eta\|T_t\|). \tag{47}$$

Proof. We prove $\phi(z_t^\dagger) = 0$ by induction. The base case holds by initialization. Suppose that $\phi(z_t^\dagger) = 0$, and set $z' := z_t^\dagger - \eta T_t$ and $d := z_{t+1}^\dagger - z_t^\dagger$. Since $z_t^\dagger \in \mathcal{M}$ is a feasible competitor in the projection defining $z_{t+1}^\dagger$, we have

$$\|d+\eta T_t\| = \|z_{t+1}^\dagger - z'\| \le \|z_t^\dagger - z'\| = \eta\|T_t\|, \qquad \|d\| \le 2\eta\|T_t\|. \tag{48}$$

Assumption 6.3(3) gives

$$\alpha_t = \Theta(\rho_2\|T_t\|), \qquad \eta^2\rho_2\|T_t\| = \Theta(\epsilon_t^2) = O(\epsilon^2).$$

For sufficiently small ϵ, therefore, $2\eta\|T_t\| \le \Delta/(4\eta\rho_2)$. Thus the ball of radius $2\eta\|T_t\|$ about $z_t^\dagger$ lies in $\mathcal{S}_t$, and Lemma D.7 applies throughout the step.

We first estimate the constraint gradients on this ball. Let $h(z) := u(z)^H T(z)$. The Lipschitz bound for T gives $\|T(z)\| = O(\|T_t\|)$. Differentiating h yields

$$\begin{aligned}\nabla h(z) &= (DT(z))^\top\overline{u(z)} + (Du(z))^H T(z)\\ &= \Lambda_t\bar{u}_t + O\Big(\rho_2\|z-z_t^\dagger\| + \eta\rho_2\|T(z)\|\Big)\\ &= \Lambda_t\bar{u}_t + O(\eta\rho_2\|T_t\|).\end{aligned}$$

Likewise,

$$\nabla\phi(z) = \nabla\phi_t + O(\eta\rho_2^2\|z-z_t^\dagger\|) = \nabla\phi_t + O(\eta^2\rho_2^2\|T_t\|).$$

The leading equality-constraint gradients span U_t with scale $|\Lambda_t| = \Omega(\eta^{-1})$, whereas $\|\nabla\phi_t^{\perp}\| = \Theta(\rho_2)$. Hence the KKT conditions give multipliers ξ and $\nu \geq 0$ such that

$$\begin{aligned} d = -\eta T_t - \mathrm{Re}\Big(\bar{\xi}\,\nabla h(z_{t+1}^{\dagger})\Big) - \nu\nabla\phi(z_{t+1}^{\dagger}), \\ \nu\phi(z_{t+1}^{\dagger}) = 0. \end{aligned} \tag{49}$$

Here ξ is real in the real case; otherwise, its real and imaginary parts are the multipliers for $\mathrm{Re}\, h = 0$ and $\mathrm{Im}\, h = 0$. Substituting the gradient estimates gives

$$d = \ -\eta T_t - \mathrm{Re}(\bar{\xi}\Lambda_t\bar{u}_t) - \nu\nabla\phi_t + O\left(\epsilon^2\left(\frac{|\xi|}{\eta} + \nu\rho_2\right)\right). \tag{50}$$

We next bound the multipliers directly. Normality and the normalization of u_t give

$$\mathrm{Re}(\bar{\xi}\Lambda_t\bar{u}_t) \in U_t, \qquad \|\,\mathrm{Re}(\bar{\xi}\Lambda_t\bar{u}_t)\| = |\xi|\,|\Lambda_t|.$$

Projecting the sum $\mathrm{Re}(\bar{\xi}\Lambda_t\bar{u}_t) + \nu\nabla\phi_t$ onto $U_t^{\perp}$ first bounds $\nu\rho_2$; subtracting $\nu\nabla\phi_t$ then bounds $|\xi|/\eta$. Using $\|\nabla\phi_t^{\perp}\| = \Theta(\rho_2)$, $\|\nabla\phi_t\| = O(\rho_2)$, and equation 48–equation 50, we obtain

$$\begin{aligned} \frac{|\xi|}{\eta} + \nu\rho_2 &\leq O(1)\|\,\mathrm{Re}(\bar{\xi}\Lambda_t\bar{u}_t) + \nu\nabla\phi_t\| \\ &\leq O(\eta\|T_t\|) + O(\epsilon^2)\left(\frac{|\xi|}{\eta} + \nu\rho_2\right). \end{aligned}$$

Absorbing the last term yields

$$\frac{|\xi|}{\eta} + \nu\rho_2 = O(\eta\|T_t\|).$$

Thus the remainder in equation 50 is $O(\epsilon^2\eta\|T_t\|)$.

We now use the equality constraint to bound $P_t d$. Expanding T and using the Lipschitz bound for u gives

$$\begin{aligned} 0 = h(z_{t+1}^{\dagger}) &= u_t^H(T_t + DT_t d) + (u(z_{t+1}^{\dagger}) - u_t)^H T(z_{t+1}^{\dagger}) + O(\rho_2\|d\|^2) \\ &= \Lambda_t u_t^H d + O\big(\rho_2\|d\|^2 + \eta\rho_2\|d\|\|T_t\|\big) \\ &= \Lambda_t u_t^H d + O(\epsilon^2\|T_t\|). \end{aligned}$$

Here $u_t^H T_t = 0$ and $u_t^H DT_t = \Lambda_t u_t^H$. Since $\|P_t d\| = |u_t^H d|$ under the stated normalization and $|\Lambda_t| = \Omega(\eta^{-1})$, we conclude that

$$\|P_t d\| = O(\epsilon^2\eta\|T_t\|). \tag{51}$$

Projecting equation 50 onto $U_t^{\perp}$, and using $P_t T_t = 0$ and equation 51, now gives

$$d = -\eta T_t - \nu\nabla\phi_t^{\perp} + O(\epsilon^2\eta\|T_t\|).$$

Finally, expand the stability constraint:

$$\begin{aligned} \phi(z_{t+1}^{\dagger}) &= \phi(z_t^{\dagger}) + (\nabla\phi_t)^{\top} d + O(\eta\rho_2^2\|d\|^2) \\ &= \eta\alpha_t - \nu\|\nabla\phi_t^{\perp}\|^2 + O(\epsilon^2\eta\rho_2\|T_t\|) \\ &= \eta\alpha_t - \nu\|\nabla\phi_t^{\perp}\|^2 + O(\epsilon^2\eta\alpha_t). \end{aligned}$$

If $\nu = 0$, this would give $\phi(z_{t+1}^{\dagger}) = \eta\alpha_t(1 + O(\epsilon^2)) > 0$, contradicting feasibility. Hence $\nu > 0$, and complementary slackness in equation 49 implies $\phi(z_{t+1}^{\dagger}) = 0$, completing the induction. Solving the same expansion for ν gives

$$\nu = \frac{\eta\alpha_t}{\|\nabla\phi_t^{\perp}\|^2}(1 + O(\epsilon^2)).$$

Since $\alpha_t = -T_t^{\top}\nabla\phi_t^{\perp}$ and $P_t T_t = 0$, substitution into the update yields

$$\begin{aligned} z_{t+1}^{\dagger} &= z_t^{\dagger} - \eta T_t - \eta\frac{\alpha_t}{\|\nabla\phi_t^{\perp}\|^2}\nabla\phi_t^{\perp} + O(\epsilon^2\eta\|T_t\|) \\ &= z_t^{\dagger} - \eta\left(T_t - \frac{T_t^{\top}\nabla\phi_t^{\perp}}{\|\nabla\phi_t^{\perp}\|^2}\nabla\phi_t^{\perp}\right) + O(\epsilon^2\eta\|T_t\|) \\ &= z_t^{\dagger} - \eta P_{U_t,\nabla\phi_t}^{\perp} T_t + O(\epsilon^2\eta\|T_t\|), \end{aligned}$$

as claimed. □

Note that Assumption 6.3(3) and the definitions of $\delta_t, \delta, \epsilon_t, \epsilon$ in Definition 6.2 gives

$$\eta\rho_2\delta = \Theta(\epsilon), \qquad \eta\|T_t\| = O(\epsilon\delta), \qquad \eta^2\rho_2\|T_t\| = O(\epsilon^2). \tag{52}$$

Next, we provide the following estimation for the later use.

