

TWO-TIMESCALE EXTRAGRADIENT FOR FINDING LOCAL MINIMAX POINTS

Jiseok Chae, Kyuwon Kim & Donghwan Kim

Department of Mathematical Sciences, KAIST

{jsch, kkw4053, donghwankim}@kaist.ac.kr

ABSTRACT

Minimax problems are notoriously challenging to optimize. However, we present that the two-timescale extragradient method can be a viable solution. By utilizing dynamical systems theory, we show that it converges to points that satisfy the second-order necessary condition of local minimax points, under mild conditions that the two-timescale gradient descent ascent fails to work. This work provably improves upon all previous results on finding local minimax points, by eliminating a crucial assumption that the Hessian with respect to the maximization variable is nondegenerate.

1 INTRODUCTION

Many noteworthy modern machine learning problems, such as generative adversarial networks (GANs) (Goodfellow et al., 2014), adversarial training (Mađry et al., 2018), and sharpness-aware minimization (SAM) (Foret et al., 2021), are instances of minimax problems, formulated as $\min_x \max_y f(x, y)$. First-order methods, such as *gradient descent ascent* (GDA) (Arrow et al., 1958) and *extragradient* (EG) (Korpelevich, 1976), are workhorses of minimax optimization in modern machine learning, but they still remain remarkably unreliable. This is in contrast to the remarkable success of *gradient descent* for minimization problems in machine learning, which is supported by theoretical results; under mild conditions, gradient descent converges to a local minimum, and almost surely avoids strict saddle points (Lee et al., 2016; 2019). Minimax optimization, however, lacks such comparable theory, and this paper is a step towards establishing it.

If the problem is convex-concave, then by Sion’s minimax theorem (Sion, 1958), the order of min and max is insignificant. However, modern machine learning problems are highly nonconvex-nonconcave, and thus their order does matter. Indeed, one of the complications in training GANs, called the *mode collapse* phenomenon, is considered to be due to finding a solution of a max-min problem, rather than that of the min-max problem (Goodfellow, 2016). The widely used notion of saddle points, also called Nash equilibria, fails to capture the ordered structure of minimax problems (Jin et al., 2020). Accordingly, Jin et al. (2020) introduced a new appropriate notion for local optimum in minimax problems, called *local minimax points*, built upon Stackelberg equilibrium from sequential game theory (von Stackelberg, 2011), which encompasses the Nash equilibrium.

Yet, how one can find such local minimax points (possibly via first-order methods) is still left as an open question. A partial answer was provided also by Jin et al. (2020); they showed that the *two-timescale* GDA (Heusel et al., 2017), *i.e.*, GDA with different step sizes for x and y , converges to certain (but not any) local minimax points. In particular, they only covered the case where the Hessian with respect to the maximization variable $\nabla_{yy}^2 f$ is nondegenerate, which disregards possibly meaningful “degenerate” local optimal points, *e.g.*, in over-parameterized training (Liu et al., 2022). In this paper, we show that the *two-timescale* EG can actually find local minimax points beyond the assumption that $\nabla_{yy}^2 f$ is nondegenerate. So, our main contribution is providing a more complete answer to the aforementioned open question. Our specific contributions can be listed as follows.

- In Section 3, we derive a second order characterization of local minimax points, without assuming that $\nabla_{yy}^2 f$ is nondegenerate. In doing so, we introduce the notion of a *restricted* Schur complement. This leads to a natural way of defining *strict non-minimax points* that we would like to avoid, analogous to the strict saddle points in minimization.

- In Section 4, we develop new tools to achieve the main results in Section 5. These tools include a new spectral analysis that does not rely on the nondegeneracy assumption on $\nabla_{\mathbf{y}\mathbf{y}}^2 f$, as well as the concept of hemicurvature to better understand the behavior of eigenvalues.
- In Section 5, from a dynamical system perspective, we show that the limit points of the two-timescale EG in the continuous time limit are the local minimax points under mild conditions, by establishing a second order property of those points. This continuous-time analysis is then used to derive a similar conclusion in the discrete-time case. In the discrete-time case, Section 5.4 further shows that two-timescale EG almost surely avoids (undesirable) *strict non-minimax points*, as desired, while Section 4.3 shows that two-timescale GDA may avoid (desirable) local minimax points where $\nabla_{\mathbf{y}\mathbf{y}}^2 f$ is *degenerate*.
- In Section 6, we extend the local result of Section 5 to a global statement: under the Minty variational inequality (MVI) condition (Minty, 1967) and additional mild conditions, the two-timescale EG *globally* converges to local minimax points.

2 PRELIMINARIES

2.1 NOTATIONS AND PROBLEM SETTING

The *spectrum* of a square matrix \mathbf{A} , denoted by $\text{spec}(\mathbf{A})$, is the set of all eigenvalues of \mathbf{A} . The *range* of a matrix \mathbf{A} , denoted by $\mathcal{R}(\mathbf{A})$, is the span of its column vectors. The *open left* (resp. *right*) *half plane*, denoted by \mathbb{C}^- (resp. \mathbb{C}_+^o), is the set of all complex numbers z such that $\text{Re } z < 0$ (resp. $\text{Re } z > 0$). The *imaginary axis* is denoted by $i\mathbb{R}$. To denote the minimization variable $\mathbf{x} \in \mathbb{R}^{d_1}$ and the maximization variable $\mathbf{y} \in \mathbb{R}^{d_2}$ at once, we use the notation $\mathbf{z} := (\mathbf{x}, \mathbf{y})$. Let C^2 be the set of twice continuously differentiable functions. The saddle-gradient operator of the objective function f will be denoted by $\mathbf{F} := (\nabla_{\mathbf{x}} f, -\nabla_{\mathbf{y}} f)$, and the derivative of \mathbf{F} will be denoted by $D\mathbf{F}$. When necessary, we will impose the following standard assumption on f .

Assumption 1. *Let $f \in C^2$, and there exists $L > 0$ such that $\|D\mathbf{F}(\mathbf{z})\| \leq L$ for all \mathbf{z} .*

2.2 LOCAL MINIMAX POINTS

Jin et al. (2020) introduced the following new notion of local optimality for minimax problems.

Definition 1 (Jin et al. (2020)). *A point $(\mathbf{x}^*, \mathbf{y}^*)$ is said to be a **local minimax point** if there exists $\delta_0 > 0$ and a function h satisfying $h(\delta) \rightarrow 0$ as $\delta \rightarrow 0$ such that, for any $\delta \in (0, \delta_0]$ and any (\mathbf{x}, \mathbf{y}) satisfying $\|\mathbf{x} - \mathbf{x}^*\| \leq \delta$ and $\|\mathbf{y} - \mathbf{y}^*\| \leq \delta$, we have*

$$f(\mathbf{x}^*, \mathbf{y}) \leq f(\mathbf{x}^*, \mathbf{y}^*) \leq \max_{\mathbf{y}': \|\mathbf{y}' - \mathbf{y}^*\| \leq h(\delta)} f(\mathbf{x}, \mathbf{y}'). \quad (1)$$

Local minimax points can be characterized by the following conditions (Jin et al., 2020).

- (First-order necessary) For $f \in C^1$, any local minimax point \mathbf{z}^* satisfies $\nabla f(\mathbf{x}^*, \mathbf{y}^*) = \mathbf{0}$.
- (Second-order necessary) For $f \in C^2$, any local minimax \mathbf{z}^* satisfies $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*) \preceq \mathbf{0}$. In addition, if $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*) \prec \mathbf{0}$, then $[\nabla_{\mathbf{x}\mathbf{x}}^2 f - \nabla_{\mathbf{x}\mathbf{y}}^2 f (\nabla_{\mathbf{y}\mathbf{y}}^2 f)^{-1} \nabla_{\mathbf{y}\mathbf{x}}^2 f](\mathbf{x}^*, \mathbf{y}^*) \succeq \mathbf{0}$.
- (Second-order sufficient) For $f \in C^2$, any stationary point \mathbf{z}^* satisfying

$$[\nabla_{\mathbf{x}\mathbf{x}}^2 f - \nabla_{\mathbf{x}\mathbf{y}}^2 f (\nabla_{\mathbf{y}\mathbf{y}}^2 f)^{-1} \nabla_{\mathbf{y}\mathbf{x}}^2 f](\mathbf{x}^*, \mathbf{y}^*) \succ \mathbf{0} \quad \text{and} \quad \nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*) \prec \mathbf{0} \quad (2)$$

is a local minimax point.

Here, the second-order necessary condition is loose when $\nabla_{\mathbf{y}\mathbf{y}}^2 f$ is degenerate. However, the existing works mostly overlook this discrepancy between that and the second-order sufficient condition (2), and focus only on the *strict*¹ local minimax points (Jin et al., 2020), also known as differential Stackelberg equilibria (Fiez et al., 2020), which are the stationary points that satisfy (2). We close this gap in Section 3, and consider a broader set of local minimax points in later sections.

¹The term *strict* used here is slightly more restrictive than that in the usual strict notion of local optimum in minimization; see e.g., (Wright & Recht, 2022, p.15).

2.3 GRADIENT DESCENT ASCENT AND EXTRAGRADIENT

This paper considers the following two standard gradient methods (with time separation introduced later). Here, η denotes the step size.

- Gradient descent ascent (GDA) (Arrow et al., 1958): $\mathbf{z}_{k+1} = \mathbf{z}_k - \eta \mathbf{F}(\mathbf{z}_k)$
- Extragradient (EG) (Korpelevich, 1976): $\mathbf{z}_{k+1} = \mathbf{z}_k - \eta \mathbf{F}(\mathbf{z}_k - \eta \mathbf{F}(\mathbf{z}_k))$

We also consider their continuous time limits, as the analyses in continuous time is more accessible than their discrete counterparts. The analyses in discrete time, which is more of our interest, can be easily translated from the continuous time limit analyses. Taking the limit as $\eta \rightarrow 0$, both GDA and EG share the same continuous time limit $\dot{\mathbf{z}}(t) = -\mathbf{F}(\mathbf{z}(t))$. However, it is also known that even when f is convex-concave, EG converges to an optimum, while GDA and $\dot{\mathbf{z}}(t) = -\mathbf{F}(\mathbf{z}(t))$ may not (Mescheder et al., 2018). This distinction suggested a need for ODEs with a higher order of approximation. In particular, Lu (2022) derived the ODE $\dot{\mathbf{z}}(t) = -\mathbf{F}(\mathbf{z}(t)) + (\eta/2)D\mathbf{F}(\mathbf{z}(t))\mathbf{F}(\mathbf{z}(t))$ as the $O(\eta)$ -approximation of EG, in the sense that it captures the dynamics of EG up to order $O(\eta)$ as $\eta \rightarrow 0$. In our dynamical system analyses, we found the following to be particularly more useful.

Lemma 2.1. *Under Assumption 1, let $s := \eta/2$ and $0 < s < 1/L$. Then, the ordinary differential equation $\dot{\mathbf{z}}(t) = -(\mathbf{I} + sD\mathbf{F}(\mathbf{z}(t)))^{-1}\mathbf{F}(\mathbf{z}(t))$ is a $O(s)$ -approximation of EG.*

2.4 DYNAMICAL SYSTEMS

This paper analyzes (two-timescale) GDA and EG methods through dynamical systems, in both continuous-time and discrete-time. In particular, we determine whether an equilibrium point \mathbf{z}^* is *asymptotically stable*, which means that a trajectory starting at a point that is sufficiently close to \mathbf{z}^* will remain close enough and eventually converge to \mathbf{z}^* . We adopt the following widely used criteria, which we refer to as *strict linear stability*, as a sufficient condition for an equilibrium to be asymptotically stable; see, e.g., (Khalil, 2002, Theorem 4.7) and (Galor, 2007, Theorem 4.8).

Definition 2 (Linear stability of dynamical systems).

- (Continuous system) For a C^1 mapping ϕ , an equilibrium point \mathbf{z}^* of $\dot{\mathbf{z}}(t) = \phi(\mathbf{z}(t))$, such that $\phi(\mathbf{z}^*) = \mathbf{0}$, is a **strict linearly stable point** if its Jacobian matrix at \mathbf{z}^* has all eigenvalues with negative real parts, i.e., $\text{spec}(D\phi(\mathbf{z}^*)) \subset \mathbb{C}_-^\circ$.
- (Discrete system) For a C^1 mapping \mathbf{w} , an equilibrium point \mathbf{z}^* of $\mathbf{z}_{k+1} = \mathbf{w}(\mathbf{z}_k)$, such that $\mathbf{z}^* = \mathbf{w}(\mathbf{z}^*)$, is a **strict linearly stable point** if its Jacobian matrix at \mathbf{z}^* has spectral radius smaller than 1, i.e., $\rho(D\mathbf{w}(\mathbf{z}^*)) < 1$.

The set of *unstable* equilibrium points in a discrete dynamical system is similarly defined.

Definition 3. *Given a C^1 mapping \mathbf{w} , the set $\mathcal{A}^*(\mathbf{w}) := \{\mathbf{z}^* : \mathbf{z}^* = \mathbf{w}(\mathbf{z}^*), \rho(D\mathbf{w}(\mathbf{z}^*)) > 1\}$ is the set of **strict linearly unstable equilibrium points**.*

We also study whether the methods avoid strictly non-optimal points in minimax problems, using the following theorem, built upon the stable manifold theorem (Shub, 1987). This was used by Lee et al. (2019) to show that gradient descent in minimization almost surely escapes strict saddle points.

Theorem 2.2 (Lee et al. (2019, Theorem 2)). *Let \mathbf{w} be a C^1 mapping such that $\det(D\mathbf{w}(\mathbf{z})) \neq \mathbf{0}$ for all \mathbf{z} . Then the set of initial points that converge to a strict linearly unstable equilibrium point has (Lebesgue) measure zero, i.e., $\mu(\{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(\mathbf{z}_0) \in \mathcal{A}^*(\mathbf{w})\}) = 0$.*

3 ON THE NECESSARY CONDITION OF LOCAL MINIMAX POINTS

In minimization, both second-order necessary and sufficient conditions of local minimum are simply characterized by the Hessian of the objective function; see e.g., (Wright & Recht, 2022, Theorems 2.4 and 2.5). However, not a similar simple correspondence between conditions for a local minimax point was previously known, making it difficult to further develop a theory for minimax optimization comparable to that in minimization. In this section, we show that such correspondence in fact exists.

3.1 RESTRICTED SCHUR COMPLEMENT

For the clarity of presentation, we first define a variant of Schur complement, named *restricted Schur complement*, which we will soon relate to second-order conditions of a local minimax point.

From now on, we assume that f is twice continuously differentiable. Let us denote the second derivatives of f by $\mathbf{A} = \nabla_{\mathbf{x}\mathbf{x}}^2 f \in \mathbb{R}^{d_1 \times d_1}$, $\mathbf{B} = \nabla_{\mathbf{y}\mathbf{y}}^2 f \in \mathbb{R}^{d_2 \times d_2}$, and $\mathbf{C} = \nabla_{\mathbf{x}\mathbf{y}}^2 f \in \mathbb{R}^{d_1 \times d_2}$. In particular we can express the Jacobian matrix of the saddle-gradient \mathbf{F} as

$$\mathbf{H} := D\mathbf{F} = \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ -\mathbf{C}^\top & -\mathbf{B} \end{bmatrix} \in \mathbb{R}^{(d_1+d_2) \times (d_1+d_2)}. \quad (3)$$

Although \mathbf{A} , \mathbf{B} , and \mathbf{C} are, strictly speaking, matrix valued *functions* of $(d_1 + d_2)$ -dimensional vector inputs, the point where those functions are evaluated will be clear from context, so we simply write \mathbf{A} , \mathbf{B} , and \mathbf{C} to denote the function values.

Let $r := \text{rank}(\mathbf{B})$. As \mathbf{B} is symmetric, it can be orthogonally diagonalized as $\mathbf{B} = \mathbf{P}\mathbf{\Delta}\mathbf{P}^\top$, where $\mathbf{\Delta} = \text{diag}\{\delta_1, \dots, \delta_r, 0, \dots, 0\}$ and \mathbf{P} is orthogonal. Let $\mathbf{\Gamma}$ be a submatrix of $\mathbf{C}\mathbf{P}$, which consists of the $d_2 - r$ rightmost columns of $\mathbf{C}\mathbf{P}$, let $q := \text{rank}(\mathbf{\Gamma})$, and let $\mathbf{U} \in \mathbb{R}^{d_1 \times (d_1 - q)}$ be a matrix whose columns form an orthonormal basis of $\mathcal{R}(\mathbf{\Gamma})^\perp$. Under this setting, we define the following.