Lemma D.11 (Estimates along the constrained trajectory)**.** *Under the standing assumptions in Section D.3, write $a_t := 1 - \eta\Lambda_t$ and choose the phases of u_t coherently. For sufficiently small ϵ, the following estimates hold uniformly along the constrained trajectory, whenever the displayed time indices are defined.*

1. ***Reference variation.*** *The reference quantities satisfy*

$$\begin{aligned} \|u_{t+1} - u_t\| + \|P_{t+1} - P_t\| + |a_{t+1} - a_t| &= O(\epsilon^2), \\ \|D^2T_{t+1} - D^2T_t\| + \|\nabla\Lambda_{t+1} - \nabla\Lambda_t\| + \|\nabla\phi^\perp_{t+1} - \nabla\phi^\perp_t\| &= O(\epsilon^2\rho_2). \end{aligned} \tag{53}$$

   In particular, $\|u_{t+1}\| = \|u_t\|$, and the real or non-real case does not change between consecutive reference points.

2. ***Coefficient variation.*** *The scalar coefficient satisfies*

$$\left| \frac{\beta^\perp_{t+1}}{2\|\nabla\phi^\perp_{t+1}\|^2} - \frac{\beta^\perp_t}{2\|\nabla\phi^\perp_t\|^2} \right| = O(\epsilon^2). \tag{54}$$

   For $\Lambda_t \notin \mathbb{R}$,

$$\left\| \frac{D^2T_{t+1}[u_{t+1}, u_{t+1}]}{a_{t+1}^2 - 1} - \frac{D^2T_t[u_t, u_t]}{a_t^2 - 1} \right\| = O(\epsilon^2\rho_2). \tag{55}$$

   In both the real and non-real cases,

$$\left| \frac{\eta\, u_{t+1}^H D^2T_{t+1}[u_{t+1}, u_{t+1}]}{2\|u_{t+1}\|^3(a_{t+1}^2 - a_{t+1})} - \frac{\eta\, u_t^H D^2T_t[u_t, u_t]}{2\|u_t\|^3(a_t^2 - a_t)} \right| = O(\epsilon^2\eta\rho_2). \tag{56}$$

3. ***Changes in output projections.*** *For every $h \in \mathbb{R}^d$, with R_t as in Definition D.4,*

$$P^\perp_{t+1}(I - \eta DT_t)h = R_tP^\perp_t h + O(\epsilon^2\|h\|), \tag{57}$$

$$R_tP^\perp_t h = (I - \eta DT_t)P^\perp_t h + O(\epsilon^2\|h\|). \tag{58}$$

4. ***Tensor estimates from approximate normality.*** *For $\Lambda_t \notin \mathbb{R}$,*

$$|u_t^H D^2T_t[h, \bar{u}_t]| \le O(\epsilon\rho_2)\|h\|, \qquad h \in \mathbb{C}^d. \tag{59}$$

   In both cases,

$$\|P^\perp_t D^2T_t[u_t, w]\| \le O(\epsilon\rho_2)\|w\|, \qquad w \in U^\perp_t. \tag{60}$$

Proof. 1. Lemma D.10 and equation 52 imply

$$\|z^\dagger_{t+1} - z^\dagger_t\| = O(\eta\|T_t\|) = O\left(\frac{\epsilon^2}{\eta\rho_2}\right). \tag{61}$$

For sufficiently small ϵ, this step lies in $\mathcal{S}_t$. Applying Lemma D.7(3)–(5) and Assumption 6.3(2) to equation 61 gives the bounds for u_t, P_t, a_t, D^2T_t, and $\nabla\Lambda_t$ in equation 53. For the projected gradient, the same estimates give

$$\begin{aligned} \|\nabla\phi^\perp_{t+1} - \nabla\phi^\perp_t\| &\le \|\nabla\phi_{t+1} - \nabla\phi_t\| + \|P_{t+1} - P_t\|\|\nabla\phi_t\| \\ &= O(\epsilon^2\rho_2), \end{aligned}$$

using $\|\nabla\phi_t\| = O(\rho_2)$ from equation 46 and Lemma D.7(3). The normalization $\|u_t\| \in \{1, \sqrt{2}\}$ and equation 53 imply $\|u_{t+1}\| = \|u_t\|$; thus the real or non-real case is unchanged.

2. By equation 53, Assumptions 6.3(2) and 6.3(3), and the definition of $\beta_t^\perp$ in Definition 6.2,

$$\begin{aligned}|\beta_{t+1}^\perp - \beta_t^\perp| &\le O(\rho_2)\big(\|\nabla\phi_{t+1}^\perp - \nabla\phi_t^\perp\| + \|D^2T_{t+1} - D^2T_t\|\big)\\ &\quad + O(\rho_2^2)\|u_{t+1} - u_t\|\\ &= O(\epsilon^2\rho_2^2),\\ \big|\|\nabla\phi_{t+1}^\perp\|^2 - \|\nabla\phi_t^\perp\|^2\big| &\le \big(\|\nabla\phi_{t+1}^\perp\| + \|\nabla\phi_t^\perp\|\big)\|\nabla\phi_{t+1}^\perp - \nabla\phi_t^\perp\|\\ &= O(\epsilon^2\rho_2^2).\end{aligned} \tag{62}$$

Subtracting the two scalar quotients gives

$$\begin{aligned}&\frac{\beta_{t+1}^\perp}{2\|\nabla\phi_{t+1}^\perp\|^2} - \frac{\beta_t^\perp}{2\|\nabla\phi_t^\perp\|^2}\\ &\quad= \frac{\beta_{t+1}^\perp - \beta_t^\perp}{2\|\nabla\phi_{t+1}^\perp\|^2} + \frac{\beta_t^\perp\big(\|\nabla\phi_t^\perp\|^2 - \|\nabla\phi_{t+1}^\perp\|^2\big)}{2\|\nabla\phi_{t+1}^\perp\|^2\|\nabla\phi_t^\perp\|^2} = O(\epsilon^2),\end{aligned} \tag{63}$$

where Assumption 6.3(3) and equation 62 bound the numerators and denominators. This proves equation 54.

Lemma D.10 gives $|a_t| = 1$. Together with Assumption 6.3(1) and the standing condition $\Lambda_t \neq 0$, this yields

$$|a_t^2 - a_t| = \eta|\Lambda_t| = \Omega(1), \tag{64}$$

$$|a_t^2 - 1| = 2\eta|\operatorname{Im}\Lambda_t| > \Delta \qquad (\Lambda_t \notin \mathbb{R}). \tag{65}$$

For the first bound, use $\eta\Lambda_t = 2$ in the real case and $\eta|\Lambda_t| \ge \eta|\operatorname{Im}\Lambda_t| > \Delta/2$ otherwise. The tensor numerator and the denominator of equation 55 satisfy

$$\begin{aligned}&\|D^2T_{t+1}[u_{t+1},u_{t+1}] - D^2T_t[u_t,u_t]\|\\ &\quad\le O(1)\|D^2T_{t+1} - D^2T_t\| + O(\rho_2)\|u_{t+1} - u_t\| = O(\epsilon^2\rho_2),\\ &|a_{t+1}^2 - a_t^2| \le 2|a_{t+1} - a_t| = O(\epsilon^2),\end{aligned} \tag{66}$$

by equation 53, Assumption 6.3(2), and $|a_t| = 1$. The same quotient subtraction as in equation 63, using equation 65 and equation 66, bounds the difference in equation 55 by

$$\frac{O(\epsilon^2\rho_2)}{\Delta} + \frac{O(\rho_2)O(\epsilon^2)}{\Delta^2} = O(\epsilon^2\rho_2).$$

For the coefficient in equation 56, equation 53 and Assumption 6.3(2) give

$$\begin{aligned}&\big|u_{t+1}^H D^2T_{t+1}[u_{t+1},u_{t+1}] - u_t^H D^2T_t[u_t,u_t]\big|\\ &\quad\le O(1)\big(\|D^2T_{t+1} - D^2T_t\| + \rho_2\|u_{t+1} - u_t\|\big) = O(\epsilon^2\rho_2),\\ &|(a_{t+1}^2 - a_{t+1}) - (a_t^2 - a_t)| \le 3|a_{t+1} - a_t| = O(\epsilon^2).\end{aligned} \tag{67}$$

Since $\|u_{t+1}\| = \|u_t\|$, the quotient subtraction in equation 63, now using equation 64 and equation 67, bounds the difference in equation 56 by

$$O(\eta)O(\epsilon^2\rho_2) + O(\eta\rho_2)O(\epsilon^2) = O(\epsilon^2\eta\rho_2).$$

3. Normality in Assumption 6.1 implies that P_t commutes with $I - \eta DT_t$. Using the definition of R_t in Definition D.4, we therefore obtain

$$\begin{aligned}P_{t+1}^\perp(I - \eta DT_t)h - R_tP_t^\perp h &= (P_t - P_{t+1})(I - \eta DT_t)P_th,\\ R_tP_t^\perp h - (I - \eta DT_t)P_t^\perp h &= (P_t - P_{t+1})(I - \eta DT_t)P_t^\perp h.\end{aligned} \tag{68}$$

Combining equation 68 with equation 53 and equation 44 proves equation 57 and equation 58.

4. Suppress t and take h real. The product rule gives

$$\begin{aligned}D\big(DT(DT)^\top - (DT)^\top DT\big)[h,v] &= D^2T[h,(DT)^\top v] + DT\,D((DT)^\top)[h,v]\\ &\quad - D((DT)^\top)[h,DTv] - (DT)^\top D^2T[h,v].\end{aligned} \tag{69}$$

Normality in Assumption 6.1 implies

$$\begin{aligned} (DT)^\top \bar{u} &= \Lambda \bar{u}, & DT\bar{u} &= \bar{\Lambda}\bar{u}, \\ u^H DT &= \Lambda u^H, & u^H (DT)^\top &= \bar{\Lambda} u^H. \end{aligned} \tag{70}$$

Also,

$$u^H D((DT)^\top)[h,\bar{u}] = \left(D^2T[h,\bar{u}]\right)^\top \bar{u} = u^H D^2T[h,\bar{u}].$$

Using this identity and equation 70 in equation 69, with $v = \bar{u}$, gives

$$\begin{aligned} &u^H D\big(DT(DT)^\top - (DT)^\top DT\big)\,[h,\bar{u}] \\ &\quad = \Lambda u^H D^2T[h,\bar{u}] + \Lambda u^H D((DT)^\top)[h,\bar{u}] \\ &\qquad - \bar{\Lambda} u^H D((DT)^\top)[h,\bar{u}] - \bar{\Lambda} u^H D^2T[h,\bar{u}] \\ &\quad = 2(\Lambda - \bar{\Lambda}) u^H D^2T[h,\bar{u}]. \end{aligned}$$

For non-real Λ, Assumptions 6.4(1) and 6.3(1) therefore imply

$$\begin{aligned} |u^H D^2T[h,\bar{u}]| &= \frac{\left|u^H D\big(DT(DT)^\top - (DT)^\top DT\big)\,[h,\bar{u}]\right|}{2|\Lambda - \bar{\Lambda}|} \\ &\le \frac{O(\epsilon\rho_2/\eta)\|h\|}{2\Delta/\eta} = O(\epsilon\rho_2)\|h\|. \end{aligned}$$

Complex linearity extends this estimate to complex h, proving equation 59.

For equation 60, take $w \in U^\perp$ real. Apply equation 69 with $h = w$ and $v = u$, project onto $U^\perp$, and use Assumption 6.4(1):

$$\begin{aligned} &(\bar{\Lambda} I - (DT|_{U^\perp})^H) P^\perp D^2T[w,u] \\ &\quad + (DT|_{U^\perp} - \Lambda I) P^\perp D((DT)^\top)[w,u] = O(\epsilon\rho_2/\eta)\|w\|. \end{aligned} \tag{71}$$

Assumptions 6.1 and 6.3(1) imply

$$\begin{aligned} \|(\bar{\Lambda} I - (DT|_{U^\perp})^H)^{-1}\| &= O(\eta), \\ \|(\bar{\Lambda} I - (DT|_{U^\perp})^H)^{-1}(DT|_{U^\perp} - \Lambda I)\| &\le 1. \end{aligned} \tag{72}$$

Using equation 72 in equation 71, and then applying Assumption 6.3(4) to $D^2T[w,h]$, gives

$$\begin{aligned} \|P^\perp D^2T[w,u]\| &\le \|P^\perp D((DT)^\top)[w,u]\| + O(\epsilon\rho_2)\|w\| \\ &= \sup_{\substack{h\in\mathbb{C}^d,\ Ph=0 \\ \|h\|\le 1}} |u^\top D^2T[w,h]| + O(\epsilon\rho_2)\|w\| \\ &= O(\epsilon\rho_2)\|w\|. \end{aligned}$$

Here Assumption 6.3(4) is extended to complex h by complex linearity. This proves equation 60.

□

D.5 Proofs of Theorem

We first bound the predicted deviation.

Lemma D.12. *If $P_0^{\perp} v_0 = 0$ and $\|v_0\| = O(\delta)$, then $\|v_t^\star\| = O(\delta)$ for all $t \le \mathcal{T}$.*

Proof. Define

$$w_t := P_t^{\perp} v_t^\star, \qquad b_t := \frac{\beta_t^{\perp}}{2\|\nabla \phi_t^{\perp}\|^2}, \qquad Q_t := \|w_t\|^2 + b_t |x_t^\star|^2. \tag{73}$$

Normality and the definition of $x_t^\star$ give

$$\|v_t^\star\|^2 = \|w_t\|^2 + \|u_t\|^2 |x_t^\star|^2.$$

By Assumption 6.3(3), $b_t = \Theta(1)$, so $Q_t = \Theta(\|v_t^\star\|^2)$ and bounding Q_t proves the lemma. For the contribution of $D^2 T_t[u_t, u_t]$ in Definition D.2, define

$$c_t := \begin{cases} \dfrac{D^2 T_t[u_t, u_t]}{(1-\eta\Lambda_t)^2 - 1}, & \Lambda_t \notin \mathbb{R}, \\ 0, & \Lambda_t \in \mathbb{R}, \end{cases} \tag{74}$$

$$H_t := \eta \operatorname{Re}\left((x_t^\star)^2 w_t^\top c_t\right), \qquad \mathcal{E}_t := Q_t + H_t. \tag{75}$$

We will show that $\mathcal{E}_t$ is comparable to Q_t and satisfies an inequality that yields the required bound by induction. Lemma D.10 gives $|1 - \eta\Lambda_t| = 1$. Together with Assumptions 6.3(1) and 6.3(2), this bounds the coefficient in equation 74 and the correction in equation 75 as

$$\|c_t\| = O(\rho_2), \qquad |H_t| \le O(\eta\rho_2)\|v_t^\star\|^3. \tag{76}$$

Bounding b_t and c_t. Applying Lemma D.11(2) to equation 73 and equation 74 gives

$$|b_{t+1} - b_t| = O(\epsilon^2), \qquad \|c_{t+1} - c_t\| = O(\epsilon^2 \rho_2). \tag{77}$$

The second bound is immediate in the real case, where $c_t = c_{t+1} = 0$.

Bounding $R_t w_t - w_t$. The definition of R_t in Definition D.4 and equation 44 give

$$\|R_t\| \le \|I - \eta D T_t\| \le 1 + O(\epsilon). \tag{78}$$

Using equation 44, equation 43, and the quadratic-form bound in Assumption 6.3(4), we obtain

$$\begin{aligned} \eta^2 \|D T_t w_t\|^2 &= \|(I - \eta D T_t) w_t\|^2 - \|w_t\|^2 + 2\eta \langle w_t, D T_t w_t \rangle \\ &\le O(\epsilon)\|w_t\|^2. \end{aligned}$$

Combining this with equation 58, applied to $h = w_t$, gives

$$\|R_t w_t - w_t\| = O(\sqrt{\epsilon})\|w_t\|. \tag{79}$$

Prove by Induction. Assume that, for some fixed $M \ge 1$,

$$\|v_s^\star\| \le M\delta \qquad (0 \le s \le t). \tag{80}$$

Here M is independent of ϵ and will be chosen at the end of the proof. All implicit constants below are independent of M.