Definition 4. For a matrix \mathbf{H} in the block form of (3), using the definition² of \mathbf{U} given above, the *restricted Schur complement*³ is defined as $\mathbf{S}_{\text{res}}(\mathbf{H}) := \mathbf{U}^\top(\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top)\mathbf{U} \in \mathbb{R}^{(d_1 - q) \times (d_1 - q)}$.

The positive semidefiniteness of the restricted Schur complement \mathbf{S}_{res} is related to the (generalized) Schur complement $\mathbf{S} := \mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top$ being positive semidefinite on a certain subspace, as follows. We remark that a 0×0 matrix is considered to be ‘‘vacuously’’ positive semidefinite.

Proposition 3.1. *The restricted Schur complement $\mathbf{S}_{\text{res}}(\mathbf{H})$ is positive semidefinite if and only if $\mathbf{v}^\top \mathbf{S} \mathbf{v} \geq 0$ for any $\mathbf{v} \in \mathbb{R}^{d_1}$ satisfying $\mathbf{C}^\top \mathbf{v} \in \mathcal{R}(\mathbf{B})$.*

3.2 REFINING THE SECOND-ORDER NECESSARY CONDITION

We then have the following second-order necessary condition in terms of the restricted Schur complement, which is our first main contribution. Notice that when \mathbf{B} is invertible, the restricted Schur complement becomes the usual Schur complement $\mathbf{A} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^\top$, which appears in the second-order sufficient condition (2). Thus, the *restricted Schur complement* provides a unified point of view in considering both the necessary and sufficient conditions.

Proposition 3.2 (Refined second-order necessary condition). *Let $f \in C^2$, then any local minimax point $(\mathbf{x}^*, \mathbf{y}^*)$ satisfies $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*) \preceq \mathbf{0}$. In addition, if the function $h(\delta)$ in Definition 1 satisfies $\limsup_{\delta \rightarrow 0^+} h(\delta)/\delta < \infty$, then $\mathbf{S}_{\text{res}}(D\mathbf{F}(\mathbf{x}^*, \mathbf{y}^*)) \succeq \mathbf{0}$.*

Remark 3.3. *The additional condition on h in Proposition 3.2 slightly limits the choice of h , which determines the size of the set we take the maximum over in Definition 1. Hence, this can be interpreted as restricting the local minimax points we are interested in, or alternatively, refining the definition of local minimax points itself; c.f., (Ma et al., 2023). Note, however, that such refined definition is yet much broader than the definition of strict local minimax points.*

Based on the refined second-order necessary condition, we define *strict non-minimax* points as below. These are the points we hope to avoid, and the definition is analogous to that of strict saddle points in minimization problems (Lee et al., 2016). In Section 5.4, we show that two-timescale EG almost surely avoids strict non-minimax points.

Definition 5 (Strict non-minimax point; \mathcal{T}^*). *A stationary point \mathbf{z}^* is said to be a **strict non-minimax point** of f if $\lambda_{\min}(\mathbf{S}_{\text{res}}(D\mathbf{F}(\mathbf{z}^*))) < 0$ or $\lambda_{\min}(-\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)) < 0$, where $\lambda_{\min}(\mathbf{A})$ denotes the smallest eigenvalue of \mathbf{A} . We denote the set of strict non-minimax points by \mathcal{T}^* .*

When necessary, we will impose the following assumption at stationary points \mathbf{z}^* that is weaker than the nondegeneracy condition on $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$, which was crucial in all existing literatures.

²A matrix \mathbf{U} is not unique in general, but there are ways to fix the choice of \mathbf{U} , e.g., applying a predetermined QR factorization algorithm on $\mathbf{\Gamma}$. However, because only the spectrum of $\mathbf{S}_{\text{res}}(\mathbf{H})$ will be important in the subsequent analyses, we do not specify the choice of \mathbf{U} .

³A similar matrix has also been considered by Zhang et al. (2022, Theorem 4.4), but as a part of a local optimality condition for *quadratic* problems only.

Assumption 2. For a stationary point \mathbf{z}^* in consideration, at least one of the matrices $\mathbf{S}_{\text{res}}(D\mathbf{F}(\mathbf{z}^*))$ and $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$ is nondegenerate.

Example 1. We provide two simple representative examples of “non-strict” local minimax points satisfying Assumption 2, especially with nondegenerate $\mathbf{S}_{\text{res}}(D\mathbf{F})$ and degenerate $\nabla_{\mathbf{y}\mathbf{y}}^2 f$.

- $f_{\text{b}}(x, y) = xy$ has a unique local minimax point $\mathbf{0} = (0, 0)$;
- $f_{\text{q}}(x_1, x_2, y_1, y_2, y_3) = \frac{1}{2}x_1^2 - \frac{1}{2}y_1^2 - \frac{1}{2}y_3^2 + x_2y_2 + y_1y_3$ has non-unique and non-isolated local minimax points $(0, 0, t, 0, t)$ for all $t \in \mathbb{R}$.

See Appendix D.3 for the details and proofs. Later in Examples 2, 3 and 4, we show that our proposed two-timescale EG finds these optima, while GDA (and its timescaled variant) avoids them.

Only when characterizing *strict* linear stability of equilibrium points in Sections 5.2 and 5.3, we will use the following slightly stronger version of Assumption 2.

Assumption 2’. For \mathbf{z}^* in consideration, in addition to Assumption 2, $D\mathbf{F}(\mathbf{z}^*)$ is nondegenerate.

4 CHARACTERIZING TIMESCALE SEPARATION WITHOUT NONDEGENERACY CONDITION ON $\nabla_{\mathbf{y}\mathbf{y}}^2 f$

In this section we begin with reviewing the close relationship between two-timescale GDA and strict local minimax points, under a nondegeneracy condition on $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$ (Jin et al., 2020). Then, we extend this to a more general setting without such nondegeneracy condition, and observe a negative result stating that two-timescale GDA avoids some degenerate local minimax points.

4.1 TIMESCALE SEPARATION IN GDA AND ITS RELATION TO STABILITY

Let us recall the *two-timescale GDA* method, which is GDA with a timescale separation with timescale parameter $\tau \geq 1$:

$$\mathbf{z}_{k+1} = \tilde{\mathbf{w}}_{\tau}(\mathbf{z}_k) := \mathbf{z}_k - \eta \mathbf{\Lambda}_{\tau} \mathbf{F}(\mathbf{z}_k) \quad \text{where} \quad \mathbf{\Lambda}_{\tau} := \text{diag}\{(1/\tau)\mathbf{I}, \mathbf{I}\}, \quad (4)$$

and its continuous time limit $\dot{\mathbf{z}}(t) = -\mathbf{\Lambda}_{\tau} \mathbf{F}(\mathbf{z}(t))$. The stability of (4) and its continuous limit depends on the spectrum of a timescaled matrix

$$\mathbf{H}_{\tau} := \mathbf{\Lambda}_{\tau} \mathbf{H} = \begin{bmatrix} \frac{1}{\tau} \mathbf{A} & \frac{1}{\tau} \mathbf{C} \\ -\mathbf{C}^{\top} & -\mathbf{B} \end{bmatrix}.$$

Jin et al. (2020) studied the following asymptotic behavior of its spectrum as $\epsilon = 1/\tau \rightarrow 0$, under a nondegeneracy condition on $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$.

Lemma 4.1 (Jin et al. (2020, Lemma 40)). *Suppose that $\mathbf{B} = \nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$ is nondegenerate. Then, the $d_1 + d_2$ complex eigenvalues λ_j of \mathbf{H}_{τ} have the following asymptotics as $\epsilon = 1/\tau \rightarrow 0+$:*

$$|\lambda_j - \epsilon \mu_j| = o(\epsilon), \quad j = 1, \dots, d_1, \quad |\lambda_{j+d_1} - \nu_j| = o(1), \quad j = 1, \dots, d_2,$$

where μ_j and ν_j are the eigenvalues of $\mathbf{A} - \mathbf{C}\mathbf{B}^{-1}\mathbf{C}^{\top}$ and $-\mathbf{B}$, respectively.

The eigenvalues of \mathbf{H}_{τ} become related to the definition of the strict local minimax point (2) as $\epsilon \rightarrow 0$, and this was used to show that GDA with sufficiently large timescale separation converges to a strict local minimax point in Fiez & Ratliff (2021, Theorem 1), and also in Jin et al. (2020, Theorem 28).

Theorem 4.2 (Fiez & Ratliff (2021, Theorem 1)). *For $f \in C^2$ and its stationary point \mathbf{z}^* with $\det(\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)) \neq 0$, the point \mathbf{z}^* is a strict local minimax point if and only if there exists a constant $\tau^* > 0$ such that the point \mathbf{z}^* is strict linearly stable for two-timescale GDA with any $\tau > \tau^*$.*

This seems to be complete, but this tells us nothing about when the nondegeneracy condition on $\nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{z}^*)$ is removed. So, as our next step, we generalize Lemma 4.1, the spectral analysis of the timescaled matrix \mathbf{H}_{τ} , and investigate the stability of the two-timescale GDA without such nondegeneracy condition.

4.2 TIMESCALE SEPARATION WITHOUT THE NONDEGENERACY CONDITION ON $\nabla_{\mathbf{y}\mathbf{y}}^2 f$

Mimicking how we constructed the restricted Schur complement (in Definition 4 and above), we can reduce \mathbf{H}_τ into a simpler form without changing the spectrum. More precisely, for a block diagonal matrix $\mathbf{Q} = \text{diag}\{\mathbf{I}, \mathbf{P}^\top\}$, the matrix $\mathbf{Q}\mathbf{H}_\tau\mathbf{Q}^\top = \begin{bmatrix} \frac{1}{\tau}\mathbf{A} & \frac{1}{\tau}\mathbf{C}\mathbf{P} \\ -(\mathbf{C}\mathbf{P})^\top & -\mathbf{\Delta} \end{bmatrix}$ is similar to \mathbf{H}_τ , hence has the same spectrum with \mathbf{H}_τ . Thus, by replacing \mathbf{C} with $\mathbf{C}\mathbf{P}$ if necessary, we may assume without loss of generality that \mathbf{B} is a diagonal matrix. Moreover, with $r = \text{rank}(\mathbf{B})$, we may further assume that there exists a diagonal $\mathbf{D} \in \mathbb{R}^{r \times r}$ with nonzero diagonal entries where \mathbf{B} is of the form $\mathbf{B} = \begin{bmatrix} -\mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$. Then a subdivision $\mathbf{C} = [\mathbf{C}_1 \ \mathbf{C}_2]$ arises naturally, where \mathbf{C}_1 is constructed from \mathbf{C} by taking the r leftmost columns, and \mathbf{C}_2 takes the rest. Note that $q = \text{rank}(\mathbf{\Gamma}) = \text{rank}(\mathbf{C}_2)$.

Using these notations, we can characterize the asymptotic behaviors of the eigenvalues of \mathbf{H}_τ when $\tau \rightarrow \infty$, as in the following theorem. This reduces to Lemma 4.1 when $r = d_2$ (and thus $q = 0$).

Theorem 4.3. *Let $\mathbf{A} \in \mathbb{R}^{d_1 \times d_1}$ be a square matrix, and let $\mathbf{B} \in \mathbb{R}^{d_2 \times d_2}$ and $\mathbf{C} \in \mathbb{R}^{d_1 \times d_2}$ be matrices in the form that is discussed above. Assume that $f \in C^2$. Then, under Assumption 2, for $\tau \geq 1$ and $\epsilon := 1/\tau$, it is possible to construct continuous functions $\lambda_j(\epsilon)$, $j = 1, \dots, d_1 + d_2$ so that they are the $d_1 + d_2$ complex eigenvalues of \mathbf{H}_τ , with the following asymptotics as $\epsilon \rightarrow 0+$:*

- (i) $|\lambda_j - i\sigma_j\sqrt{\epsilon}| = o(\sqrt{\epsilon}), \quad |\lambda_{j+d_1} + i\sigma_j\sqrt{\epsilon}| = o(\sqrt{\epsilon}), \quad j = 1, \dots, q,$
- (ii) $|\lambda_{j+q} - \epsilon\mu_j| = o(\epsilon), \quad j = 1, \dots, d_1 - q,$
- (iii) $|\lambda_{j+d_1+q} - \nu_j| = o(1), \quad j = 1, \dots, r,$

with being nonzero whenever $\epsilon > 0$ while the remaining $d_2 - q - r$ are constantly 0. Here, μ_j are the eigenvalues of the restricted Schur complement $\mathbf{S}_{\text{res}}(\mathbf{H})$, ν_j are the nonzero eigenvalues of $-\mathbf{B}$, and σ_j are the singular values of \mathbf{C}_2 .

As we remove the nondegeneracy condition on \mathbf{B} , the $d_1 - q$ eigenvalues of type (ii) are now related to the eigenvalues of the restricted Schur complement \mathbf{S}_{res} , rather than the Schur complement as in Lemma 4.1. This illustrates a close connection to the refined second-order necessary condition. Nevertheless, the overall form of the asymptotic behavior of the eigenvalues of types (ii) and (iii) remains unchanged from that of Lemma 4.1. What differentiates Theorem 4.3 from Lemma 4.1 the most is the existence of the unprecedented $2q$ eigenvalues of type (i).

Example 2. *Consider the bilinear function $f_b(x, y) = xy$ in Example 1. A straightforward computation gives $\text{spec}(\mathbf{H}_\tau) = \{\pm i\sqrt{\epsilon}\}$. Notice that this spectrum cannot be explained by Lemma 4.1, and Theorem 4.2 does not apply. In fact, $\text{spec}(-\mathbf{H}_\tau) \not\subset \mathbb{C}_-^<$ for any $\tau > 0$, so from Definition 2(i) one can see that a unique (non-strict) local minimax point $\mathbf{0}$ is not strict linearly stable for continuous-time GDA, even with timescale separation. Moreover, Definition 3 implies that the point $\mathbf{0}$ is in fact rather a strict linearly unstable point for discrete-time two-timescale GDA.*

On the contrary, Theorem 4.3 exactly explains why $\text{spec}(\mathbf{H}_\tau)$ is $\{\pm i\sqrt{\epsilon}\}$; this is an example where $\text{spec}(\mathbf{H}_\tau)$ contains only the eigenvalues of type (i). As we will see in Proposition 5.1 and 5.5, for an equilibrium point to be strict linearly stable for EG with timescale separation, the spectrum $\text{spec}(\mathbf{H}_\tau)$ must lie on a certain target set, which contains a deleted neighborhood of the origin in the imaginary axis. Thus, $\mathbf{0}$ becomes strict linearly stable for two-timescale EG, as desired.

However, the general situation in terms of the type (i) eigenvalues is of course not always as simple as those in Example 2. This necessitated us to introduce the term *hemicurvature* of the eigenvalue function $\lambda_j(\epsilon)$, as in next subsection, to help identify whether $\text{spec}(\mathbf{H}_\tau)$ lie in a certain *target* set.

4.2.1 HEMICURVATURE OF THE EIGENVALUE FUNCTION $\lambda_j(\epsilon)$

Let us focus on the eigenvalues of type (i), i.e., λ_j with $j \in \mathcal{I} := \{1, \dots, q\} \cup \{d_1 + 1, \dots, d_1 + q\}$. Identifying the complex plane with \mathbb{R}^2 , we may consider $\lambda_j(\epsilon)$ as a plane curve $(\text{Re } \lambda_j(\epsilon), \text{Im } \lambda_j(\epsilon))$ parametrized by ϵ . All that Theorem 4.3 tells us about $\lambda_j(\epsilon)$ is that its leading term is of order $\sqrt{\epsilon}$, which determines the tangential direction of λ_j . More precisely, $\lambda_j(\epsilon)$ converges to 0 as $\epsilon \rightarrow 0+$ asymptotically in the direction along the imaginary axis. The complication here is that the target sets, as mentioned in Example 2, for both GDA and EG happen to have the imaginary axis also as their tangential line at the origin. That is, the tangential information (Theorem 4.3) alone is not sufficient for our stability analysis. For further details, with figures, on this issue, see Appendix E.3.