Definition D.2, Lemma D.7(3), and Assumption 6.3(3) give

$$\eta |y_t^\star| = O(\epsilon M), \qquad \|w_{t+1} - R_t w_t\| = O(\epsilon M^2 \delta). \tag{81}$$

The two output components are orthogonal, and the critical component changes in norm by the factor $|1 + \eta y_t^\star|$. Using equation 78, equation 79, and equation 81, we obtain, for sufficiently small ϵ depending on M,

$$\begin{aligned} \|v_{t+1}^\star\| &\le (1 + O(\epsilon M))\|v_t^\star\| + O(\epsilon M^2 \delta) \le 2M\delta, \\ \|w_{t+1} - w_t\| &\le O(\sqrt{\epsilon}\, M\delta + \epsilon M^2 \delta). \end{aligned} \tag{82}$$

Together with equation 80, this gives $\|v_s^\star\| \le 2M\delta$ for $0 \le s \le t+1$. Equation equation 76 and Assumption 6.3(3) consequently imply $|H_s| \le O(\epsilon M)Q_s$ throughout this range, and hence

$$\frac{1}{2}Q_s \le \mathcal{E}_s \le \frac{3}{2}Q_s \qquad (0 \le s \le t+1). \tag{83}$$

We first consider the non-real case. Definition D.2 and equation 46 give $\operatorname{Re} y_t^\star = (\nabla\phi_t^\perp)^\top w_t$. Using this identity, $|1-\eta\Lambda_t| = 1$ from Lemma D.10, equation 77, equation 81, and equation 82, we expand equation 75 as

$$\begin{aligned}
\mathcal{E}_{t+1} - \mathcal{E}_t &= \|w_{t+1}\|^2 - \|w_t\|^2 + b_{t+1}|x_{t+1}^\star|^2 - b_t|x_t^\star|^2 + H_{t+1} - H_t \\
&= \|R_t w_t\|^2 - \|w_t\|^2 + 2\eta b_t \delta_t^2 (\nabla\phi_t^\perp)^\top R_t w_t \\
&\quad + 2\eta b_t |x_t^\star|^2 (\nabla\phi_t^\perp)^\top (w_t - R_t w_t) \\
&\quad - \eta \operatorname{Re}\big((x_t^\star)^2 (R_t w_t)^\top D^2 T_t[u_t, u_t]\big) \\
&\quad + \eta \operatorname{Re}\big(\big((x_{t+1}^\star)^2 - (x_t^\star)^2\big) w_t^\top c_t\big) + O(\epsilon^{3/2} M^4 \delta^2) \\
&= \|R_t w_t\|^2 - \|w_t\|^2 + 2\eta b_t \delta_t^2 (\nabla\phi_t^\perp)^\top R_t w_t \\
&\quad + 2\eta b_t |x_t^\star|^2 (\nabla\phi_t^\perp)^\top (w_t - R_t w_t) \\
&\quad + \eta \operatorname{Re}\big((x_t^\star)^2 (w_t - R_t w_t)^\top D^2 T_t[u_t, u_t]\big) + O(\epsilon^{3/2} M^4 \delta^2).
\end{aligned} \tag{84}$$

In the second equality of equation 84, the output projectors disappear because $R_t w_t \in U_{t+1}^\perp$. The changes in w_t and c_t within H_t are bounded by equation 82 and equation 77; their contribution is $O(\epsilon^{3/2} M^4 \delta^2)$. The squares of the perpendicular update correction and of $\eta y_t^\star$, together with the change in b_t, contribute $O(\epsilon^2 M^4 \delta^2)$ by equation 81 and equation 77. The third equality of equation 84 uses the critical recurrence from Definition D.2: $(x_{t+1}^\star)^2 = (1-\eta\Lambda_t)^2 (x_t^\star)^2 + O(\epsilon M^3 \delta^2)$, and the identity $((1-\eta\Lambda_t)^2 - 1)c_t = D^2 T_t[u_t, u_t]$ from equation 74.

We now bound the final expression in equation 84. By equation 78, the difference $\|R_t w_t\|^2 - \|w_t\|^2$ is at most $O(\epsilon)\|w_t\|^2$. The term $2\eta b_t \delta_t^2 (\nabla\phi_t^\perp)^\top R_t w_t$ is bounded by $O(\eta\rho_2\delta^2)\|w_t\| \le O(\epsilon)(Q_t + \delta^2)$, using Assumption 6.3(3), equation 73, and $2\delta\|w_t\| \le \delta^2 + \|w_t\|^2$. The two terms containing $w_t - R_t w_t$ in equation 84 are $O(\epsilon^{3/2} M^3 \delta^2)$ by equation 79. Consequently,

$$\begin{aligned}
\mathcal{E}_{t+1} - \mathcal{E}_t &\le O(\epsilon)(Q_t + \delta^2) + O(\epsilon^{3/2} M^4 \delta^2) \\
&\le O(\epsilon)\mathcal{E}_t + O(\epsilon\delta^2) + O(\epsilon^{3/2} M^4 \delta^2),
\end{aligned} \tag{85}$$

where the last inequality follows from equation 83. In the real case, $H_t = 0$ by equation 75, and the $D^2 T_t[u_t, u_t]$ oscillatory term in Definition D.2 is absent, so equation 85 also holds.

Since $w_0 = 0$, we have $H_0 = 0$ and $\mathcal{E}_0 = b_0|x_0^\star|^2 = O(\delta^2)$. Under equation 80, equation 85 holds at every step from 0 through t. Iteration and $t+1 \le \mathcal{T} = O(\epsilon^{-1})$ therefore give

$$\begin{aligned}
\mathcal{E}_{t+1} &\le e^{O(\epsilon(t+1))} \left[\mathcal{E}_0 + O((t+1)\epsilon\delta^2) + O((t+1)\epsilon^{3/2} M^4 \delta^2)\right] \\
&\le O(\delta^2) + O(\sqrt{\epsilon}\, M^4 \delta^2).
\end{aligned}$$

By equation 83, the induction hypothesis thus implies

$$\|v_{t+1}^\star\|^2 \le O(\delta^2) + O(\sqrt{\epsilon}\, M^4 \delta^2).$$

Choose M large enough that $\|v_0\| \le M\delta$ and the first term is at most $M^2\delta^2/4$. For this fixed M, sufficiently small ϵ makes the second term at most $M^2\delta^2/4$ and ensures all preceding local bounds. Consequently,

$$\|v_{t+1}^\star\|^2 \le \frac{1}{2} M^2 \delta^2 < M^2 \delta^2.$$

This proves the induction step; the base case holds by the choice of M. Hence $\|v_t^\star\| \le M\delta = O(\delta)$ for all $t \le \mathcal{T}$. □

Lemma D.13. *There exist corrected deviations $\widetilde{v}_t$, given explicitly by equation 103, such that every step $0 \le t < \mathcal{T}$ with $\|v_t\| = O(\delta)$ satisfies*

$$\|\widetilde{v}_t - v_t\| + \|\widetilde{v}_{t+1} - v_{t+1}\| = O(\epsilon\delta),$$
$$\|\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t)\| = O(\epsilon^2\delta).$$

These estimates are uniform in t when the implicit constant in the input bound $\|v_t\| = O(\delta)$ is fixed.

Proof. Fix a step with $\|v_t\| = O(\delta)$ and write $a_t := 1 - \eta\Lambda_t$ as in Lemma D.11. Lemma D.10 gives $|a_t| = 1$; we also use Assumption 6.3(3). All implicit constants are uniform in t and independent of ϵ, with the input bound and the gap constant $\Delta > 0$ fixed.

The actual update and Lemma D.10 give

$$z_{t+1} = z_t - \eta T(z_t),$$
$$z_{t+1}^\dagger = z_t^\dagger - \eta P^\perp_{U_t,\nabla\phi_t} T_t + O(\epsilon^2\eta\|T_t\|).$$

Subtracting these updates and using $z_t = z_t^\dagger + v_t$ yields

$$\begin{aligned} v_{t+1} &= v_t - \eta T(z_t^\dagger + v_t) + \eta P^\perp_{U_t,\nabla\phi_t} T_t + O(\epsilon^2\eta\|T_t\|) \\ &= v_t - \eta\big(T(z_t^\dagger + v_t) - T_t\big) + \eta(P^\perp_{U_t,\nabla\phi_t} T_t - T_t) + O(\epsilon^2\eta\|T_t\|). \end{aligned} \tag{86}$$

Feasibility gives $P_t T_t = 0$ and $T_t^\top \nabla\phi_t^\perp = T_t^\top \nabla\phi_t = -\alpha_t$. The projector identity

$$P^\perp_{U_t,\nabla\phi_t} = P_t^\perp - \frac{\nabla\phi_t^\perp(\nabla\phi_t^\perp)^\top}{\|\nabla\phi_t^\perp\|^2} \tag{87}$$

therefore gives

$$P^\perp_{U_t,\nabla\phi_t} T_t - T_t = -\frac{T_t^\top\nabla\phi_t^\perp}{\|\nabla\phi_t^\perp\|^2}\nabla\phi_t^\perp = \frac{\alpha_t}{\|\nabla\phi_t^\perp\|^2}\nabla\phi_t^\perp. \tag{88}$$

Taylor expansion at $z_t^\dagger$ gives

$$T(z_t^\dagger + v_t) - T_t = DT_t v_t + \frac{1}{2}D^2T_t[v_t, v_t] + O(\rho_3\delta^3). \tag{89}$$

Assumptions 6.3(2) and 6.3(3) imply

$$\eta\rho_3\delta^3 = O\big((\eta\rho_2\delta)^2\delta\big) = O(\epsilon^2\delta),$$
$$\epsilon^2\eta\|T_t\| = O(\epsilon^3\delta) = O(\epsilon^2\delta). \tag{90}$$

Substituting equation 88 and equation 89 into equation 86, with the remainders bounded by equation 90, gives

$$v_{t+1} = (I - \eta DT_t)v_t + \eta\frac{\alpha_t}{\|\nabla\phi_t^\perp\|^2}\nabla\phi_t^\perp - \frac{\eta}{2}D^2T_t[v_t, v_t] + O(\epsilon^2\delta). \tag{91}$$

By equation 91, equation 44, and Assumptions 6.3(2) and 6.3(3),

$$\begin{aligned} \|v_{t+1}\| &\le \|I - \eta DT_t\|\|v_t\| + \frac{\eta\alpha_t}{\|\nabla\phi_t^\perp\|} + \frac{\eta\rho_2}{2}\|v_t\|^2 + O(\epsilon^2\delta) \\ &= O(\delta) + O(\eta\|T_t\|) + O(\eta\rho_2\delta^2) + O(\epsilon^2\delta) = O(\delta), \end{aligned} \tag{92}$$

where the last equality uses Assumption 6.3(3). By equation 53, the fixed normalization of u_t, and equation 92,

$$\left|x_{t+1} - \frac{u_t^H v_{t+1}}{\|u_t\|}\right| \le O(\epsilon^2)\|v_{t+1}\| = O(\epsilon^2\delta). \tag{93}$$

The definition of x_t gives the decomposition

$$v_t = \begin{cases} x_t u_t + v_t^\perp, & \Lambda_t \in \mathbb{R}, \\ (x_t u_t + \bar{x}_t\bar{u}_t)/\sqrt{2} + v_t^\perp, & \Lambda_t \notin \mathbb{R}. \end{cases} \tag{94}$$

Normality in Assumption 6.1 and $u_t^H \nabla \phi_t^\perp = 0$ imply

$$\frac{u_t^H (I - \eta D T_t) v_t}{\|u_t\|} = (1 - \eta \Lambda_t) \frac{u_t^H v_t}{\|u_t\|} = a_t x_t,$$
$$\frac{u_t^H}{\|u_t\|} \left(\eta \frac{\alpha_t}{\|\nabla \phi_t^\perp\|^2} \nabla \phi_t^\perp \right) = 0.$$

Applying these identities to equation 91, using equation 93, and expanding according to equation 94, we obtain

$$\begin{aligned} x_{t+1} &= a_t x_t - \frac{\eta}{2\|u_t\|} u_t^H D^2 T_t[v_t, v_t] + O(\epsilon^2 \delta) \\ &= a_t x_t - \frac{\eta x_t^2}{2\|u_t\|^3} u_t^H D^2 T_t[u_t, u_t] + \underbrace{\left(-\frac{\eta x_t}{\|u_t\|^2} u_t^H D^2 T_t[u_t, v_t^\perp] \right)}_{= \eta a_t \overline{y_t} x_t} \\ &\quad - \frac{\eta}{2\|u_t\|} u_t^H D^2 T_t[v_t^\perp, v_t^\perp] \\ &\quad - \frac{\eta \mathbf{1}_{\{\Lambda_t \notin \mathbb{R}\}}}{2\|u_t\|} u_t^H \Big(\frac{\bar{x}_t^2}{\|u_t\|^2} D^2 T_t[\bar{u}_t, \bar{u}_t] + \frac{2|x_t|^2}{\|u_t\|^2} D^2 T_t[u_t, \bar{u}_t] \\ &\qquad\qquad + \frac{2\bar{x}_t}{\|u_t\|} D^2 T_t[v_t^\perp, \bar{u}_t] \Big) + O(\epsilon^2 \delta). \end{aligned} \tag{95}$$

We treat the four contributions in equation 95 in order. The identity under the brace is proved in equation 98.

- **The x_t^2 term.** Assumptions 6.3(2) and 6.3(3) give

$$\left| \frac{\eta x_t^2}{2\|u_t\|^3} u_t^H D^2 T_t[u_t, u_t] \right| \le O(\eta \rho_2) |x_t|^2 = O(\epsilon \delta). \tag{96}$$

  We retain this contribution in equation 102.

- **The $x_t \overline{y_t}$ term.** The definition of y_t and $|a_t| = 1$ give

$$\begin{aligned} y_t &= -a_t (\nabla \Lambda_t^\perp)^H v_t, \\ \overline{y_t} &= -\bar{a}_t (\nabla \Lambda_t)^\top v_t^\perp, \\ (\nabla \Lambda_t)^\top v_t^\perp &= -a_t \overline{y_t}. \end{aligned} \tag{97}$$

  The eigenvalue derivative formula equation 45 at the normal base point and equation 97 yield

$$\begin{aligned} -\frac{\eta x_t}{\|u_t\|^2} u_t^H D^2 T_t[u_t, v_t^\perp] &= -\eta x_t (\nabla \Lambda_t)^\top v_t^\perp \\ &= \eta a_t \overline{y_t} x_t. \end{aligned} \tag{98}$$

  Lemma D.7(3) and Assumption 6.3(3) also give

$$\begin{aligned} \eta |y_t| &\le O(\eta \rho_2) \|v_t^\perp\| = O(\epsilon), \\ |\eta a_t \overline{y_t} x_t| &= O(\epsilon \delta). \end{aligned} \tag{99}$$

- **The $D^2 T_t[v_t^\perp, v_t^\perp]$ term.** Assumptions 6.3(4) and 6.3(3) give

$$\left| \frac{\eta}{2\|u_t\|} u_t^H D^2 T_t[v_t^\perp, v_t^\perp] \right| \le O(\eta \epsilon \rho_2) \|v_t^\perp\|^2 = O(\epsilon^2 \delta). \tag{100}$$

- **The $\bar{x}_t^2$, $|x_t|^2$, and $\bar{x}_t$ terms.** These are the three terms inside parentheses in equation 95, present only when $\Lambda_t \notin \mathbb{R}$. Lemma D.11(4) provides equation 59. Applying equation 59

with $h = \bar{u}_t, u_t, v_t^{\perp}$, respectively, and using Assumption 6.3(3), we obtain

$$\begin{aligned}
\frac{\eta|x_t|^2}{2\|u_t\|^3}|u_t^H D^2 T_t[\bar{u}_t, \bar{u}_t]| &\le O(\eta\epsilon\rho_2)|x_t|^2 = O(\epsilon^2\delta),\\
\frac{\eta|x_t|^2}{\|u_t\|^3}|u_t^H D^2 T_t[u_t, \bar{u}_t]| &\le O(\eta\epsilon\rho_2)|x_t|^2 = O(\epsilon^2\delta),\\
\frac{\eta|x_t|}{\|u_t\|^2}|u_t^H D^2 T_t[v_t^{\perp}, \bar{u}_t]| &\le O(\eta\epsilon\rho_2)|x_t|\|v_t^{\perp}\| = O(\epsilon^2\delta).
\end{aligned} \tag{101}$$

Substituting equation 98, equation 100, and equation 101 into equation 95 gives

$$x_{t+1} = a_t(1+\eta\overline{y_t})x_t - \frac{\eta}{2\|u_t\|^3}u_t^H D^2 T_t[u_t, u_t]x_t^2 + O(\epsilon^2\delta). \tag{102}$$

Define

$$\widetilde{v}_s := v_s + \|u_s\|\,\mathrm{Re}(r_s x_s^2 u_s), \qquad 0 \le s \le \mathcal{T}, \tag{103}$$

where

$$r_s := \frac{\eta\, u_s^H D^2 T_s[u_s, u_s]}{2\|u_s\|^3(a_s^2 - a_s)}. \tag{104}$$