The local canonical form of curves suggests investigating the curvature information of λ_j . To this end, let us identify the coordinate functions $\operatorname{Re}(\lambda_j)$ and $\operatorname{Im}(\lambda_j)$, when $j \in \mathcal{I}$. By reindexing the singular values of \mathbf{C}_2 for notational convenience, we already know that $\operatorname{Im}(\lambda_j) = \pm i\sigma_j\sqrt{\epsilon} + o(\sqrt{\epsilon})$. To characterize $\operatorname{Re}(\lambda_j)$, by identifying $\mathbf{H}_\tau = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ -\mathbf{C}^\top & -\mathbf{B} \end{bmatrix} + \epsilon \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$ as a perturbed linear operator, one can use a classical result from perturbation theory (see, e.g., Section II.1.2 of (Kato, 1995)) to expand $\lambda_j(\epsilon)$ into a convergent Puiseux series $\lambda_j(\epsilon) = \sum_{k=1}^{\infty} \alpha_j^{(k)} \epsilon^{k/p_j}$ for some positive integer p_j . Let k_0 be the smallest k such that $\operatorname{Re} \alpha_j^{(k)} \neq 0$, and define $\varrho_j := k_0/p_j$, $\zeta_j := \alpha_j^{(k_0)}$, and $\xi_j := \operatorname{Re} \zeta_j$. Then, we may write $\operatorname{Re}(\lambda_j) = \xi_j \epsilon^{\varrho_j} + o(\epsilon^{\varrho_j})$. Notice that $\varrho_j > 1/2$. If such k_0 does not exist, then $\operatorname{Re}(\lambda_j) = 0$ holds, so we may simply put $\varrho_j = 1$ and $\xi_j = 0$. Now, instead of the curvature itself, we define and consider the following value, named a *hemicurvature*:

$$\iota_j := \lim_{\epsilon \rightarrow 0^+} (\xi_j \epsilon^{\varrho_j - 1}) / \sigma_j^2. \quad (5)$$

In Proposition E.4 we show that the hemicurvature ι_j is equal to $-1/2$ times the curvature of the trajectory of λ_j when $\epsilon \rightarrow 0^+$. The reason we consider this quantity instead of the curvature is because not only it captures the curvature information, but also it is the limit of $\operatorname{Re}(1/\lambda_j)$ as $\epsilon \rightarrow 0$; see Proposition E.5. This simplifies the subsequent analyses, as we see in Section 5.

4.3 TWO-TIMESCALE GDA AVOIDS SOME NON-STRICT LOCAL MINIMAX POINTS

It is obviously desirable to have a method that converges to an optimal point, regardless of whether the point is strict or non-strict. This unfortunately fails for two-timescale GDA, which almost surely avoids some non-strict local minimax points. We already saw such an instance in Example 2. More generally, consider the set $\mathcal{X}_{\text{ns}}^*$ of non-strict local minimax points satisfying $\iota_j < \eta/2$ for some $j \in \mathcal{I}$. We have a negative result below that timescale separation does not help GDA to converge to such non-strict points, despite being useful in converging to *strict* local minimax points.

Theorem 4.4. *Let $\mathbf{z}^* \in \mathcal{X}_{\text{ns}}^*$. Under Assumptions 1, 2 and $0 < \eta < 1/L$, there exists a constant $\tau^* > 0$ such that for any $\tau > \tau^*$, the set of initial points converging to \mathbf{z}^* by the discrete-time two-timescale GDA $\tilde{\mathbf{w}}_\tau$ has Lebesgue measure zero, i.e., $\mu(\{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \tilde{\mathbf{w}}_\tau^k(\mathbf{z}_0) = \mathbf{z}^*\}) = 0$.*

5 LOCAL MINIMAX POINTS AND THE LIMIT POINTS OF TWO-TIMESCALE EG

In this section we show that the two-timescale EG converges to a stationary point that satisfies the refined second-order necessary condition mentioned in Proposition 3.1. We first show that such points with an invertible $D\mathbf{F}$ are the strict linearly stable points of two-timescale EG. Then, we show that two-timescale EG almost surely avoids a strict non-minimax point, even without the invertibility assumption on $D\mathbf{F}$. These results are comparable to the convergence guarantees known for gradient descent on minimization problems—namely that gradient descent locally converges to local minima, and almost surely avoids strict *saddle* points (Lee et al., 2016).

Before we begin, we would like to remark that a point not being strict linearly *stable*—for example because of $D\mathbf{F}$ being not invertible—does not imply that the point is strict linearly *unstable*. In other words, the strict linear stability is sufficient but not necessary for a method to find that point, as will soon be detailed in Section 6.

5.1 TWO-TIMESCALE EG

We now consider the EG with timescale separation, to resolve the aforementioned avoidance issue of GDA. For simplicity in the analysis, we first work with the continuous-time limit of EG, especially the high-order approximation version in Lemma 2.1, so that we have a clear distinction from GDA. Applying the timescale separation on the ODE in Lemma 2.1 using $\mathbf{\Lambda}_\tau$ as in (4), we get

$$\dot{\mathbf{z}}(t) = \phi_\tau(\mathbf{z}(t)) := -(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(\mathbf{z}(t)))^{-1} \mathbf{\Lambda}_\tau \mathbf{F}(\mathbf{z}(t)). \quad (6)$$

Its discretization with $s = \eta/2$, namely the two-timescale EG, becomes

$$\mathbf{z}_{k+1} = \mathbf{w}_\tau(\mathbf{z}_k) := \mathbf{z}_k - \eta \mathbf{\Lambda}_\tau \mathbf{F}(\mathbf{z}_k - \eta \mathbf{\Lambda}_\tau \mathbf{F}(\mathbf{z}_k)). \quad (7)$$

Their correspondence can be easily shown by following the proof of Lemma 2.1 with $\|\mathbf{\Lambda}_\tau\| \leq 1$. From now on, we will refer both of them simply as τ -EG, and denote their Jacobian matrix, either

$D\phi_\tau$ or Dw_τ , as \mathbf{J}_τ . Whether they are in terms of continuous-time or discrete-time will be clear from the context. τ -EG takes the stationary points of \mathbf{F} as its equilibria, see Appendix F.

The analyses of the stability properties are based on Definition 2, so the Jacobian matrices, of the systems (6) and (7), play a central role. Hence, in this section, for convenience, the Jacobian matrix at a given equilibrium point \mathbf{z}^* will be denoted by $\mathbf{J}_\tau^* := \mathbf{J}_\tau(\mathbf{z}^*)$. Meanwhile, similar to Theorem 4.2, the upcoming stability results share a property of being held for all τ larger than a certain threshold. So for simplicity, we introduce the following notion of ∞ -EG.

Definition 6. We say that \mathbf{z}^* is a strict linearly stable equilibrium point of ∞ -EG if, there exists some τ^* such that, for any $\tau > \tau^*$ the point \mathbf{z}^* is a strict linearly stable equilibrium point of τ -EG.

5.2 RELATION WITH TWO-TIMESCALE EG IN CONTINUOUS-TIME

For the stability analysis of continuous-time two-timescale EG (6) in terms of Definition 2(i), we need to examine $\text{spec}(\mathbf{J}_\tau^*)$, but computing \mathbf{J}_τ in general is cumbersome. However, there is a simple characterization of $\text{spec}(\mathbf{J}_\tau^*)$ in terms of $\text{spec}(\mathbf{H}_\tau^*)$, where $\mathbf{H}_\tau^* := \mathbf{H}_\tau(\mathbf{z}^*)$. For $s < 1/L$, define a subset of the complex plane $\bar{\mathcal{D}}_s := \{z \in \mathbb{C} : |z + 1/2s| \leq 1/2s\}$, which is a closed disk centered at $-1/2s$ with radius $1/2s$. Then we have the following fact, implying that the stability analysis of the ODE (6) can be done by examining $\text{spec}(\mathbf{H}_\tau^*)$. See Appendix G.2 for a visualization of this statement, with a comparison to the continuous-time GDA.

Proposition 5.1. Under Assumption 1, $\text{spec}(\mathbf{J}_\tau^*) \subset \mathbb{C}_-^\circ$ if and only if $\text{spec}(\mathbf{H}_\tau^*) \cap \bar{\mathcal{D}}_s = \emptyset$.

We are now ready to state our main stability result of two-timescale EG in terms of the refined second-order necessary condition. This states that under Assumption 2', for a step size s taken not too small, a point satisfying the refined second-order necessary condition is strict linearly stable.

Theorem 5.2. Suppose Assumptions 1 and 2'. Let $s_0 := \max_{j \in \mathcal{I}}(-\iota_j)$, for ι_j defined in (28). Then, an equilibrium point \mathbf{z}^* satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $s_0 < 1/L$ if and only if there exists a constant $0 < s^* < 1/L$ such that \mathbf{z}^* is a strict linearly stable point of ∞ -EG for any step size $s \in (s^*, 1/L)$.

Proof. Let us briefly sketch the proof, whose detail is deferred to Appendix G.4. For any $j \in \mathcal{I}$, using the hemicurvature ι_j in (5), it is possible to determine whether the eigenvalues λ_j of \mathbf{H}_τ^* are in $\bar{\mathcal{D}}_s$ or not, using Lemma G.4. Then, combining this result with the behaviors of λ_j where $j \notin \mathcal{I}$, which is established in Theorem 4.3, we can characterize exactly when does $\text{spec}(\mathbf{H}_\tau^*) \cap \bar{\mathcal{D}}_s = \emptyset$ hold. The conclusion then follows from Proposition 5.1 and Definition 2(i). \square

Although choosing step size s close enough to $1/L$ would suffice in practice, it is possible that s may not be large enough in view of Theorem 5.2. Furthermore, the condition $s_0 < 1/L$ may not hold, and even checking such condition is challenging as one needs to compute all the hemicurvatures ι_j . As a complement, we provide stability conditions that do not depend on s_0 (and consequently on s^*).

Theorem 5.3. Let \mathbf{z}^* be an equilibrium point. Suppose Assumptions 1 and 2' hold.

- (i) (Sufficient condition) Suppose that \mathbf{z}^* satisfies $\mathbf{S} \succeq \mathbf{0}$ and $\mathbf{B} \preceq \mathbf{0}$. Then for any step size $0 < s < 1/L$, the point \mathbf{z}^* is a strict linearly stable point of ∞ -EG.
- (ii) (Necessary condition) Suppose that \mathbf{z}^* is a strict linearly stable equilibrium point of ∞ -EG for any step size $0 < s < 1/L$. Then, it satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$ and $\mathbf{B} \preceq \mathbf{0}$.

If we introduce a mild additional assumption that all σ_j values are distinct, Theorem 5.3 can be refined into a tight necessary and sufficient condition. This, however, excludes local minimax points that satisfy $\mathbf{u}_j^\top \mathbf{S} \mathbf{u}_j < 0$ for some j . We leave treating such points for any given s as a future work.

Theorem 5.4. Under Assumptions 1 and 2', suppose that all σ_j values are distinct. Then, an equilibrium \mathbf{z}^* satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $\mathbf{u}_j^\top \mathbf{S} \mathbf{u}_j \geq 0$ for every left singular vector \mathbf{u}_j of \mathbf{C}_2 if and only if \mathbf{z}^* is a strict linearly stable point of ∞ -EG for any step size $0 < s < 1/L$.

5.3 RELATION WITH TWO-TIMESCALE EG IN DISCRETE-TIME

For the stability analysis of the discrete-time two-timescale EG (7), this time in terms of Definition 2(ii), we need to examine $\text{spec}(\mathbf{J}_\tau^*)$, and this can be similarly characterized in terms of

$\text{spec}(\mathbf{H}_\tau^*)$. For $\eta < 1/L$, define a peanut-shaped subset of the complex plane

$$\mathcal{P}_\eta := \{z = x + iy \in \mathbb{C} : (\eta x - 1/2)^2 + \eta^2 y^2 + 3/4 < \sqrt{1 + 3\eta^2 y^2}\}, \quad (8)$$

containing a subset of $i\mathbb{R} \setminus \{0\}$, as shown in Appendix H.1. Then we have the following. See Appendix H.2 for a visualization of this statement, with a comparison to the discrete-time GDA.

Proposition 5.5. *Under Assumption 1, $\rho(\mathbf{J}_\tau^*) < 1$ if and only if $\text{spec}(\mathbf{H}_\tau^*) \subset \mathcal{P}_\eta$.*

We then analyze the stability of (7), using arguments similar to the continuous-time counterparts.

Theorem 5.6. *Suppose Assumptions 1 and 2' hold. Then, an equilibrium point \mathbf{z}^* satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $s_0 < 1/2L$ if and only if there exists a constant $0 < \eta^* < 1/L$ such that \mathbf{z}^* is a strict linearly stable point of ∞ -EG for any step size η satisfying $\eta^* < \eta < 1/L$.*

In Appendix H.5, we also provide stability conditions that do not depend on s_0 (and consequently on η^*), similar to the continuous-time analyses in Theorems 5.3 and 5.4.

Example 3. *Consider the bilinear function $f_b(x, y) = xy$ in Example 1, and recall that $\text{spec}(\mathbf{H}_\tau) = \{\pm i\sqrt{\epsilon}\}$. Notice that $\text{spec}(\mathbf{H}_\tau) \cap \overline{\mathcal{D}}_s = \emptyset$ for any $s > 0$, and $\text{spec}(\mathbf{H}_\tau) \subset \mathcal{P}_\eta$ whenever η is sufficiently small. Therefore, as mentioned in Example 2, the local minimax point $\mathbf{0}$ is strict linearly stable for both continuous- and discrete-time two-timescale EG.*

5.4 TWO-TIMESCALE EG ALMOST ALWAYS AVOIDS STRICT NON-MINIMAX POINTS

Using Theorem 2.2, we show that two-timescale EG almost surely avoids strict non-minimax points.

Theorem 5.7. *Let \mathbf{z}^* be a strict non-minimax point, i.e., $\mathbf{z}^* \in \mathcal{T}^*$. Under Assumptions 1, 2, and $0 < \eta < (\sqrt{5}-1)/2L$, there exists $\tau^* > 0$ such that for any $\tau > \tau^*$, the set of initial points that converge to \mathbf{z}^* by two-timescale (discrete) EG \mathbf{w}_τ has measure zero. Moreover, if \mathcal{T}^* is finite⁴, then there exists $\tau^* > 0$ such that for any $\tau > \tau^*$, $\mu(\{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \mathbf{w}_\tau^k(\mathbf{z}_0) \in \mathcal{T}^*\}) = 0$.*

This, however, does not mean that strict non-minimax points are always the only points that two-timescale EG almost surely escapes, as one would desire. Considering Theorems 5.6 and 5.7, this is possible if $s_0 < 1/2L$ holds and we can choose the step size so that $\eta^* < \eta < (\sqrt{5}-1)/2L$. Nevertheless, these conditions are somewhat restrictive and we leave removing them as future work.

6 TWO-TIMESCALE EG GLOBALLY FINDS LOCAL MINIMAX POINTS

Our analysis has thus far been local, but can be extended to a global statement. One can show that any accumulation point of the iterates computed by two-timescale EG is a stationary point, for instance under the *Minty variational inequality* (MVI) condition (Minty, 1967) which encompasses nonconvex-nonconcave settings; see Appendix J. With appropriate settings on s_0 , η , and \mathcal{T}^* , Theorem 5.7 guarantees almost sure avoidance of strict non-minimax points. Hence, under these settings, with Assumptions 1 and 2 and assuming that the two-timescale EG converges and there are no non-strict non-minimax points, akin to the case of gradient descent as in (Lee et al., 2016, Theorem 11), we can conclude that the two-timescale EG globally and almost surely finds local minimax points.

As a consequential result, we now have that two-timescale EG is also capable of finding local minimax points even with a singular $D\mathbf{F}$, i.e., some λ_j of \mathbf{H}_τ in Theorem 4.3 are zero, as demonstrated in the following example with non-isolated local minimax points.

Example 4. *Consider the quadratic function f_q in Example 1. Similar to the case of Example 2, since $\text{spec}(\mathbf{H}_\tau) = \{0, 2, \epsilon, \pm i\sqrt{\epsilon}\}$, all (non-strict) local minimax points are strict linearly unstable for discrete-time two-timescale GDA. Yet, while $0 \in \text{spec}(\mathbf{H}_\tau)$, both $\mathbb{C} \setminus \overline{\mathcal{D}}_s$ and \mathcal{P}_η never contain 0. Thus, unlike Example 3, the optima are not strict linearly stable for two-timescale EG. However, recall that a point not being strict linearly stable does not imply that the point is strict linearly unstable. Indeed, for this f_q we have $s_0 = 0$ because $\{\lambda_j : j \in \mathcal{I}\} = \{\pm i\sqrt{\epsilon}\}$, and thus $\eta^* \leq 0$. Therefore, the local minimax points are not avoided by two-timescale EG. In fact, any stationary point of this f_q is a local optimal point that satisfies the MVI condition (see Lemma J.3), hence the local optimal points are rather found by discrete-time two-timescale EG.*

We defer the final discussions and a review on the related works to the appendices.

⁴If we assume that all points in \mathcal{T}^* are isolated and the iterates are bounded, then \mathcal{T}^* is essentially finite.

ACKNOWLEDGMENTS

This work was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2019R1A5A1028324, 2022R1C1C1003940), and the Samsung Science & Technology Foundation grant (No. SSTF-BA2101-02).

REFERENCES

- K. J. Arrow, L. Hurwicz, and H. Uzawa. *Studies in Linear and Non-Linear Programming*, volume 2 of *Stanford Mathematical Studies in the Social Sciences*. Stanford University Press, 1958.
- H. Attouch and B. F. Svaiter. A continuous dynamical Newton-like approach to solving monotone inclusions. *SIAM J. Control Optim.*, 49(2):574–98, 2011.
- A. Beck. *First-Order Methods in Optimization*. SIAM, 2017.
- J. Chae, K. Kim, and D. Kim. Open problem: Is there a first-order method that only converges to local minimax optima? In *Conference on Learning Theory*, pp. 5957–5964. PMLR, 2023.
- Z. Chen, Z. Hu, Q. Li, Z. Wang, and Y. Zhou. A cubic regularization approach for finding local minimax points in nonconvex minimax optimization. *Transactions on Machine Learning Research*, 2023.
- C. Daskalakis and I. Panageas. The limit points of (optimistic) gradient descent in min-max optimization. In *Neural Info. Proc. Sys.*, 2020.
- J. Diakonikolas, C. Daskalakis, and M. Jordan. Efficient methods for structured nonconvex-nonconcave min-max optimization. In *International Conference on Artificial Intelligence and Statistics*, pp. 2746–2754. PMLR, 2021.
- M. P. do Carmo. *Differential Geometry of Curves and Surfaces*. Dover Publications, revised and updated second edition, 2016.
- Y. G. Evtushenko. Some local properties of minimax problems. *USSR Comp. Math. and Math. Phys.*, 14(3):129–138, 1974.
- F. Farnia and A. Ozdaglar. Do GANs always have Nash equilibria? In *Proc. Intl. Conf. Mach. Learn.*, 2020.
- T. Fiez and L. Ratliff. Local convergence analysis of gradient descent ascent with finite timescale separation. In *Proc. Intl. Conf. on Learning Representations*, 2021.
- T. Fiez, B. Chasnov, and L. Ratliff. Implicit learning dynamics in Stackelberg games: equilibria characterization, convergence analysis, and empirical study. In *Proc. Intl. Conf. Mach. Learn.*, 2020.
- P. Foret, A. Kleiner, H. Mobahi, and B. Neyshabur. Sharpness-aware minimization for efficiently improving generalization. In *Proc. Intl. Conf. on Learning Representations*, 2021.
- O. Galor. *Discrete Dynamical Systems*. Springer, 2007.
- I. Goodfellow. NIPS 2016 Tutorial: Generative adversarial networks, 2016. arXiv 1701.00160.
- I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative adversarial networks. In *Neural Info. Proc. Sys.*, 2014.
- E. Gorbunov, A. B. Taylor, S. Horváth, and G. Gidel. Convergence of proximal point and extragradient-based methods beyond monotonicity: the case of negative comonotonicity. In *Proc. Intl. Conf. Mach. Learn.*, 2023.
- M. Heusel, H. Ramsauer, T. Unterthiner B. Nessler, and S. Hochreiter. GANs trained by a two time-scale update rule converge to a local Nash equilibrium. In *Neural Info. Proc. Sys.*, 2017.
- N. J. Higham. *Functions of Matrices*. SIAM, 2008.

- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2nd edition, 2012.
- C. Jin, P. Netrapalli, and M. I. Jordan. What is local optimality in nonconvex-nonconcave minimax optimization? In *Proc. Intl. Conf. Mach. Learn*, 2020.
- T. Kato. *Perturbation Theory for Linear Operators*. Springer, 2nd edition, 1995.
- H. K. Khalil. *Nonlinear Systems*. Pearson, 3rd edition, 2002.
- G. M. Korpelevich. An extragradient method for finding saddle points and other problems. *Ekonomika i Mateaticheskie Metody*, 12(4):747–56, 1976.
- J. D. Lee, M. Simchowitz, M. I. Jordan, and B. Recht. Gradient descent only converges to minimizers. In *Conference on Learning Theory*, pp. 1246–1257. PMLR, 2016.
- J. D. Lee, I. Panageas, G. Piliouras, M. Simchowitz, M. I. Jordan, and B. Recht. First-order methods almost always avoid strict saddle points. *Mathematical Programming*, 176:311–337, 2019.
- C. Liu, L. Zhu, and M. Belkin. Loss landscapes and optimization in over-parameterized non-linear systems and neural networks. *Applied and Computational Harmonic Analysis*, 59:85–116, 2022.
- H. Lu. An $O(s^r)$ -resolution ODE framework for understanding discrete-time algorithms and applications to the linear convergence of minimax problems. *Mathematical Programming*, 194:1061–1112, 2022.
- X. Ma, W. Yao, J. J. Ye, and J. Zhang. Calm local optimality for nonconvex-nonconcave minimax problems, 2023. arXiv 2306.17443.
- P. Mahdavinia, Y. Deng, H. Li, and M. Mahdavi. Tight analysis of extra-gradient and optimistic gradient methods for nonconvex minimax problems. In *Neural Info. Proc. Sys.*, 2022.
- L. Mescheder, A. Geiger, and S. Nowozin. Which training methods for GANs do actually converge? In *Proc. Intl. Conf. Mach. Learn*, 2018.
- G. J. Minty. On the generalization of a direct method of the calculus of variations. *Bull. Amer. Math. Soc*, 73:314–321, 1967.
- A. Mądry, A. Makelov, L. Schmdit, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. In *Proc. Intl. Conf. on Learning Representations*, 2018.
- E. K. Ryu and W. Yin. *Large-Scale Convex Optimization*. Cambridge University Press, 2022.
- E. K. Ryu, K. Yuan, and W. Yin. ODE analysis of stochastic gradient methods with optimism and anchoring for minimax problems and GANs, 2019. arXiv 1905.10899.
- M. Shub. *Global Stability of Dynamical Systems*. Springer, 1987.
- M. Sion. On general minimax theorems. *Pacific J. Math.*, 8(1):171–6, 1958.
- J. Stoer and R. Bulirsch. *Introduction to Numerical Analysis*. Springer, 3rd edition, 2002.
- T. Tao. *Analysis II*. Springer, third edition, 2016.
- H. von Stackelberg. *Market Structure and Equilibrium*. Springer, 2011.
- Y. Wang, G. Zhang, and J. Ba. On solving minimax optimization locally: A follow-the-ridge approach. In *Proc. Intl. Conf. on Learning Representations*, 2020.
- S. J. Wright and B. Recht. *Optimization for Data Analysis*. Cambridge University Press, 2022.
- M. Zedek. Continuity and location of zeros of linear combinations of polynomials. *Proceedings of the American Mathematical Society*, 16(1):78–84, 1965.
- G. Zhang, K. Wu, P. Poupart, and Y. Yu. Newton-type methods for minimax optimization, 2020. arXiv 2006.14592.
- G. Zhang, P. Poupart, and Y. Yu. Optimality and stability in non-convex smooth games. *J. Mach. Learn. Res.*, 23:35–1, 2022.

CONTENTS

1	Introduction	1
2	Preliminaries	2
2.1	Notations and problem setting	2
2.2	Local minimax points	2
2.3	Gradient descent ascent and extragradient	3
2.4	Dynamical systems	3
3	On the necessary condition of local minimax points	3
3.1	Restricted Schur complement	4
3.2	Refining the second-order necessary condition	4
4	Characterizing timescale separation without nondegeneracy condition on $\nabla_{yy}^2 f$	5
4.1	Timescale separation in GDA and its relation to stability	5
4.2	Timescale separation without the nondegeneracy condition on $\nabla_{yy}^2 f$	6
4.2.1	Hemicurvature of the eigenvalue function $\lambda_j(\epsilon)$	6
4.3	Two-timescale GDA avoids some non-strict local minimax points	7
5	Local minimax points and the limit points of two-timescale EG	7
5.1	Two-timescale EG	7
5.2	Relation with two-timescale EG in continuous-time	8
5.3	Relation with two-timescale EG in discrete-time	8
5.4	Two-timescale EG almost always avoids strict non-minimax points	9
6	Two-timescale EG globally finds local minimax points	9
A	Discussions	13
B	Related works	14
C	Proofs and Missing Details for Section 2	14
C.1	Proof of Lemma 2.1	14
C.2	Further discussions on the new ODE	14
D	Proofs for Section 3	15
D.1	Proof of Proposition 3.1	15
D.2	Proof of Proposition 3.2	15
D.3	Proof of Example 1	16
E	Proofs and Missing Details for Section 4	16
E.1	Proof of Theorem 4.3	16

E.2	Proofs on the properties of the hemicurvature	21
E.3	On the necessity of considering the hemicurvature	22
E.4	Proof of Theorem 4.4	23
F	Proof for Section 5.1	23
G	Proofs and Missing Details for Section 5.2	24
G.1	Proof of Proposition 5.1	24
G.2	Target sets of continuous-time methods	25
G.3	A consequence of Assumption 2'	26
G.4	Proof of Theorem 5.2	26
G.5	Proof of Theorem 5.3	27
G.6	Proof of Theorem 5.4	30
H	Proofs and Missing Details for Section 5.3	33
H.1	On the peanut-shaped region \mathcal{P}_η	33
H.2	Target sets of discrete-time methods	34
H.3	Proof of Proposition 5.5	34
H.4	Proof of Theorem 5.6	35
H.5	Additional stability conditions of two-timescale EG in discrete-time	36
H.6	Proof of Theorem H.3	36
H.7	Proof of Theorem H.4	37
I	Proof for Section 5.4	37
I.1	Proof of Theorem 5.7	37
J	Global convergence analysis of two-timescale EG	38

A DISCUSSIONS

We demonstrated that two-timescale extragradient converges to local minimax points without a non-degeneracy condition on Hessian with respect to the maximization variable, from a dynamical system viewpoint. We have then shown that two-timescale extragradient almost surely avoids strict non-minimax points, using a theorem built upon the stable manifold theorem. To the best of our knowledge, this is the first result in minimax optimization that attempts to fully characterize the stability behavior near stationary points, comparable to that of gradient descent in [Lee et al. \(2016\)](#).

Still, our work is just a takeoff in exploring first-order methods for finding non-strict local minimax points, leaving several unanswered questions. While our results state that they hold whenever the timescale parameter τ is sufficiently large, we do not discuss how large exactly should τ be. A slight gap remains between the second order necessary and sufficient conditions of local minimax points, due to the additional condition on h given in Proposition 3.2. The analysis of the limit points of two-timescale EG involves conditions on \mathcal{S}_{res} and s_0 . For further discussion, we advise the reader to refer to ([Chae et al., 2023](#)). Investigations on these issues will lead to interesting future works.

B RELATED WORKS

Equilibria in minimax problems Nash equilibrium is a widely used notion of (local) optimality in minimax optimization, especially when minimax problems are interpreted as simultaneous games, where both players act simultaneously (Jin et al., 2020). However, there are problems that do not have a local Nash equilibrium point including realistic GAN applications (Farnia & Ozdaglar, 2020). In fact, the correct way of interpreting minimax problems, especially in modern machine learning applications, are as sequential games (Jin et al., 2020). These necessitated a new notion of *local* optimality that reflects the sequential nature of minimax optimization. The notion of local optimality in sequential games, built upon Stackelberg equilibrium (von Stackelberg, 2011), was first introduced by Evtushenko (1974), but this was later found to be not *truly* local (Jin et al., 2020). Instead, Fiez et al. (2020) and Jin et al. (2020) considered another notion of (true) local optimum, called *strict* local minimax points, which are exactly the points satisfying the sufficient condition of local optimum proposed by Evtushenko (1974). This, however, only covers the cases where $\nabla_{\mathbf{y}\mathbf{y}}^2 f$ is nondegenerate, hence disregards possibly meaningful “non-strict” local optimal points, *e.g.*, in over-parameterized training (Liu et al., 2022).

So, Jin et al. (2020) also introduced a more general definition of local optimum, named local minimax points; see Section 2.2 for more details. To the best of our knowledge, all existing methods in Evtushenko (1974); Fiez et al. (2020); Wang et al. (2020); Zhang et al. (2020); Chen et al. (2023), even that in Jin et al. (2020), focus on finding only *strict* local minimax points.

Two-timescale methods Plain GDA can converge to nonoptimal points (Daskalakis & Panageas, 2020; Jin et al., 2020). So, Jin et al. (2020) adopted a (computationally cheap) timescale separation into GDA, *i.e.*, using different timescales, or step sizes, for each player’s action. It is proven that two-timescale GDA with a sufficiently large timescale separation converges to strict local minimax points (Jin et al., 2020; Fiez & Ratliff, 2021). Recently, two-timescale EG was studied in Zhang et al. (2022); Mahdavinia et al. (2022), but it was only shown to behave similar to the two-timescale GDA; this paper demonstrates the particular usefulness of two-timescale EG when $\nabla_{\mathbf{y}\mathbf{y}}^2 f$ is degenerate.

C PROOFS AND MISSING DETAILS FOR SECTION 2

C.1 PROOF OF LEMMA 2.1

Letting $s = \eta/2$, the ODE derived by Lu (2022), mentioned in Section 2.3, can be written as

$$\begin{aligned}\dot{\mathbf{z}}(t) &= -\mathbf{F}(\mathbf{z}(t)) + \frac{\eta}{2}D\mathbf{F}(\mathbf{z}(t))\mathbf{F}(\mathbf{z}(t)) \\ &= -(\mathbf{I} - sD\mathbf{F}(\mathbf{z}(t)))\mathbf{F}(\mathbf{z}(t)).\end{aligned}\tag{9}$$

As we assume that $0 < s < \frac{1}{L}$, by Assumption 1, the matrix $\mathbf{I} + sD\mathbf{F}(z)$ is invertible. Applying the approximation $(\mathbf{I} + sD\mathbf{F}(\mathbf{z}(t)))^{-1} = \mathbf{I} - sD\mathbf{F}(\mathbf{z}(t)) + O(s^2)$, the ODE

$$\dot{\mathbf{z}}(t) = -(\mathbf{I} + sD\mathbf{F}(\mathbf{z}(t)))^{-1}\mathbf{F}(\mathbf{z}(t))\tag{10}$$

can then be seen as an approximation of (9).

C.2 FURTHER DISCUSSIONS ON THE NEW ODE

Let $\mathbf{g}(t) := \mathbf{F}(\mathbf{z}(t))$. Using the chain rule, we have

$$D\mathbf{F}(\mathbf{z}(t))\dot{\mathbf{z}}(t) = \dot{\mathbf{g}}(t).$$

Therefore, (10) is equivalent to

$$\begin{aligned}-\mathbf{g}(t) &= -\mathbf{F}(\mathbf{z}(t)) = (\mathbf{I} + sD\mathbf{F}(\mathbf{z}(t)))\dot{\mathbf{z}}(t) \\ &= \dot{\mathbf{z}}(t) + s\dot{\mathbf{g}}(t).\end{aligned}\tag{11}$$

We would like to mention that special case of (11) when $s = 1$ has originally been studied as a continuous-time limit of the Newton’s method by Attouch & Svaiter (2011), and also as that of the optimistic GDA by Ryu et al. (2019).

D PROOFS FOR SECTION 3

D.1 PROOF OF PROPOSITION 3.1

We use \mathbf{P} , $\mathbf{\Gamma}$ and \mathbf{U} that are introduced above Definition 4. Interpreting \mathbf{P} as a change-of-basis matrix, we have $\mathbf{C}^\top \mathbf{v} \in \mathcal{R}(\mathbf{B})$ if and only if the last $(d_2 - r)$ components of $\mathbf{P}^\top \mathbf{C}^\top \mathbf{v}$ are zero. The latter holds if and only if \mathbf{v} is orthogonal to the last $(d_2 - r)$ columns of $\mathbf{C}\mathbf{P}$, or in other words, if and only if $\mathbf{v} \perp \mathcal{R}(\mathbf{\Gamma})$.

Let $\mathcal{W} := \mathcal{R}(\mathbf{\Gamma})^\perp$ and $w := \dim \mathcal{W}$. If $w = 0$, then this means that the zero vector is the only vector such that $\mathbf{C}^\top \mathbf{v} \in \mathcal{R}(\mathbf{B})$, so there is nothing to show. Assuming $w \geq 1$, by how \mathbf{U} is constructed, we have $\mathbf{v} \in \mathcal{W}$ if and only if there exists $\mathbf{w} \in \mathbb{R}^w$ such that $\mathbf{v} = \mathbf{U}\mathbf{w}$. Hence, the condition $\mathbf{v}^\top (\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger \mathbf{C}^\top) \mathbf{v} \geq 0$ for any \mathbf{v} satisfying $\mathbf{C}^\top \mathbf{v} \in \mathcal{R}(\mathbf{B})$ is equivalent to having

$$\mathbf{w}^\top \mathbf{U}^\top (\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger \mathbf{C}^\top) \mathbf{U}\mathbf{w} = \mathbf{w}^\top \mathbf{S}_{\text{res}} \mathbf{w} \geq 0$$

for any $\mathbf{w} \in \mathbb{R}^w$.