Lemma D.10 and Assumption 6.3(1) give $|a_s^2 - a_s| = \Omega(1)$, using $\Lambda_s \neq 0$ in the real case. Thus Assumption 6.3(2) and equation 104 imply

$$|r_s| = O(\eta\rho_2). \tag{105}$$

The input bound and equation 92, together with equation 105 and Assumption 6.3(3), give

$$\|\widetilde{v}_s - v_s\| \le 2|r_s||x_s|^2 = O(\eta\rho_2\delta^2) = O(\epsilon\delta), \qquad s = t, t+1. \tag{106}$$

By equation 103 and the normalization of u_s,

$$\begin{aligned}
\widetilde{x}_s := \frac{u_s^H\widetilde{v}_s}{\|u_s\|} &= x_s + r_s x_s^2,\\
P_s^{\perp}\widetilde{v}_s &= v_s^{\perp},\\
-a_s(\nabla\Lambda_s^{\perp})^H\widetilde{v}_s &= y_s.
\end{aligned} \tag{107}$$

It remains to prove equation 119. Using Definition D.2, equation 102, and equation 107, we obtain

$$\begin{aligned}
&\frac{u_{t+1}^H}{\|u_{t+1}\|}\big(\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t)\big)\\
&\quad= x_{t+1} + r_{t+1}x_{t+1}^2 - a_t(1+\eta\overline{y_t})(x_t + r_t x_t^2)\\
&\quad= \underbrace{\left[-\frac{\eta}{2\|u_t\|^3}u_t^H D^2 T_t[u_t, u_t] + r_t(a_t^2 - a_t)\right]}_{=0 \text{ by equation } 104} x_t^2\\
&\qquad+ (r_{t+1} - r_t)x_{t+1}^2 + r_t(x_{t+1}^2 - a_t^2 x_t^2) - \eta a_t\overline{y_t}r_t x_t^2 + O(\epsilon^2\delta).
\end{aligned} \tag{108}$$

Equations equation 102, equation 96, and equation 99 imply

$$\begin{aligned}
x_{t+1} &= a_t x_t + O(\epsilon\delta),\\
x_{t+1}^2 - a_t^2 x_t^2 &= (x_{t+1} - a_t x_t)(x_{t+1} + a_t x_t) = O(\epsilon\delta^2).
\end{aligned} \tag{109}$$

For the coefficient defined in equation 104, Lemma D.11(2) also gives

$$|r_{t+1} - r_t| = O(\epsilon^2\eta\rho_2). \tag{110}$$

Combining equation 105, equation 110, equation 109, and equation 99 with Assumption 6.3(3), we obtain

$$\begin{aligned}
|(r_{t+1} - r_t)x_{t+1}^2| &= O(\epsilon^3\delta),\\
|r_t(x_{t+1}^2 - a_t^2 x_t^2)| &= O(\epsilon^2\delta),\\
|\eta a_t\overline{y_t}r_t x_t^2| &= O(\epsilon^2\delta).
\end{aligned} \tag{111}$$

Substituting equation 111 into equation 108 proves

$$\frac{u_{t+1}^H}{\|u_{t+1}\|}\big(\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t)\big) = O(\epsilon^2\delta). \tag{112}$$

We next estimate $P_{t+1}^\perp(\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t))$. Lemma D.11(4) provides the mixed estimate equation 60.

For the expansion of $P_t^\perp D^2 T_t[v_t, v_t]$ using equation 94, Assumption 6.3(4), equation 60, and complex conjugation give

$$\begin{aligned}
\|P_t^\perp D^2 T_t[v_t^\perp, v_t^\perp]\| &\le O(\epsilon\rho_2)\|v_t^\perp\|^2 = O(\epsilon\rho_2\delta^2),\\
\frac{2|x_t|}{\|u_t\|}\|P_t^\perp D^2 T_t[u_t, v_t^\perp]\| &\le O(\epsilon\rho_2)|x_t|\|v_t^\perp\| = O(\epsilon\rho_2\delta^2),\\
\frac{2|\bar{x}_t|}{\|u_t\|}\|P_t^\perp D^2 T_t[\bar{u}_t, v_t^\perp]\| &= \frac{2|x_t|}{\|u_t\|}\|P_t^\perp D^2 T_t[u_t, v_t^\perp]\| = O(\epsilon\rho_2\delta^2).
\end{aligned} \tag{113}$$

The third line is needed only when $\Lambda_t \notin \mathbb{R}$. Expanding with equation 94 and using equation 113, we obtain the first equality in

$$\begin{aligned}
P_t^\perp D^2 T_t[v_t, v_t] &= \begin{cases} x_t^2 P_t^\perp D^2 T_t[u_t, u_t], & \Lambda_t \in \mathbb{R},\\ |x_t|^2 P_t^\perp D^2 T_t[u_t, \bar{u}_t] + P_t^\perp \mathrm{Re}(x_t^2 D^2 T_t[u_t, u_t]), & \Lambda_t \notin \mathbb{R}\end{cases} + O(\epsilon\rho_2\delta^2)\\
&= \frac{\beta_t^\perp |x_t|^2}{\|\nabla\phi_t^\perp\|^2}\nabla\phi_t^\perp + P_t^\perp \mathrm{Re}(x_t^2 D^2 T_t[u_t, u_t])\mathbf{1}_{\{\Lambda_t\notin\mathbb{R}\}} + O(\epsilon\rho_2\delta^2).
\end{aligned} \tag{114}$$

The second equality in equation 114 uses Assumption 6.4(2), the definition of $\beta_t^\perp$ in Definition 6.2, and equation 87. Applying equation 57 to $h = v_t$ gives

$$P_{t+1}^\perp(I - \eta DT_t)v_t = R_t v_t^\perp + O(\epsilon^2\delta). \tag{115}$$

Substitute equation 114 and equation 115 into equation 91, using equation 53 to change the tensor projection from $P_t^\perp$ to $P_{t+1}^\perp$. With $\alpha_t = \beta_t^\perp\delta_t^2/2$ from Definition 6.2 and the scaling in Assumption 6.3(3), this yields

$$\begin{aligned}
P_{t+1}^\perp v_{t+1} &= R_t v_t^\perp + \frac{\eta\beta_t^\perp}{2\|\nabla\phi_t^\perp\|^2}P_{t+1}^\perp\nabla\phi_t^\perp(\delta_t^2 - |x_t|^2)\\
&\quad - \frac{\eta}{2}P_{t+1}^\perp \mathrm{Re}(x_t^2 D^2 T_t[u_t, u_t])\mathbf{1}_{\{\Lambda_t\notin\mathbb{R}\}} + O(\epsilon^2\delta)\\
&= P_{t+1}^\perp \mathrm{step}_t(v_t) + O(\epsilon^2\delta),
\end{aligned} \tag{116}$$

where the last equality is Definition D.2. By equation 107, replacing v_t by $\widetilde{v}_t$ leaves $v_t^\perp$ fixed and replaces $|x_t|^2, x_t^2$ by $|\widetilde{x}_t|^2, \widetilde{x}_t^2$ in the $P_{t+1}^\perp$ component of Definition D.2. Assumptions 6.3(2) and 6.3(3), together with equation 106, give

$$\begin{aligned}
&\|P_{t+1}^\perp(\mathrm{step}_t(\widetilde{v}_t) - \mathrm{step}_t(v_t))\|\\
&\quad\le O(\eta\rho_2)\left(\big||\widetilde{x}_t|^2 - |x_t|^2\big| + |\widetilde{x}_t^2 - x_t^2|\right)\\
&\quad\le O(\eta\rho_2\delta)|\widetilde{x}_t - x_t| = O(\epsilon^2\delta).
\end{aligned} \tag{117}$$

Using equation 107 at $s = t+1$ and combining equation 116 with equation 117, we obtain

$$P_{t+1}^\perp(\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t)) = O(\epsilon^2\delta). \tag{118}$$

Finally, equation 112 and equation 118, together with orthogonality, imply

$$\widetilde{v}_{t+1} = \mathrm{step}_t(\widetilde{v}_t) + O(\epsilon^2\delta), \tag{119}$$

as required. □

Lemma D.14. *If $P_0^\perp v_0 = 0$ and $\|v_0\| = O(\delta)$, then*

$$\|v_t - v_t^\star\| = O(\epsilon\delta), \qquad t \le \mathcal{T}.$$

Proof. Fix $H > 0$ independent of ϵ, with $\epsilon\mathcal{T} \le H$. By Lemma D.12, choose $C_\star$ such that $\|v_t^\star\| \le C_\star\delta$ for all $t \le \mathcal{T}$. Use the corrected deviations from Lemma D.13 and let $a_t := 1 - \eta\Lambda_t$. Lemma D.10 gives $|a_t| = 1$, and Assumption 6.3(3) gives $\eta\rho_2\delta = O(\epsilon)$. All implicit constants are uniform in the time indices and independent of ϵ, with H, $C_\star$, and $\Delta > 0$ fixed.