D.2 PROOF OF PROPOSITION 3.2

Proof. Let $\mathbf{A} := \nabla_{\mathbf{x}\mathbf{x}}^2 f(\mathbf{x}^*, \mathbf{y}^*)$, $\mathbf{B} := \nabla_{\mathbf{y}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*)$, and $\mathbf{C} := \nabla_{\mathbf{x}\mathbf{y}}^2 f(\mathbf{x}^*, \mathbf{y}^*)$. Since \mathbf{y}^* is a local maximum of $f(\mathbf{x}^*, \cdot)$, it holds that $\mathbf{B} \preceq \mathbf{0}$.

Because $(\mathbf{x}^*, \mathbf{y}^*)$ is a local minimax point, there exists a function h such that (1) holds. Moreover, by the assumption we made on h , there exist δ_1 and a constant M such that $\frac{h(\delta)}{\delta} \leq M$ whenever $0 < \delta \leq \delta_1$. We may then replace δ_0 in Definition 1 by $\min\{\delta_0, \delta_1\}$ without affecting the local optimality. In other words, without loss of generality we may assume that

$$\frac{h(\delta)}{\delta} \leq M \quad \text{whenever} \quad 0 < \delta \leq \delta_0 \quad (12)$$

holds in Definition 1. Let us denote $h'(\delta) = \max\{h(\delta), 2\|\mathbf{B}^\dagger \mathbf{C}^\top\|\delta\}$.

Using the first order necessary condition, we have

$$f(\mathbf{x}^* + \delta_{\mathbf{x}}, \mathbf{y}^* + \delta_{\mathbf{y}}) = f(\mathbf{x}^*, \mathbf{y}^*) + \frac{1}{2} \delta_{\mathbf{x}}^\top \mathbf{A} \delta_{\mathbf{x}} + \delta_{\mathbf{x}}^\top \mathbf{C} \delta_{\mathbf{y}} + \frac{1}{2} \delta_{\mathbf{y}}^\top \mathbf{B} \delta_{\mathbf{y}} + R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}}) \quad (13)$$

where $R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}})$ is the remainder term of a degree 2 Taylor series expansion. Now let us assume that $\|\delta_{\mathbf{x}}\| = \delta \in (0, \delta_0]$ and $\mathbf{C}^\top \delta_{\mathbf{x}} \in \mathcal{R}(\mathbf{B})$. Then, by the second order necessary assumption $\mathbf{B} \preceq \mathbf{0}$, it holds that

$$-\mathbf{B}^\dagger \mathbf{C}^\top \delta_{\mathbf{x}} = \underset{\delta_{\mathbf{y}}}{\operatorname{argmax}} f(\mathbf{x}^*, \mathbf{y}^*) + \frac{1}{2} \delta_{\mathbf{x}}^\top \mathbf{A} \delta_{\mathbf{x}} + \delta_{\mathbf{x}}^\top \mathbf{C} \delta_{\mathbf{y}} + \frac{1}{2} \delta_{\mathbf{y}}^\top \mathbf{B} \delta_{\mathbf{y}}.$$

Meanwhile, let $\epsilon > 0$ be arbitrary. Then, noticing that

$$R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}}) = o(\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2),$$

we know that there exists $\delta_2 > 0$ such that

$$\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2 < \delta_2 \implies \left| \frac{R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}})}{\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2} \right| < \epsilon.$$

Here, if $\|\delta_{\mathbf{y}}\| \leq h(\delta)$, then by (12) and that $\delta \leq \delta_0$ we have

$$\frac{\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2}{\|\delta_{\mathbf{x}}\|^2} = 1 + \frac{\|\delta_{\mathbf{y}}\|^2}{\|\delta_{\mathbf{x}}\|^2} \leq 1 + \left(\frac{h(\delta)}{\delta} \right)^2 \leq 1 + M^2 \quad (14)$$

which then implies

$$\left| \frac{R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}})}{\|\delta_{\mathbf{x}}\|^2} \right| = \frac{\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2}{\|\delta_{\mathbf{x}}\|^2} \left| \frac{R(\delta_{\mathbf{x}}, \delta_{\mathbf{y}})}{\|\delta_{\mathbf{x}}\|^2 + \|\delta_{\mathbf{y}}\|^2} \right| \leq (1 + M^2)\epsilon. \quad (15)$$

Note that, by (14), the condition $\|\delta_x\|^2 + \|\delta_y\|^2 < \delta_2$ holds whenever $\|\delta_x\| \leq \frac{\delta_2}{\sqrt{1+M^2}}$. Thus, since ϵ was chosen arbitrarily, we conclude that $R(\delta_x, \delta_y) = o(\|\delta_x\|^2)$, as long as $\|\delta_y\| \leq h(\delta)$.

Therefore, again invoking that (x^*, y^*) is a local minimax point, we have

$$\begin{aligned} 0 &\leq \max_{\|\delta_y\| \leq h(\delta)} f(x^* + \delta_x, y^* + \delta_y) - f(x^*, y^*) \\ &\leq \max_{\|\delta_y\| \leq h(\delta)} \left(\frac{1}{2} \delta_x^\top A \delta_x + \delta_x^\top C \delta_y + \frac{1}{2} \delta_y^\top B \delta_y \right) + \max_{\|\delta_y\| \leq h(\delta)} R(\delta_x, \delta_y) \\ &\leq \max_{\|\delta_y\| \leq h'(\delta)} \left(\frac{1}{2} \delta_x^\top A \delta_x + \delta_x^\top C \delta_y + \frac{1}{2} \delta_y^\top B \delta_y \right) + o(\|\delta_x\|^2) \\ &= \frac{1}{2} \delta_x^\top (A - CB^\dagger C^\top) \delta_x + o(\|\delta_x\|^2). \end{aligned}$$

This inequality holds for any δ_x satisfying $C^\top \delta_x \in \mathcal{R}(B)$, so we are done. \square

D.3 PROOF OF EXAMPLE 1

- $f_b(x, y) = xy$: The saddle-gradient of f_b is $F(x, y) = (x, -y)$, where $(x^*, y^*) = (0, 0)$ is a unique stationary point. Since $f_b(x^*, y) = f_b(x^*, y^*) = f_b(x, y^*) = 0$ for any (x, y) , we can say that $(0, 0)$ is a unique local minimax point by definition.

Regarding Assumption 2, the Jacobian matrix of the saddle-gradient F is

$$DF = \begin{bmatrix} A & C \\ -C^\top & -B \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix},$$

where B is degenerate with $r = \text{rank}(B) = 0$. Since $d_1 = d_2 = 1$, $\Gamma = [1] \in \mathbb{R}^{1 \times 1}$ and $q = \text{rank}(\Gamma) = 1$, the matrix U is of size 1×0 . So, $S_{\text{res}}(DF)$ is then a 0×0 matrix. A linear map from a zero dimensional vector space to a zero dimensional vector space is trivially an identity map, so the corresponding 0×0 matrix is *trivially* nondegenerate.

- $f_q(x_1, x_2, y_1, y_2, y_3) = \frac{1}{2}x_1^2 - \frac{1}{2}y_1^2 - \frac{1}{2}y_2^2 + x_2y_2 + y_1y_3$: The saddle-gradient of f_q is $F(x_1, x_2, y_1, y_2, y_3) = (x_1, y_2, y_1 - y_3, -x_2, y_3 - y_1)$, where $\{t \in \mathbb{R} : (0, 0, t, 0, t)\}$ is a set of stationary points. Since the inequality

$$f_q(x^*, y) = -\frac{1}{2}(y_1 - y_3)^2 \leq f_q(x^*, y^*) = 0 \leq f_q(x, y^*) = \frac{1}{2}x_1^2$$

holds for any $(x^*, y^*) \in \{t \in \mathbb{R} : (0, 0, t, 0, t)\}$ and for any (x, y) , the set of local minimax points is by definition $\{t \in \mathbb{R} : (0, 0, t, 0, t)\}$.

Regarding Assumption 2, the Jacobian matrix of the saddle-gradient F is

$$DF = \begin{bmatrix} A & C \\ -C^\top & -B \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & -1 \\ 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 1 \end{bmatrix},$$

where B is degenerate with $r = \text{rank}(B) = 2$. Since $d_1 = 2$, $d_2 = 3$, $\Gamma = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \in \mathbb{R}^{2 \times 1}$, and $q = \text{rank}(\Gamma) = 1$, we have $U = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \in \mathbb{R}^{2 \times 1}$, and thus $S_{\text{res}}(DF) = [1] \in \mathbb{R}^{1 \times 1}$ is nondegenerate. Notice that for this example, the classical generalized Schur complement $S = A - CB^\dagger C^\top = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \in \mathbb{R}^{2 \times 2}$ is degenerate.

E PROOFS AND MISSING DETAILS FOR SECTION 4

E.1 PROOF OF THEOREM 4.3

This theorem exactly reduces to Lemma 4.1 when $\nabla_{yy}^2 f(z^*)$ is nondegenerate, it is only left to show this theorem for the case where $S_{\text{res}}(DF(z^*))$ is nondegenerate.

We begin with the following observation. Suppose that restricted Schur complement \mathbf{S}_{res} is invertible. Then, one can show that $D\mathbf{F}$ is similar to

$$\mathbf{G} = \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{C}_2 \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} \\ -\mathbf{C}_2^\top & \mathbf{0} & \mathbf{0} \end{bmatrix},$$

but \mathbf{C}_2 may not be of full column rank. Let $\text{rank}(\mathbf{C}_2) := q$, and let $\mathbf{C}_2 = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ be the (full) singular value decomposition of \mathbf{C}_2 where for some invertible diagonal matrix $\mathbf{\Sigma}_q$ it holds that

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_q & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Then, by defining $\mathbf{L}_q = \mathbf{U} \begin{bmatrix} \mathbf{\Sigma}_q \\ \mathbf{0} \end{bmatrix}$, notice that $\mathbf{U}\mathbf{\Sigma} = [\mathbf{L}_q \ \mathbf{0}]$. Thus, for $\tilde{\mathbf{Q}} = \text{diag}\{\mathbf{I}, \mathbf{I}, \mathbf{V}^\top\}$ we have

$$\tilde{\mathbf{Q}}\mathbf{G}\tilde{\mathbf{Q}}^\top = \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{C}_2\mathbf{V} \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} \\ -\mathbf{V}^\top\mathbf{C}_2^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{L}_q & \mathbf{0} \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} & \mathbf{0} \\ -\mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}$$

and therefore

$$\det(\lambda\mathbf{I} - D\mathbf{F}) = \det(\lambda\mathbf{I} - \tilde{\mathbf{Q}}\mathbf{G}\tilde{\mathbf{Q}}^\top) = \lambda^{d_2-r-q} \det \left(\lambda\mathbf{I} - \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{L}_q \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} \\ -\mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} \right).$$

So, the eigenvalues of $D\mathbf{F}$ are either 0 or the eigenvalues of

$$\Phi := \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{L}_q \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} \\ -\mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad (16)$$

and analogously the eigenvalues of \mathbf{H}_τ are either 0 or the eigenvalues of $\Phi_\tau := \Lambda_\tau\Phi$.

Moreover, by how \mathbf{L}_q is defined, it holds that the singular values of \mathbf{L}_q are exactly the singular values of \mathbf{C}_2 , and $\mathcal{R}(\mathbf{L}_q) = \mathcal{R}(\mathbf{C}_2)$. In particular, there is no essential change in $\mathbf{S}_{\text{res}}(\mathbf{H})$ when \mathbf{C}_2 is replaced by \mathbf{L}_q . Also, notice that \mathbf{L}_q is of full column rank.

Therefore, characterizing the eigenvalues of Φ_τ is equivalent to characterizing the nonzero eigenvalues of \mathbf{H}_τ .

To prove Theorem 4.3 we need the following result, which shows a relation between the eigenvalues of the restricted Schur complement and a determinantal equation that resembles the eigenvalue equation of Φ_τ .

Lemma E.1. *A (possibly complex) number μ is an eigenvalue of the restricted Schur complement if and only if it is a root of the equation*

$$\det \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} & -\mathbf{C}_1 & -\mathbf{L}_q \\ \mathbf{C}_1^\top & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} = 0. \quad (17)$$

Proof. By the property of Moore-Penrose pseudoinverses, it holds that

$$\mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top = [\mathbf{C}_1 \ \mathbf{C}_2] \begin{bmatrix} -\mathbf{D}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{C}_1^\top \\ \mathbf{C}_2^\top \end{bmatrix} = -\mathbf{C}_1\mathbf{D}^{-1}\mathbf{C}_1^\top.$$

Multiplying matrices of determinant 1 does not change the determinant, so by observing

$$\begin{aligned}
& \begin{bmatrix} \mathbf{I} & -\mathbf{C}_1\mathbf{D}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} & -\mathbf{C}_1 & -\mathbf{L}_q \\ \mathbf{C}_1^\top & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{D}^{-1}\mathbf{C}_1^\top & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \\
&= \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} - \mathbf{C}_1\mathbf{D}^{-1}\mathbf{C}_1^\top & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} \\
&= \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} + \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} \\
&= \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}
\end{aligned}$$

we get that (17) is equivalent to the determinantal equation $\det \mathbf{T} = 0$ where

$$\mathbf{T} := \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}. \quad (18)$$

Since permuting the rows and the columns change the determinant only up to a sign, we have

$$\begin{aligned}
|\det \mathbf{T}| &= \left| \det \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} \right| \\
&= \left| \det \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & \mathbf{0} & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{D} & \mathbf{0} \end{bmatrix} \right| \\
&= \left| \det \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & -\mathbf{L}_q & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & -\mathbf{D} \end{bmatrix} \right|.
\end{aligned}$$

As \mathbf{D} is invertible by assumption, we obtain that $\det \mathbf{T} = 0$ is further equivalent to

$$\det \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \mathbf{0} \end{bmatrix} = 0. \quad (19)$$

Now we choose an orthogonal matrix \mathbf{Q} such that it has a block matrix form

$$\mathbf{Q} = [\mathbf{U}_1 \quad \mathbf{U}_2]$$

where \mathbf{U}_2 is the $d_1 \times q$ matrix of left singular vectors appearing in the reduced singular value decomposition of $\mathbf{L}_q = \mathbf{U}_2\mathbf{\Sigma}_q\mathbf{V}_2^\top$, and \mathbf{U}_1 is consequently a $d_1 \times (d_1 - q)$ matrix with the columns consisting an orthonormal basis of $\mathcal{R}(\mathbf{L}_q)^\perp$. Note that \mathbf{U}_1 corresponds to the matrix \mathbf{U} that appears in the restricted Schur complement, and considering how \mathbf{L}_q is defined, we may take $\mathbf{V}_2 = \mathbf{I}$. Applying a similarity transformation, we have

$$\begin{bmatrix} \mathbf{Q}^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mu\mathbf{I} - \mathbf{S} & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{Q} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} = \begin{bmatrix} \mu\mathbf{I} - \mathbf{Q}^\top\mathbf{S}\mathbf{Q} & -\mathbf{Q}^\top\mathbf{L}_q \\ \mathbf{L}_q^\top\mathbf{Q} & \mathbf{0} \end{bmatrix}. \quad (20)$$

Exploiting the block structure of \mathbf{U} , observe that

$$\mathbf{Q}^\top\mathbf{L}_q = \begin{bmatrix} \mathbf{U}_1^\top \\ \mathbf{U}_2^\top \end{bmatrix} \mathbf{L}_q = \begin{bmatrix} \mathbf{U}_1^\top\mathbf{L}_q \\ \mathbf{U}_2^\top\mathbf{L}_q \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{\Sigma}_q\mathbf{V}_2^\top \end{bmatrix} = \begin{bmatrix} \mathbf{0} \\ \mathbf{\Sigma}_q \end{bmatrix}$$

by the orthogonality of $\mathcal{R}(U_1)$ and $\mathcal{R}(L_q)$, and

$$Q^\top S Q = \begin{bmatrix} U_1^\top \\ U_2^\top \end{bmatrix} S [U_1 \ U_2] = \begin{bmatrix} U_1^\top S U_1 & U_1^\top S U_2 \\ U_2^\top S U_1 & U_2^\top S U_2 \end{bmatrix}.$$

Thus, the matrix in (20) can be also expressed as

$$M := \begin{bmatrix} \mu I - U_1^\top S U_1 & -U_1^\top S U_2 & \mathbf{0} \\ -U_2^\top S U_1 & \mu I - U_2^\top S U_2 & -\Sigma_q \\ \mathbf{0} & \Sigma_q & \mathbf{0} \end{bmatrix}.$$

Similarity transformations preserve determinants, so (19) holds if and only if $\det M = 0$. Using again that the fact that permuting the rows and the columns change the determinant only up to a sign, we obtain

$$\begin{aligned} |\det M| &= \left| \det \begin{bmatrix} \mu I - U_1^\top S U_1 & -U_1^\top S U_2 & \mathbf{0} \\ -U_2^\top S U_1 & \mu I - U_2^\top S U_2 & -\Sigma_q \\ \mathbf{0} & \Sigma_q & \mathbf{0} \end{bmatrix} \right| \\ &= \left| \det \begin{bmatrix} \mathbf{0} & \Sigma_q & \mathbf{0} \\ \mu I - U_1^\top S U_1 & -U_1^\top S U_2 & \mathbf{0} \\ -U_2^\top S U_1 & \mu I - U_2^\top S U_2 & -\Sigma_q \end{bmatrix} \right| \\ &= \left| \det \begin{bmatrix} \Sigma_q & \mathbf{0} & \mathbf{0} \\ -U_1^\top S U_2 & \mu I - U_1^\top S U_1 & \mathbf{0} \\ \mu I - U_2^\top S U_2 & -U_2^\top S U_1 & -\Sigma_q \end{bmatrix} \right|. \end{aligned}$$

Notice that both $\mu I - U_1^\top S U_1$ and Σ_q are all square matrices, hence

$$\det M = \det(\mu I - U_1^\top S U_1) \det(\Sigma_q) \det(-\Sigma_q).$$

Because L_q is of full rank, we know that Σ_q is invertible. Therefore, $\det M = 0$ if and only if $\det(\mu I - U_1^\top S U_1) = 0$.