For each update, split the error into three terms:

$$\begin{aligned} v_{t+1} - v_{t+1}^\star = {} & \underbrace{v_{t+1} - \widetilde{v}_{t+1}}_{\text{correction}} \\ & + \underbrace{\widetilde{v}_{t+1} - \mathrm{step}_t(\widetilde{v}_t)}_{\text{update error}} \\ & + \underbrace{\mathrm{step}_t(\widetilde{v}_t) - \mathrm{step}_t(v_t^\star)}_{\text{propagated error}}. \end{aligned} \tag{120}$$

Define the linear propagator

$$A_{s\to t} := (I - \eta D T_{t-1}) \cdots (I - \eta D T_s), \qquad A_{t\to t} := I, \qquad 0 \le s \le t \le \mathcal{T}. \tag{121}$$

By equation 44 and equation 121,

$$\|A_{s\to t}\| \le (1 + C\epsilon)^{t-s} \le e^{CH}. \tag{122}$$

We will prove

$$\left\| \sum_{s=0}^{t-1} \big[A_{s\to t}(\widetilde{v}_s - v_s) - A_{s+1\to t}(\widetilde{v}_{s+1} - v_{s+1}) \big] \right\| \le C_H \epsilon\delta, \tag{123}$$

$$\|\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s)\| \le C\epsilon^2\delta, \tag{124}$$

$$\|\,\mathrm{step}_s(\widetilde{v}_s) - \mathrm{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star)\| \le C\epsilon\|v_s - v_s^\star\| + C\epsilon^2\delta. \tag{125}$$

Assuming equation 123–equation 125, we first derive the conclusion on any initial interval where these bounds hold. We begin with an exact identity: subtract $A_{s\to s+1}(v_s - v_s^\star)$ from equation 120 at time s, then add and subtract $A_{s\to s+1}(\widetilde{v}_s - v_s^\star)$ to obtain

$$\begin{aligned} & (v_{s+1} - v_{s+1}^\star) - A_{s\to s+1}(v_s - v_s^\star) \\ & \quad = A_{s\to s+1}(\widetilde{v}_s - v_s) - (\widetilde{v}_{s+1} - v_{s+1}) \\ & \qquad + \big[\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s)\big] \\ & \qquad + \big[\mathrm{step}_s(\widetilde{v}_s) - \mathrm{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star)\big]. \end{aligned}$$

Multiply this identity by $A_{s+1\to t}$ and sum over $s < t$. By equation 121, $A_{s+1\to t}A_{s\to s+1} = A_{s\to t}$. Since $v_0 = v_0^\star$, the left-hand side telescopes, giving

$$\begin{aligned} v_t - v_t^\star &= \sum_{s=0}^{t-1} \big[A_{s+1\to t}(v_{s+1} - v_{s+1}^\star) - A_{s\to t}(v_s - v_s^\star) \big] \\ &= \sum_{s=0}^{t-1} \big[A_{s\to t}(\widetilde{v}_s - v_s) - A_{s+1\to t}(\widetilde{v}_{s+1} - v_{s+1}) \big] \\ &\quad + \sum_{s=0}^{t-1} A_{s+1\to t} \big[\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s) \big] \\ &\quad + \sum_{s=0}^{t-1} A_{s+1\to t} \big[\mathrm{step}_s(\widetilde{v}_s) - \mathrm{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star) \big]. \end{aligned} \tag{126}$$

Taking norms in equation 126 and applying the triangle inequality to its three sums gives

$$\begin{aligned}
&\|v_t - v_t^\star\| \\
&\quad\le \left\|\sum_{s=0}^{t-1}\big[A_{s\to t}(\widetilde{v}_s - v_s) - A_{s+1\to t}(\widetilde{v}_{s+1} - v_{s+1})\big]\right\| \\
&\qquad + \left\|\sum_{s=0}^{t-1} A_{s+1\to t}\big[\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s)\big]\right\| \\
&\qquad + \left\|\sum_{s=0}^{t-1} A_{s+1\to t}\big[\mathrm{step}_s(\widetilde{v}_s) - \mathrm{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star)\big]\right\|.
\end{aligned}$$

Keep the first sum inside a single norm and apply equation 123. For the second and third sums, apply the triangle inequality and $\|Av\| \le \|A\|\|v\|$, followed by equation 122, equation 124, and equation 125. This yields

$$\begin{aligned}
&\|v_t - v_t^\star\| \\
&\quad\le C_H\epsilon\delta + \sum_{s=0}^{t-1}\|A_{s+1\to t}\|\|\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s)\| \\
&\qquad + \sum_{s=0}^{t-1}\|A_{s+1\to t}\|\|\mathrm{step}_s(\widetilde{v}_s) - \mathrm{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star)\| \\
&\quad\le C_H\epsilon\delta + e^{CH}\sum_{s=0}^{t-1} C\epsilon^2\delta \\
&\qquad + e^{CH}\sum_{s=0}^{t-1}\big(C\epsilon\|v_s - v_s^\star\| + C\epsilon^2\delta\big) \\
&\quad= C_H\epsilon\delta + 2Ce^{CH}t\epsilon^2\delta + Ce^{CH}\epsilon\sum_{s=0}^{t-1}\|v_s - v_s^\star\| \\
&\quad\le \big(C_H + 2CHe^{CH}\big)\epsilon\delta + Ce^{CH}\epsilon\sum_{s=0}^{t-1}\|v_s - v_s^\star\| \\
&\quad\le C_H\epsilon\delta + C_H\epsilon\sum_{s=0}^{t-1}\|v_s - v_s^\star\|. \qquad (127)
\end{aligned}$$

The penultimate inequality uses $t\epsilon^2\delta = (\epsilon t)\epsilon\delta \le H\epsilon\delta$. In the last inequality, we increase C_H once to bound both coefficients uniformly in t and ϵ.

The discrete Gronwall inequality applied to equation 127 gives

$$\begin{aligned}
\|v_t - v_t^\star\| &\le C_H\epsilon\delta(1 + C_H\epsilon)^t \\
&\le C_H\epsilon\delta e^{C_H\epsilon t} \\
&\le C_H e^{C_H H}\epsilon\delta = O(\epsilon\delta). \qquad (128)
\end{aligned}$$

Here we used $1 + C_H\epsilon \le e^{C_H\epsilon}$ and $\epsilon t \le H$.

Bounds for the induction. We now prove $\|v_t - v_t^\star\| \le \delta$ for all $t \le \mathcal{T}$ by induction. The base case holds because $v_0 = v_0^\star$. For the induction step, suppose that

$$\|v_s - v_s^\star\| \le \delta, \qquad 0 \le s \le t < \mathcal{T}. \qquad (129)$$

Lemma D.12 and equation 129 give

$$\begin{aligned}
\|v_s\| &\le \|v_s^\star\| + \|v_s - v_s^\star\| \\
&\le C_\star\delta + \delta = (C_\star + 1)\delta, \qquad 0 \le s \le t. \qquad (130)
\end{aligned}$$

We verify equation 123 for every terminal time $1 \le r \le t+1$, and equation 124–equation 125 for every update index $0 \le s \le t$. All constants in this verification depend only on the fixed input bound in equation 130 and the fixed parameters.

Bounding the sum of the correction terms. For $1 \le r \le t+1$, the correction sum in equation 126 telescopes:

$$\begin{aligned}\sum_{s=0}^{r-1}\big[A_{s\to r}(\widetilde{v}_s - v_s) - A_{s+1\to r}(\widetilde{v}_{s+1} - v_{s+1})\big] \\ = A_{0\to r}(\widetilde{v}_0 - v_0) - (\widetilde{v}_r - v_r).\end{aligned} \tag{131}$$

By equation 130, Lemma D.13 applies at the input 0 and at the input $r-1$. Thus equation 106 bounds the corrections at initialization and at the output r by $O(\epsilon\delta)$. Together with equation 122, this gives

$$\begin{aligned}&\|A_{0\to r}(\widetilde{v}_0 - v_0) - (\widetilde{v}_r - v_r)\| \\ &\quad\le \|A_{0\to r}\|\|\widetilde{v}_0 - v_0\| + \|\widetilde{v}_r - v_r\| \\ &\quad\le e^{CH}C\epsilon\delta + C\epsilon\delta \\ &\quad\le C_H\epsilon\delta.\end{aligned}$$

Combining this with equation 131 proves equation 123 through time $t+1$.

Bounding the update error. For each $0 \le s \le t$, equation 130 allows us to apply Lemma D.13. Its conclusion equation 119 gives

$$\|\widetilde{v}_{s+1} - \mathrm{step}_s(\widetilde{v}_s)\| \le C\epsilon^2\delta.$$

This proves equation 124 for every $0 \le s \le t$.

Bounding the remaining difference of predicted updates. Fix $0 \le s \le t$. By equation 130 and equation 106,

$$\begin{aligned}\|\widetilde{v}_s\| &\le \|v_s\| + \|\widetilde{v}_s - v_s\| \\ &\le (C_\star + 1)\delta + C\epsilon\delta \\ &\le (C_\star + 2)\delta\end{aligned}$$

for sufficiently small ϵ. Also, $\|v_s^\star\| \le C_\star\delta$ by Lemma D.12. For inputs v, v' in the ball of radius $(C_\star+2)\delta$, write their coordinates at time s as (x, y) and (x', y'). Definition D.2 and Lemma D.7(3) give

$$\begin{aligned}|x - x'| \le \|v - v'\|, \qquad |y - y'| \le O(\rho_2)\|v - v'\|, \\ \big||x|^2 - |x'|^2\big| + |x^2 - (x')^2| \le O(\delta)\|v - v'\|, \\ |x\overline{y} - x'\overline{y'}| = |\overline{y}(x - x') + x'(\overline{y} - \overline{y'})| \le O(\rho_2\delta)\|v - v'\|.\end{aligned} \tag{132}$$

The affine forcing in $\mathrm{step}_s(v) - \mathrm{step}_s(v')$ cancels. The linear part differs from $(I - \eta DT_s)(v - v')$ because its output coordinates use P_{s+1} and u_{s+1}. For the complementary component, apply equation 58 to $h = v - v'$. For the critical component, the difference is $a_s(x - x')(u_{s+1} - u_s)$ in the real case and $\sqrt{2}\,\mathrm{Re}(a_s(x - x')(u_{s+1} - u_s))$ in the non-real case. Equations equation 58 and equation 53, together with equation 132, therefore give

$$\begin{aligned}&\|R_sP_s^\perp(v - v') - (I - \eta DT_s)P_s^\perp(v - v')\| \le O(\epsilon^2)\|v - v'\|, \\ &|a_s|\,|x - x'|\,\|u_{s+1} - u_s\| \le O(\epsilon^2)\|v - v'\|.\end{aligned} \tag{133}$$

In Definition D.2, the coefficient of $x\overline{y}$ has norm $O(\eta)$, while those of $|x|^2$ and x^2 have norm $O(\eta\rho_2)$. Here Assumption 6.3(3) gives $\beta_s^\perp/\|\nabla\phi_s^\perp\| = O(\rho_2)$, and Assumption 6.3(2) bounds $\|D^2T_s[u_s, u_s]\|$ by $O(\rho_2)$. Combining equation 133 with equation 132 therefore yields

$$\begin{aligned}&\|\,\mathrm{step}_s(v) - \mathrm{step}_s(v') - (I - \eta DT_s)(v - v')\| \\ &\quad\le O(\epsilon^2)\|v - v'\| + O(\eta)|x\overline{y} - x'\overline{y'}| \\ &\qquad + O(\eta\rho_2)\left(\big||x|^2 - |x'|^2\big| + |x^2 - (x')^2|\right) \\ &\quad\le O(\epsilon^2)\|v - v'\| + O(\eta\rho_2\delta)\|v - v'\| \le C\epsilon\|v - v'\|,\end{aligned} \tag{134}$$

where the last inequality uses Assumption 6.3(3). Apply equation 134 with $v = \widetilde{v}_s$ and $v' = v_s^\star$. Using equation 121 and equation 106, we obtain

$$\begin{aligned}
&\| \operatorname{step}_s(\widetilde{v}_s) - \operatorname{step}_s(v_s^\star) - A_{s\to s+1}(\widetilde{v}_s - v_s^\star)\| \\
&\quad \le C\epsilon \|\widetilde{v}_s - v_s^\star\| \\
&\quad \le C\epsilon \|v_s - v_s^\star\| + C\epsilon \|\widetilde{v}_s - v_s\| \\
&\quad \le C\epsilon \|v_s - v_s^\star\| + C\epsilon^2\delta.
\end{aligned}$$

This proves equation 125 for every $0 \le s \le t$.

Closing the induction. Under equation 129, we have verified equation 123 for every terminal time up to $t+1$, and equation 124–equation 125 for every update index from 0 through t. Consequently, equation 128 applies at time $t+1$:

$$\begin{aligned}
\|v_{t+1} - v_{t+1}^\star\| &\le C_H \epsilon\delta (1 + C_H\epsilon)^{t+1} \\
&\le C_H e^{C_H H}\epsilon\delta \\
&\le \delta,
\end{aligned}$$

where we used $\epsilon(t+1) \le H$ and chose ϵ sufficiently small that $C_H e^{C_H H}\epsilon \le 1$. The constants are independent of the induction step, so the same choice of ϵ applies throughout the trajectory. This closes the induction. The same estimate gives $\|v_t - v_t^\star\| = O(\epsilon\delta)$ for all $t \le \mathcal{T}$. □

Proof of Theorem 6.5. Lemma D.14 gives $v_t = v_t^\star + O(\epsilon\delta)$, proving the deviation estimate in Theorem 6.5. Together with Lemma D.12, this gives $\|v_t\| = O(\delta)$ and places z_t in $\mathcal{S}_t$ for sufficiently small ϵ by Assumption 6.3(3).

For the eigenvalue estimate, consider $\Lambda_t \notin \mathbb{R}$, so $\|u_t\| = \sqrt{2}$. Taylor expansion, Lemma D.7(3),(5), and Assumption 6.3(3), together with equation 94 applied to $v_t^\star$, yield

$$\begin{aligned}
\Lambda(z_t) &= \Lambda_t + (\nabla\Lambda_t)^\top v_t^\star + O(\rho_2\epsilon\delta + \eta\rho_2^2\delta^2) \\
&= \Lambda_t + (\nabla\Lambda_t)^\top \left(\frac{x_t^\star u_t}{\sqrt{2}} + \frac{\overline{x_t^\star}\,\bar{u}_t}{\sqrt{2}} + v_t^{\star,\perp} \right) + O(\epsilon^2/\eta) \\
&= \Lambda_t + \frac{(\nabla\Lambda_t)^\top u_t}{\sqrt{2}} x_t^\star + \frac{(\nabla\Lambda_t)^\top \bar{u}_t}{\sqrt{2}} \overline{x_t^\star} \\
&\quad + (\nabla\Lambda_t)^\top v_t^{\star,\perp} + O(\epsilon^2/\eta).
\end{aligned} \tag{135}$$

By Lemma D.10, $|1 - \eta\Lambda_t| = 1$. The definition of $y_t^\star$ in Definition D.2 therefore gives

$$(\nabla\Lambda_t)^\top v_t^{\star,\perp} = -(1 - \eta\Lambda_t)\overline{y_t^\star}. \tag{136}$$

For the conjugate term in equation 135, equation 45 at the normal reference point, equation 59 from Lemma D.11(4), and symmetry of the two tensor inputs give

$$(\nabla\Lambda_t)^\top \bar{u}_t = \frac{u_t^H D^2 T_t[\bar{u}_t, u_t]}{\|u_t\|^2} = O(\epsilon\rho_2). \tag{137}$$

Since $|x_t^\star| = O(\delta)$ by Lemma D.12, equation 137 and Assumption 6.3(3) bound the conjugate term in equation 135 by $O(\epsilon\rho_2\delta) = O(\epsilon^2/\eta)$. Substituting equation 136 and equation 137 into equation 135, and using $\|u_t\| = \sqrt{2}$, gives

$$\Lambda(z_t) = \Lambda_t - (1 - \eta\Lambda_t)\overline{y_t^\star} + \frac{(\nabla\Lambda_t)^\top u_t}{\|u_t\|} x_t^\star + O(\epsilon^2/\eta), \tag{138}$$

as stated in Theorem 6.5. The real case follows similarly using $v_t^\star = x_t^\star u_t + v_t^{\star,\perp}$ and $\|u_t\| = 1$, with the conjugate term absent. □