Combining all of the chain of equivalent equations established, we conclude that

$$\det \begin{bmatrix} \mu I - A & -C_1 & -L_q \\ C_1^\top & -D & \mathbf{0} \\ L_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix} = 0 \quad \text{if and only if} \quad \det(\mu I - U_1^\top S U_1) = 0. \quad (21)$$

Noting that $U_1^\top S U_1$ is the restricted Schur complement, we are done. \square

The following is an immediate corollary of the previous theorem.

Corollary E.2. *Suppose that the restricted Schur complement is invertible. Then equation (17) does not have $\mu = 0$ as a solution. In particular, Φ is invertible.*

Proof. The first part directly follows from the equivalence (21). For the second part, because substituting $\mu = 0$ in (17) does not make the equation hold, we must have $\det(-\Phi) \neq 0$. \square

We also use the following lemma, which tells us that the roots of a polynomial is continuous with respect to its coefficients. For the proof, see the original reference.

Lemma E.3 (Zedek (1965, Theorem 1)). *Given a polynomial $p_n(z) := \sum_{k=0}^n a_k z^k$, $a_n \neq 0$, an integer $m \geq n$ and a number $\epsilon > 0$, there exists a number $\delta > 0$ such that whenever the $m + 1$ complex numbers b_k , $0 \leq k \leq m$, satisfy the inequalities*

$$|b_k - a_k| < \delta \quad \text{for} \quad 0 \leq k \leq n, \quad \text{and} \quad |b_k| < \delta \quad \text{for} \quad n + 1 \leq k \leq m,$$

then the roots β_k , $1 \leq k \leq m$, of the polynomial $q_m(z) := \sum_{k=0}^m b_k z^k$ can be labeled in such a way as to satisfy, with respect to the zeros α_k , $1 \leq k \leq n$, of $p_n(z)$, the inequalities

$$|\beta_k - \alpha_k| < \epsilon \quad \text{for} \quad 1 \leq k \leq n, \quad \text{and} \quad |\beta_k| > \frac{1}{\epsilon} \quad \text{for} \quad n + 1 \leq k \leq m.$$

Proof of Theorem 4.3. The eigenvalues of Φ_τ are the solutions of the equation

$$0 = p_\epsilon(\lambda) := \det \begin{bmatrix} \lambda \mathbf{I} - \epsilon \mathbf{A} & -\epsilon \mathbf{C}_1 & -\epsilon \mathbf{L}_q \\ \mathbf{C}_1^\top & \lambda \mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \lambda \mathbf{I} \end{bmatrix} = \det(\lambda \mathbf{I} - \Phi_\tau). \quad (22)$$

By Lemma E.3, constructing the functions $\lambda_j(\epsilon)$ so that they are continuous is possible. Also by that lemma, letting $\epsilon \rightarrow 0$, we see that the eigenvalues converges to the solutions of the equation

$$p_0(\lambda) = \det \begin{bmatrix} \lambda \mathbf{I} & \mathbf{0} & \mathbf{0} \\ \mathbf{C}_1^\top & \lambda \mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \lambda \mathbf{I} \end{bmatrix} = 0.$$

Hence, the r eigenvalues of \mathbf{H}_τ converges to the r nonzero eigenvalues of $-\mathbf{B}$, and the other $d_1 + q$ eigenvalues converges to zero, as $\epsilon \rightarrow 0$.

To investigate the eigenvalues that converges to zero further, we begin by observing that, whenever $|\lambda|$ is small enough so that $\lambda \mathbf{I} - \mathbf{D}$ is invertible, it holds that

$$\begin{aligned} \det \begin{bmatrix} \lambda \mathbf{I} - \epsilon \mathbf{A} & -\epsilon \mathbf{C}_1 & -\epsilon \mathbf{L}_q \\ \mathbf{C}_1^\top & \lambda \mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \lambda \mathbf{I} \end{bmatrix} &= \det \begin{bmatrix} \lambda \mathbf{I} - \epsilon \mathbf{A} + \epsilon \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top & \mathbf{0} & -\epsilon \mathbf{L}_q \\ \mathbf{0} & \lambda \mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \lambda \mathbf{I} \end{bmatrix} \\ &= \det(\lambda \mathbf{I} - \mathbf{D}) \det \begin{bmatrix} \lambda \mathbf{I} - \epsilon (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\epsilon \mathbf{L}_q \\ \mathbf{L}_q^\top & \lambda \mathbf{I} \end{bmatrix}. \end{aligned}$$

That is, any λ_j that converges to zero as $\epsilon \rightarrow 0$ is a solution of the equation

$$0 = \det \begin{bmatrix} \lambda \mathbf{I} - \epsilon (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\epsilon \mathbf{L}_q \\ \mathbf{L}_q^\top & \lambda \mathbf{I} \end{bmatrix}. \quad (23)$$

Now let us reparametrize (23) by $\lambda = \kappa \sqrt{\epsilon}$ to get

$$\begin{aligned} 0 &= \det \begin{bmatrix} \kappa \sqrt{\epsilon} \mathbf{I} - \epsilon (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\epsilon \mathbf{L}_q \\ \mathbf{L}_q^\top & \kappa \sqrt{\epsilon} \mathbf{I} \end{bmatrix} \\ &= \sqrt{\epsilon}^{d_1} \det \begin{bmatrix} \kappa \mathbf{I} - \sqrt{\epsilon} (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\sqrt{\epsilon} \mathbf{L}_q \\ \mathbf{L}_q^\top & \kappa \sqrt{\epsilon} \mathbf{I} \end{bmatrix} \\ &= \sqrt{\epsilon}^{d_1 + d_2} \det \begin{bmatrix} \kappa \mathbf{I} - \sqrt{\epsilon} (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \kappa \mathbf{I} \end{bmatrix}. \end{aligned}$$

Since $(\lambda \mathbf{I} - \mathbf{D})^{-1}$ converges as $\epsilon \rightarrow 0$, we have that if $\lambda_j \rightarrow 0$ as $\epsilon \rightarrow 0$ then λ_j should be a solution of the equation

$$0 = \det \begin{bmatrix} \kappa \mathbf{I} - \sqrt{\epsilon} (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \kappa \mathbf{I} \end{bmatrix} \quad (24)$$

By Lemma E.3, we see that such eigenvalues divided by $\sqrt{\epsilon}$ converge to the roots of

$$0 = \det \begin{bmatrix} \kappa \mathbf{I} & -\mathbf{L}_q \\ \mathbf{L}_q^\top & \kappa \mathbf{I} \end{bmatrix}. \quad (25)$$

Following the first few statements in the proof of Lemma E.1, we get that \mathbf{L}_q must be of full column rank. This implies that \mathbf{L}_q has exactly q singular values. Thus, the nonzero solutions of (25) are exactly $\pm i \sigma_k$, $k = 1, \dots, q$ where σ_k are the singular values of \mathbf{L}_q , and therefore there are $2q$ instances among λ_j where each $\lambda_j / \sqrt{\epsilon}$ converges to each one of $\pm i \sigma_k$ as $\epsilon \rightarrow 0$.

So far, we have shown that r eigenvalues have magnitude $\Theta(1)$, and $2q$ eigenvalues have magnitude $\Theta(\sqrt{\epsilon})$. Meanwhile, because we assume that the restricted Schur complement is invertible, from

$$\det \Phi_\tau = \epsilon^{d_1} \det \begin{bmatrix} \mathbf{A} & \mathbf{C}_1 & \mathbf{L}_q \\ -\mathbf{C}_1^\top & \mathbf{D} & \mathbf{0} \\ -\mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}$$

and that the determinant on the right hand side is nonzero by Corollary E.2, we know that the product of all λ_j should be of order $\Theta(\epsilon^{d_1})$. From these two observations, one can deduce that the product of the remaining $d_1 - q$ eigenvalues should be of order $\Theta(\epsilon^{d_1 - q})$. We claim that these $d_1 - q$ eigenvalues are all exactly of order $\Theta(\epsilon)$. To this end, let us examine what properties would the eigenvalues of order $O(\epsilon)$ have. Reparametrizing $\lambda = \mu\epsilon$ in (22), we get

$$\begin{aligned} 0 &= \begin{bmatrix} \mu\epsilon\mathbf{I} - \epsilon\mathbf{A} & -\epsilon\mathbf{C}_1 & -\epsilon\mathbf{L}_q \\ \mathbf{C}_1^\top & \mu\epsilon\mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mu\epsilon\mathbf{I} \end{bmatrix} \\ &= \epsilon^{d_1} \det \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} & -\mathbf{C}_1 & -\mathbf{L}_q \\ \mathbf{C}_1^\top & \mu\mathbf{I} - \mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mu\mathbf{I} \end{bmatrix}. \end{aligned}$$

Then, by Lemma E.3, μ converges to a root of the equation

$$0 = \det \begin{bmatrix} \mu\mathbf{I} - \mathbf{A} & \mathbf{C}_1 & -\mathbf{L}_q \\ \mathbf{C}_1^\top & -\mathbf{D} & \mathbf{0} \\ \mathbf{L}_q^\top & \mathbf{0} & \mathbf{0} \end{bmatrix}. \quad (26)$$

Again by Corollary E.2, $\mu = 0$ cannot be a root of (26). This in particular implies that no λ_j can be of order $o(\epsilon)$, or in other words, all eigenvalues of \mathbf{H}_τ are of order $\Omega(\epsilon)$. Therefore, for a product of $d_1 - q$ eigenvalues of $\mathbf{\Phi}_\tau$ to be of order $\Theta(\epsilon^{d_1 - q})$, each of those eigenvalues must be of order exactly $\Theta(\epsilon)$, and the claim follows.

Notice that the discussions previous paragraph further imply the following two facts:

- If λ is an eigenvalue of order $O(\epsilon)$ then λ/ϵ converges to a root of (26) as $\epsilon \rightarrow 0$.
- The right hand side of (26) is a polynomial of degree $d_1 - q$ in μ , whose roots are nonzero.

It is now immediate that the $d_1 - q$ eigenvalues of $\mathbf{\Phi}_\tau$ that are of order $\Theta(\epsilon)$ is of the form $\lambda(\epsilon) = \mu\epsilon + o(\epsilon)$ for a (nonzero) root μ of (26).

The claim that $\lambda_j(\epsilon) \neq 0$ whenever $\epsilon > 0$ for any of the $d_1 + q + r$ eigenvalues that are not constantly zero is a direct consequence of the simple fact $\det(\mathbf{\Phi}_\tau) = \det(\mathbf{\Lambda}_\tau) \det(\mathbf{\Phi}) \neq 0$. \square

E.2 PROOFS ON THE PROPERTIES OF THE HEMICURVATURE

In Section 4.2.1, for $\lambda_j(\epsilon)$ with

$$\begin{aligned} \operatorname{Re}(\lambda_j(\epsilon)) &= \xi_j \epsilon^{\varrho_j} + o(\epsilon^{\varrho_j}), \\ \operatorname{Im}(\lambda_j(\epsilon)) &= \sigma_j \sqrt{\epsilon} + o(\sqrt{\epsilon}), \end{aligned} \quad (27)$$

and $\varrho_j > 1/2$, we defined its *hemicurvature* to be the quantity

$$\iota_j := \lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^{\varrho_j - 1}}{\sigma_j^2}.$$

As briefly mentioned therein, the term is named after the following observation.

Proposition E.4. *Let $\alpha(\epsilon) = (\operatorname{Re} \lambda_j(\epsilon), \operatorname{Im} \lambda_j(\epsilon))$ be a plane curve whose components are given as (27). Let $\kappa(\epsilon)$ be the (signed) curvature of α , then it holds that*

$$\iota_j = -\frac{1}{2} \lim_{\epsilon \rightarrow 0^+} \kappa(\epsilon).$$

Proof. For convenience, let us denote $x(\epsilon) = \operatorname{Re}(\lambda_j(\epsilon))$ and $y(\epsilon) = \operatorname{Im}(\lambda_j(\epsilon))$. Let us also omit the subscript j , as there is no confusion by doing so. We use the well known formula (do Carmo, 2016, p. 27)

$$\kappa = \frac{x'y'' - x''y'}{((x')^2 + (y')^2)^{3/2}}$$

to compute the curvature of a plane curve. To this end, observe that

$$\begin{aligned} x'(\epsilon) &= \xi \rho \epsilon^{\rho-1} + o(\epsilon^{\rho-1}), & x''(\epsilon) &= \xi \rho(\rho-1) \epsilon^{\rho-2} + o(\epsilon^{\rho-2}), \\ y'(\epsilon) &= \frac{1}{2} \sigma \epsilon^{-1/2} + o(\epsilon^{-1/2}), & y''(\epsilon) &= -\frac{1}{4} \sigma \epsilon^{-3/2} + o(\epsilon^{-3/2}). \end{aligned}$$

Substituting, and using the fact that $\rho > 1/2$ hence $2\rho - 2 > -1$, we get

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \kappa(\epsilon) &= \lim_{\epsilon \rightarrow 0^+} \frac{-\frac{1}{4} \xi \rho \sigma \epsilon^{\rho-5/2} - \frac{1}{2} \xi \rho(\rho-1) \sigma \epsilon^{\rho-5/2} + o(\epsilon^{\rho-5/2})}{(\xi^2 \rho^2 \epsilon^{2\rho-2} + \frac{1}{4} \sigma^2 \epsilon^{-1} + o(\epsilon^{-1}))^{3/2}} \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{-2\xi \rho(2\rho-1) \sigma \epsilon^{\rho-1} + o(\epsilon^{\rho-1})}{(4\xi^2 \rho^2 \epsilon^{2\rho-1} + \sigma^2 + o(1))^{3/2}} \\ &= -2\rho(2\rho-1) \lim_{\epsilon \rightarrow 0^+} \frac{\xi \epsilon^{\rho-1}}{\sigma^2}. \end{aligned}$$

If $1/2 < \rho < 1$, then the limit is ∞ , so the constant factor $\rho(2\rho-1)$ does not affect the limit. Similarly, if $\rho > 1$, then the limit is 0, so the constant factor $\rho(2\rho-1)$ does not affect the limit also. Finally, if $\rho = 1$, then $\rho(2\rho-1) = 1$, so we are done. \square

In Section 4.2.1, we have also mentioned that one noteworthy property of the hemicurvature is that it is related to $\text{Re}(1/\lambda_j)$. A precise statement on this is as follows.

Proposition E.5. *For $\lambda_j(\epsilon)$ as in (27), it holds that*

$$\iota_j = \lim_{\epsilon \rightarrow 0^+} \text{Re} \left(\frac{1}{\lambda_j(\epsilon)} \right). \quad (28)$$

Proof. We use the simple fact that if $x, y \in \mathbb{R}$ then

$$\text{Re} \left(\frac{1}{x+iy} \right) = \text{Re} \left(\frac{x-iy}{x^2+y^2} \right) = \frac{x}{x^2+y^2}. \quad (29)$$

Substituting (27) into the above leads to

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \text{Re} \left(\frac{1}{\lambda_j(\epsilon)} \right) &= \lim_{\epsilon \rightarrow 0^+} \frac{\text{Re} \lambda_j(\epsilon)}{(\text{Re} \lambda_j(\epsilon))^2 + (\text{Im} \lambda_j(\epsilon))^2} \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^{\rho_j} + o(\epsilon^{\rho_j})}{\sigma_j^2 \epsilon + o(\epsilon)} \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^{\rho_j-1}}{\sigma_j^2} = \iota_j \end{aligned}$$

where in the second line we use the fact that $\rho > 1/2$ implies $\epsilon^{2\rho} = o(\epsilon)$. \square

E.3 ON THE NECESSITY OF CONSIDERING THE HEMICURVATURE

Let us present a depiction on why the tangential information alone may not be sufficient for our stability analyses: see Figure 1. Consider two disks $\bar{D}_{s/2}$ and \bar{D}_s on the complex plane, both tangent to the imaginary axis at 0, but with different radii; $\frac{1}{s}$ and $\frac{1}{2s}$ respectively. The respective boundaries $\partial \bar{D}_{s/2}$ and $\partial \bar{D}_s$ are circles, so assuming without loss of generality that they are positively oriented, it is straightforward that the hemicurvatures of each boundary curves are each $-\frac{s}{2}$ and $-s$, respectively. Now suppose that some $\lambda_j(\epsilon)$ satisfies (27) with $\rho_j = 1$ and $\frac{\xi_j}{\sigma_j^2} = -\frac{3s}{4}$. The shape of the trajectory of such λ_j will be as drawn in the figure, and in particular, this trajectory is also tangent to the imaginary axis at 0. As $\epsilon \rightarrow 0^+$ so that $\lambda_j(\epsilon) \rightarrow 0$, on one hand, we can visually see that the inclusion $\lambda_j(\epsilon) \in \bar{D}_{s/2}$ eventually becomes true. On the other hand, although the fact that $\bar{D}_{s/2}$ and \bar{D}_s share the same tangent line at 0 implies that they are indistinguishable locally at 0 if only the tangential information is used, somewhat interestingly it holds that $\lambda_j(\epsilon) \notin \bar{D}_s$ for all (sufficiently small) $\epsilon > 0$.




Figure 1: Two disks and a curve, all tangent to the same line at the same point.

E.4 PROOF OF THEOREM 4.4

To prove Theorem 4.4 we need the following results.

Proposition E.6. *Under Assumption 1 and $0 < \eta < \frac{1}{L}$, then $\det(D\tilde{\mathbf{w}}_\tau(\mathbf{z})) \neq 0$ for all \mathbf{z} .*

Proof. Assumption 1 implies that any eigenvalue λ of $D\mathbf{F}(\mathbf{z})$ satisfies $|\lambda| \leq L$. So, assuming $0 < \eta < \frac{1}{L}$, any eigenvalue μ of $D\tilde{\mathbf{w}}_\tau(\mathbf{z}) = \mathbf{I} - \eta\mathbf{\Lambda}_\tau D\mathbf{F}(\mathbf{z})$ satisfies $0 < |\mu| < 2$. Hence, we have $\det(D\tilde{\mathbf{w}}_\tau(\mathbf{z})) > 0$. \square

Proposition E.7. *For any $\mathbf{z}^* \in \mathcal{X}_{\text{ns}}^*$, there exists a positive constant $\tau^* > 0$ such that $\mathbf{z}^* \in \mathcal{A}^*(\tilde{\mathbf{w}}_\tau)$ for any $\tau > \tau^*$.*

Proof. Since $D\tilde{\mathbf{w}}_\tau = \mathbf{I} - \eta\mathbf{H}_\tau$, we have

$$\begin{aligned} \mathcal{A}^*(\tilde{\mathbf{w}}_\tau) &= \{\mathbf{z}^* : \mathbf{z}^* = \tilde{\mathbf{w}}_\tau(\mathbf{z}^*), \rho(D\tilde{\mathbf{w}}_\tau(\mathbf{z}^*)) > 1\} \\ &= \{\mathbf{z}^* : \mathbf{z}^* = \tilde{\mathbf{w}}_\tau(\mathbf{z}^*), \exists \lambda \in \text{spec}(\mathbf{H}_\tau(\mathbf{z}^*)) \text{ s.t. } |1 - \eta\lambda| > 1\}. \end{aligned}$$

Observing that $\{z \in \mathbb{C} : |1 - \eta z| > 1\}$ can be equivalently formulated as $\{z \in \mathbb{C} : \text{Re}(\frac{1}{z}) < \frac{\eta}{2}\}$, there exists some $\tau^* > 0$ such that $\mathbf{z}^* \in \mathcal{A}^*(\tilde{\mathbf{w}}_\tau)$ for any $\tau > \tau^*$, since $\mathbf{z}^* \in \mathcal{X}_{\text{ns}}^*$ satisfies $\lim_{\epsilon \rightarrow 0^+} \text{Re}(1/\lambda_j) = \nu_j < \frac{\eta}{2}$ for some $j \in \mathcal{I}$. \square

Proof of Theorem 4.4. Notice that $\mathbf{z}^* \in \mathcal{A}^*(\tilde{\mathbf{w}}_\tau)$ implies

$$\{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(\mathbf{z}_0) = \mathbf{z}^*\} \subset \{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(\mathbf{z}_0) \in \mathcal{A}^*(\tilde{\mathbf{w}}_\tau)\}.$$

Therefore, it follows from Theorem 2.2 that there exists a positive constant $\tau^* > 0$ such that $\mu(\{\mathbf{z}_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(\mathbf{z}_0) = \mathbf{z}^*\}) = 0$ for any $\tau > \tau^*$. \square

F PROOF FOR SECTION 5.1

Proposition F.1. *Under Assumption 1, a point \mathbf{z}^* is an equilibrium point of (6) for $0 < s < 1/L$, or respectively of (7) for $0 < \eta < 1/L$, if and only if $\mathbf{F}(\mathbf{z}^*) = \mathbf{0}$.*

Proof. The ‘‘if’’ part is straightforward from both equations (6) and (7). To show the reverse direction, suppose that \mathbf{z}^* is an equilibrium point. For the case (6), we must have

$$(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(\mathbf{z}^*))^{-1} \mathbf{\Lambda}_\tau \mathbf{F}(\mathbf{z}^*) = \mathbf{0}.$$

Under Assumption 1 and $0 < s < \frac{1}{L}$, $(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(z^*))^{-1}$ is invertible, this is equivalent to $\mathbf{F}(z^*) = \mathbf{0}$. For the case (7), we must have

$$\mathbf{F}(z^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^*)) = \mathbf{0}.$$

For the sake of contradiction, suppose that $\mathbf{F}(z^*) \neq \mathbf{0}$, and let $\mathbf{w}^* := z^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^*)$. Note that $\mathbf{w}^* \neq z^*$ by assumption. Then,

$$\begin{aligned} z^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^*) &= z^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^*)) - \eta\mathbf{\Lambda}_\tau\mathbf{F}(z^*) \\ &= \mathbf{w}^* - \eta\mathbf{\Lambda}_\tau\mathbf{F}(\mathbf{w}^*), \end{aligned}$$

hence we have $z^* - \mathbf{w}^* = \eta\mathbf{\Lambda}_\tau(\mathbf{F}(z^*) - \mathbf{F}(\mathbf{w}^*))$. Under Assumption 1, which is equivalent to the L -Lipschitz continuity of \mathbf{F} , and $0 < \eta < \frac{1}{L}$, we have

$$\begin{aligned} \|z^* - \mathbf{w}^*\| &= \eta \|\mathbf{\Lambda}_\tau(\mathbf{F}(z^*) - \mathbf{F}(\mathbf{w}^*))\| \\ &\leq \eta L \|\mathbf{\Lambda}_\tau\| \|z^* - \mathbf{w}^*\| \\ &\leq \eta L \|z^* - \mathbf{w}^*\| \\ &< \|z^* - \mathbf{w}^*\| \end{aligned}$$

which is absurd. Therefore, we conclude that $\mathbf{F}(z^*) = \mathbf{0}$. \square

G PROOFS AND MISSING DETAILS FOR SECTION 5.2

G.1 PROOF OF PROPOSITION 5.1

The following lemmata will be used in the proof.

Lemma G.1. *At an equilibrium point z^* of ODE (6), it holds that*

$$\mathbf{J}_\tau(z^*) = -(\mathbf{I} + s\mathbf{H}_\tau^*)^{-1}\mathbf{H}_\tau^*.$$

Proof. Differentiating the right hand side of (6) with respect to z_i , one gets that the i^{th} column of the Jacobian matrix \mathbf{J}_τ is equal to

$$\begin{aligned} [\mathbf{J}_\tau]_{:,i} &= \frac{d}{dz_i} \left(-(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(z))^{-1} \mathbf{\Lambda}_\tau \mathbf{F}(z) \right) \\ &= -(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(z))^{-1} \left(\mathbf{\Lambda}_\tau D_i \mathbf{F}(z) - \left(s\mathbf{\Lambda}_\tau \frac{d}{dz_i} D\mathbf{F}(z) \right) (\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(z))^{-1} \mathbf{\Lambda}_\tau \mathbf{F}(z) \right) \end{aligned}$$

where $D_i \mathbf{F}$ denotes the i^{th} partial derivative of \mathbf{F} . Recall that $\mathbf{F}(z^*) = \mathbf{0}$ by Proposition F.1. So, when we evaluate the expression above at z^* , the last term vanishes, giving us what corresponds to the i^{th} column of the matrix identity

$$\mathbf{J}_\tau(z^*) = -(\mathbf{I} + s\mathbf{\Lambda}_\tau D\mathbf{F}(z^*))^{-1} \mathbf{\Lambda}_\tau D\mathbf{F}(z^*).$$

Note that this is exactly the claimed. \square

Lemma G.2. *A complex number $\lambda \in \mathbb{C} \setminus \{-\frac{1}{s}\}$ an eigenvalue of \mathbf{H}_τ^* if and only if $\mu = -\frac{\lambda}{1+s\lambda}$ is an eigenvalue of \mathbf{J}_τ^* .*

Proof. The ‘‘only if’’ part is a direct consequence of applying Theorem 1.13 in (Higham, 2008) on the function $\lambda \mapsto -(1+s\lambda)^{-1}\lambda$. For the ‘‘if’’ part, note that $\mu \neq -\frac{1}{s}$ as long as $\lambda \in \mathbb{C}$, hence $\mathbf{I} + s\mathbf{J}_\tau^*$ is always invertible. The assertion then follows from the fact that the map $X \mapsto -(1+sX)^{-1}X$ is an involution. \square

Proof of Proposition 5.1. It is clear that $f : \mathbb{C} \cup \{\infty\} \rightarrow \mathbb{C} \cup \{\infty\}$ defined as $f(\lambda) = -\frac{\lambda}{1+s\lambda}$ is continuous. Also, it is an involution, hence a bijection. Therefore, f is a homeomorphism.

Let us identify the image of the imaginary axis $i\mathbb{R}$ under the map f . Let $t \in \mathbb{R}$ be any real number, then observe that

$$f(it) + \frac{1}{2s} = \frac{-it}{1+ist} + \frac{1}{2s} = \frac{i+st}{2s(i-st)}.$$

As $st \in \mathbb{R}$, it follows that

$$\left| f(it) + \frac{1}{2s} \right| = \frac{1}{2s} \quad \forall t \in \mathbb{R}.$$

Moreover, $\lim_{t \rightarrow \infty} f(it) = -\frac{1}{s}$. This shows that $f(i\mathbb{R} \cup \{\infty\}) = \partial\bar{\mathcal{D}}_s$, where $\partial\bar{\mathcal{D}}_s$ denotes the boundary of $\bar{\mathcal{D}}_s$, a circle of radius $\frac{1}{2s}$ centered at $-\frac{1}{2s}$ in the complex plane.

In $\mathbb{C} \setminus i\mathbb{R}$ we have two connected components \mathbb{C}_-° and \mathbb{C}_+° , and in $f(\mathbb{C} \setminus i\mathbb{R})$ we have two connected components $\bar{\mathcal{D}}_s \setminus \partial\bar{\mathcal{D}}_s$ and $\mathbb{C} \setminus \bar{\mathcal{D}}_s$. To see which is mapped to which, notice that

$$f(1) = -\frac{1}{1+s}$$

and thus $-\frac{1}{s} < f(1) < 0$. It follows that $f(1) \in \bar{\mathcal{D}}_s \setminus \partial\bar{\mathcal{D}}_s$, and therefore,

$$\begin{aligned} f(\mathbb{C}_-^\circ) &= \mathbb{C} \setminus \bar{\mathcal{D}}_s, \\ f(\mathbb{C}_+^\circ) &= \bar{\mathcal{D}}_s \setminus \partial\bar{\mathcal{D}}_s. \end{aligned}$$


In particular, f is a bijective mapping between \mathbb{C}_-° and $\mathbb{C} \setminus \bar{\mathcal{D}}_s$. So, by Lemma G.2, the spectrum of \mathbf{J}_τ^* lies inside the open left half plane if and only if $\text{spec}(\mathbf{H}_\tau^*) \subset \mathbb{C} \setminus \bar{\mathcal{D}}_s$, which is equivalent to the claimed statement. \square

We would like to remark that $\text{spec}(D\mathbf{F}) \cap \bar{\mathcal{D}}_s = \emptyset$, appearing in the statement of Proposition 5.1, is equivalent to the so-called *negative comonotonicity* condition on \mathbf{F} , which is a condition that guarantees the convergence of (discrete) EG, as demonstrated by Gorbunov et al. (2023). While their discussions rely on algebraic manipulations, our work offers a dynamical systems perspective that sheds light on the significance of this condition.

G.2 TARGET SETS OF CONTINUOUS-TIME METHODS

Using Proposition 5.1, we can depict the *target sets* of two-timescale EG. By *target sets* we mean the region in the complex plane that the spectrum of $\text{spec}(\mathbf{H}_\tau^*)$ must be contained in, in order to make $\text{spec}(\mathbf{J}_\tau^*) \subset \mathbb{C}_-^\circ$ hold. See Figure 2a for a depiction. To additionally demonstrate how the target set of τ -EG depends on the parameter s , we overlaid two target sets on the same plot.

For continuous-time τ -GDA, as studied by Fiez & Ratliff (2021), the target set is the open right half plane. To contrast this set to the set in the case of τ -EG, we depicted the target set for τ -GDA in Figure 2b. An interesting thing to note here is that the target set of τ -GDA can be seen as the (set-theoretic) limit of the target set of τ -EG when $s \rightarrow 0+$, since the limit of $\bar{\mathcal{D}}_s$ as $s \rightarrow 0+$ is the open left half plane. (Here we consider $\bar{\mathcal{D}}_s$ as the closure of \mathcal{D}_s , where the latter denotes the open disk obtained by taking the interior of $\bar{\mathcal{D}}_s$.)



(a) Target set of continuous-time τ -EG



(b) Target set of continuous-time τ -GDA

Figure 2: A depiction of the target sets of continuous two-timescale methods. Notice that for τ -EG we have overlaid two target sets with different choices of s on the same plot.

G.3 A CONSEQUENCE OF ASSUMPTION 2'

As all of the remaining theorems in Section 5.2 employ Assumption 2', let us begin with studying what do we get by additionally assuming the invertibility of DF .

Proposition G.3. *Suppose that Assumption 2' holds. Then $S_{\text{res}}(DF(z^*))$ is necessarily invertible.*

Proof. Because Assumption 2' implies Assumption 2, either one of $S_{\text{res}}(DF(z^*))$ or $\nabla_{\mathbf{y}\mathbf{y}}^2 f(z^*)$ is nondegenerate. For the sake of contradiction, suppose that $S_{\text{res}}(DF(z^*))$ is singular, which asserts that $B = \nabla_{\mathbf{y}\mathbf{y}}^2 f(z^*)$ is nondegenerate. As DF is invertible, a classical result on Schur complements tells us that $S = A - CB^{-1}C^\top$ is also invertible; see, e.g., (Horn & Johnson, 2012, Section 0.8.5). Meanwhile, if B is invertible then $r = d_2$, so Γ is a matrix with zero columns. Hence, $\mathcal{R}(\Gamma)^\perp = \mathbb{R}^{d_2}$, and so the matrix U that defines the restricted Schur complement through the relation $S_{\text{res}}(DF) = U^\top S U$ becomes a square orthogonal matrix, which is in particular invertible. However, this implies that the restricted Schur complement $S_{\text{res}}(DF(z^*))$ is a product of invertible matrices, which cannot be singular. Therefore, our hypothesis must be false, and $S_{\text{res}}(DF(z^*))$ must be invertible. \square

Thus, assuming the invertibility of DF ensures that all μ_j are nonzero in Theorem 4.3. Moreover, because $H_\tau = \Lambda_\tau DF$ is then also invertible hence cannot have a 0 as its eigenvalue, we must have $d_2 - q - r = 0$, and all eigenvalues of H_τ fall into one of type (i), (ii), (iii) eigenvalues characterized in Theorem 4.3.

G.4 PROOF OF THEOREM 5.2

Lemma G.4. *Let $s_0 := \max_{j \in \mathcal{I}}(-\iota_j)$. If $s > s_0$, then there exists some $\tau^* > 0$ such that whenever $\tau > \tau^*$, all the eigenvalues λ_j with $j \in \mathcal{I}$ lie outside of the disk \bar{D}_s . Conversely, if $s < s_0$, then for any sufficiently large τ , there exists some $j \in \mathcal{I}$ such that $\lambda_j(\epsilon) \in \bar{D}_s$.*

Proof. Fix any $j \in \{1, \dots, d_2 - r\}$. Observing that \bar{D}_s can be equivalently formulated as

$$\bar{D}_s = \left\{ z \in \mathbb{C} : \text{Re} \left(\frac{1}{z} \right) \leq -s \right\} \quad (30)$$

we may alternatively analyze how the real part of $1/\lambda_j$ behaves. By Proposition E.5, it holds that

$$\lim_{\epsilon \rightarrow 0^+} \text{Re} \left(\frac{1}{\lambda_j} \right) = \iota_j. \quad (31)$$

For $j \in \{d_1 + 1, \dots, d_1 + d_2 - r\}$, by simply replacing σ_j by $-\sigma_j$, we also have (31).

As we define $s_0 = \max_{j \in \mathcal{I}}(-\iota_j)$, we equivalently have $-s_0 = \min_{j \in \mathcal{I}} \iota_j$. Hence, if $-s < -s_0$, then for all $j \in \mathcal{I}$, the quantity $\text{Re}(1/\lambda_j(\epsilon))$ will eventually be larger than $-s$ as $\tau \rightarrow \infty$. That is, in view of (30), there exists some $\tau^* > 0$ such that whenever $\tau > \tau^*$, all the eigenvalues λ_j with $j \in \mathcal{I}$ lie outside of the disk \bar{D}_s .

Conversely, if $-s > -s_0$, then because \mathcal{I} is a finite set, there exists some $j \in \mathcal{I}$ such that $\text{Re}(1/\lambda_j(\epsilon)) \rightarrow -s_0$ as $\tau \rightarrow \infty$. In other words, for such j , whenever τ is sufficiently large it holds that $\text{Re}(1/\lambda_j(\epsilon)) < -s$, or equivalently, $\lambda_j(\epsilon) \in \bar{D}_s$, according to (30). \square

Remark G.5. *If $\xi_j \geq 0$ for all $j \in \mathcal{I}$, then we have $s_0 \leq 0$. In that case, since we always assume that $s > 0$, we trivially have $s > s_0$. On the other hand, if $\rho_j < 1$ and $\xi_j < 0$ for some $j \in \mathcal{I}$, then $\iota_j = -\infty$, so as a consequence $s_0 = \infty$. In that case, since s is finite, we trivially have $s < s_0$.*

Proof of Theorem 5.2. We analyze the necessary and sufficient condition for each set of eigenvalues λ_j of H_τ^* categorized in Theorem 4.3 to lie outside \bar{D}_s below, and combining them leads to the claimed statement.

- (i) Consider λ_j with $j \in \mathcal{I}$. Suppose that $s_0 < \frac{1}{L}$. Then, by Lemma G.4, for any step size s such that $s_0 < s < \frac{1}{L}$, it holds that $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$ whenever τ is sufficiently large. Conversely, suppose that $s_0 \geq \frac{1}{L}$. Then, for any s such that $0 < s < \frac{1}{L}$, we have $s < s_0$. In that case, by Lemma G.4, for any sufficiently large timescale $\tau > 0$ there exists some $j \in \mathcal{I}$ such that $\lambda_j(\epsilon) \in \overline{\mathcal{D}}_s$. Therefore, $s_0 < \frac{1}{L}$ if and only if there exists some $0 < s^* < \frac{1}{L}$ such that, for any s satisfying $s^* < s < \frac{1}{L}$, τ being sufficiently large implies $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$ for all $j \in \mathcal{I}$.
- (ii) Consider $\lambda_j = \epsilon\mu_k + o(\epsilon)$ for some k . Note that they converge to zero asymptotically along the real axis. In particular, whenever ϵ is sufficiently small, if $\mu_k > 0$ then $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$, and if $\mu_k < 0$ then $\lambda_j(\epsilon) \in \overline{\mathcal{D}}_s$. Recall that in Theorem 4.3 we have established that μ_k are the eigenvalues of the restricted Schur complement \mathbf{S}_{res} . One can then deduce that $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$ if and only if there exists some τ^* where $\tau > \tau^*$ implies that every λ_j of order $\Theta(\epsilon)$ satisfies $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$.
- (iii) Consider $\lambda_j = \nu_k + o(1)$ for some k . As we take Assumption 1, for any $\tau > 1$ it holds that

$$\|\mathbf{H}_\tau^*\| = \|\mathbf{A}_\tau \mathbf{D}\mathbf{F}(\mathbf{z}^*)\| \leq \|\mathbf{A}_\tau\| \|\mathbf{D}\mathbf{F}(\mathbf{z}^*)\| \leq L.$$

As $\lambda_j(\epsilon)$ is an eigenvalue of \mathbf{H}_τ^* , it follows that $|\lambda_j| \leq L$. Similarly, for \mathbf{I}_{d_2} denoting the $d_2 \times d_2$ identity matrix, observe that

$$\|\mathbf{B}\| = \left\| \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{B} \end{bmatrix} \right\| = \left\| \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_2} \end{bmatrix} \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ -\mathbf{C}^\top & \mathbf{B} \end{bmatrix} \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_{d_2} \end{bmatrix} \right\| \leq \|\mathbf{D}\mathbf{F}(\mathbf{z}^*)\| = L.$$

It follows that $|\nu_k| \leq L$ for the eigenvalues ν_k of $-\mathbf{B}$. Now, let $0 < s < \frac{1}{L}$, and suppose that $\mathbf{B} \not\preceq \mathbf{0}$, so that there exists some k such that $-L \leq \nu_k < 0$. Then, since the half-open interval $[-L, 0)$ is contained in the interior of $\overline{\mathcal{D}}_s$, for sufficiently large τ we will have $\lambda_j(\epsilon) \in \overline{\mathcal{D}}_s$. Conversely, if $\mathbf{B} \preceq \mathbf{0}$, then $\nu_k > 0$ for all k , so as a consequence, whenever $s > 0$ and τ is sufficiently large, we will have $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$ for all j such that $\lambda_j(\epsilon) = \nu_k + o(1)$. Therefore, $\mathbf{B} \preceq \mathbf{0}$ if and only if there exists some $0 < s < \frac{1}{L}$ such that τ being sufficiently large implies every λ_j of order $\Theta(1)$ satisfies $\lambda_j(\epsilon) \notin \overline{\mathcal{D}}_s$.

Combining all the discussions made, we conclude that $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $s_0 < \frac{1}{L}$ if and only if there exists some $0 < s^* < \frac{1}{L}$ such that, for any s satisfying $s^* < s < \frac{1}{L}$, τ being sufficiently large implies $\lambda_j(\tau) \notin \overline{\mathcal{D}}_s$ for all j . The conclusion then follows from Proposition 5.1. \square

G.5 PROOF OF THEOREM 5.3

The proof of this theorem utilizes several theorems on the locations of eigenvalues. For completeness, we hereby include their statements.

Theorem G.6 (Weyl). *Let \mathbf{A} , \mathbf{B} be Hermitian matrices, and let the respective eigenvalues of \mathbf{A} , \mathbf{B} , and $\mathbf{A} + \mathbf{B}$ be $\{\lambda_k(\mathbf{A})\}$, $\{\lambda_k(\mathbf{B})\}$, and $\{\lambda_k(\mathbf{A} + \mathbf{B})\}$, each sorted in increasing order. Then for each k , it holds that*

$$\lambda_{k-j+1}(\mathbf{A}) + \lambda_j(\mathbf{B}) \leq \lambda_k(\mathbf{A} + \mathbf{B})$$

for any $j = 1, \dots, k$.

Definition 7. *For any matrix \mathbf{A} , let \mathbf{A}^H denote the conjugate transpose of \mathbf{A} . When \mathbf{A} is a square matrix, the matrix $\frac{1}{2}(\mathbf{A} + \mathbf{A}^H)$ is called the Hermitian part of \mathbf{A} .*

Theorem G.7 (Bendixson). *Let \mathbf{A} be any square matrix, and let \mathbf{R} be its Hermitian part. Then for every eigenvalue λ of \mathbf{A} it holds that*

$$\lambda_{\min}(\mathbf{R}) \leq \text{Re } \lambda \leq \lambda_{\max}(\mathbf{R})$$

where $\lambda_{\min}(\mathbf{R})$ and $\lambda_{\max}(\mathbf{R})$ are the minimum and the maximum eigenvalues of \mathbf{R} , respectively.

Theorem G.8 (Rayleigh). *Let \mathbf{A} be a Hermitian matrix, and denote by $\lambda_{\max}(\mathbf{A})$ the maximum eigenvalue of \mathbf{A} . Then it holds that*

$$\lambda_{\max}(\mathbf{A}) = \max_{\|\mathbf{x}\|=1} \mathbf{x}^H \mathbf{A} \mathbf{x}. \quad (32)$$

Theorems G.6 and G.8 can be found in (Horn & Johnson, 2012) as Theorem 4.3.1 and Theorem 4.2.2, respectively. Theorem G.7 can be found as Theorem 6.9.15 in (Stoer & Bulirsch, 2002). For the proof of these theorems, see the references.

To be more precise, in terms of Rayleigh’s theorem, we need the following corollary rather than the theorem itself.

Corollary G.9. *Let $\mathbf{E} = \text{diag}\{\vartheta_1, \dots, \vartheta_n\}$ be a real diagonal matrix. Choose β to be a nonnegative number such that $-\beta \leq \min\{\vartheta_1, \dots, \vartheta_n\}$. Then for any real matrix $\mathbf{C} \in \mathbb{R}^{m \times n}$, it holds that*

$$\lambda_{\min}(\mathbf{C}\mathbf{E}\mathbf{C}^\top) \geq -\beta \|\mathbf{C}\|^2$$

where $\lambda_{\min}(\mathbf{C}\mathbf{E}\mathbf{C}^\top)$ denotes the minimum eigenvalue of $\mathbf{C}\mathbf{E}\mathbf{C}^\top$.

Proof. Let $\lambda_{\max}(\cdot)$ denote the largest eigenvalue of a given matrix. By how we chose β , we have $-\vartheta_k \leq \beta$ for all k . Hence, for any fixed \mathbf{x} , let $\mathbf{C}^\top \mathbf{x} = \mathbf{y} = [y_1 \dots y_m]^\top$, then

$$\begin{aligned} \mathbf{x}^\top \mathbf{C}(-\mathbf{E})\mathbf{C}^\top \mathbf{x} &= \mathbf{y}^\top (-\mathbf{E})\mathbf{y} = \sum_{k=1}^m -\vartheta_k \mathbf{y}_k^2 \\ &\leq \beta \sum_{k=1}^m \mathbf{y}_k^2 = \beta \|\mathbf{y}\|^2 = \beta \|\mathbf{C}^\top \mathbf{x}\|^2 \leq \beta \|\mathbf{C}\|^2 \|\mathbf{x}\|^2. \end{aligned}$$

Taking the maximum over $\|\mathbf{x}\| = 1$, by Rayleigh’s theorem we get

$$\lambda_{\max}(-\mathbf{C}\mathbf{E}\mathbf{C}^\top) = \max_{\|\mathbf{x}\|=1} \mathbf{x}^\top \mathbf{C}(-\mathbf{E})\mathbf{C}^\top \mathbf{x} \leq \beta \|\mathbf{C}\|^2.$$

As λ is an eigenvalue of $-\mathbf{C}\mathbf{E}\mathbf{C}^\top$ if and only if $-\lambda$ is an eigenvalue of $\mathbf{C}\mathbf{E}\mathbf{C}^\top$, it holds that $\lambda_{\max}(-\mathbf{C}\mathbf{E}\mathbf{C}^\top) = -\lambda_{\min}(\mathbf{C}\mathbf{E}\mathbf{C}^\top)$, so we are done. \square

In order to make the proof clearer, let us state the following simple but technical fact also as a separate lemma.

Lemma G.10. *Let ν be a nonnegative real number. Then, for any $\lambda \in \mathbb{C}_-$, it holds that*

$$-1 \leq \text{Re} \left(\frac{1}{\frac{\nu}{\lambda} - 1} \right) \leq 0.$$

Proof. Let $\lambda = x + iy$ for $x, y \in \mathbb{R}$, and then note that $x \leq 0$. As it holds that

$$\frac{\nu}{\lambda} - 1 = \left(\frac{\nu x}{x^2 + y^2} - 1 \right) - i \frac{\nu y}{x^2 + y^2}$$

we have $\text{Re} \left(\frac{\nu}{\lambda} - 1 \right) \leq -1$, and consequently $\left| \frac{\nu}{\lambda} - 1 \right| \geq 1$. Let $\frac{\nu}{\lambda} - 1 = \xi + iv$ for $\xi, v \in \mathbb{R}$, then as we know that $\xi \leq -1$, we obtain

$$\text{Re} \left(\frac{1}{\frac{\nu}{\lambda} - 1} \right) = \frac{\xi}{\xi^2 + v^2} < 0.$$

Meanwhile, since the size of the real part of a complex number cannot be larger than its absolute value, it also holds that

$$\left| \text{Re} \left(\frac{1}{\frac{\nu}{\lambda} - 1} \right) \right| \leq \left| \frac{1}{\frac{\nu}{\lambda} - 1} \right| \leq 1.$$

Combining the two bounds gives us the result. \square

Proof of Theorem 5.3(i). For any j such that $\lambda_j(\epsilon) = \nu_k + o(1)$, since $\mathbf{B} \preceq \mathbf{0}$ implies $\nu_k > 0$, for any sufficiently large τ we have $\lambda_j(\epsilon) \in \mathbb{C}_+^\circ$. In particular, for any $s > 0$, it holds that $\lambda_j(\epsilon) \notin \bar{\mathcal{D}}_s$ whenever τ is sufficiently large.

For any j such that $\lambda_j(\epsilon) = \epsilon \mu_k + o(\epsilon)$, notice that $\mathbf{S} \succeq \mathbf{0}$ implies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, so we have $\mu_k > 0$. Because $\frac{1}{\epsilon} \lambda_j(\epsilon)$ converges to μ_k , for any sufficiently large τ we have $\frac{1}{\epsilon} \lambda_j(\epsilon) \in \mathbb{C}_+^\circ$, which implies that $\lambda_j(\epsilon) \in \mathbb{C}_+^\circ$. Therefore, for any $s > 0$, it holds that $\lambda_j(\epsilon) \notin \bar{\mathcal{D}}_s$ whenever τ is sufficiently large.

Now for convenience, let $\theta_j(\epsilon) := \frac{1}{\epsilon} \lambda_j(\epsilon) = \tau \lambda_j(\epsilon)$. Our next step is to show that

$$\lim_{\epsilon \rightarrow 0^+} \operatorname{Re}(\theta_j(\epsilon)) \geq 0 \quad \forall j \in \mathcal{I}.$$

If that is the case, for all $j \in \mathcal{I}$ we would have

$$0 \leq \frac{1}{\sigma_j^2} \lim_{\epsilon \rightarrow 0^+} \operatorname{Re}(\theta_j(\epsilon)) = \lim_{\epsilon \rightarrow 0^+} \frac{1}{\sigma_j^2} \operatorname{Re} \left(\frac{1}{\epsilon} \lambda_j(\epsilon) \right) = \lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^\varrho + o(\epsilon^\varrho)}{\epsilon \sigma_j^2} = \iota_j,$$

and consequently $s_0 \leq 0$. Then, for any $s > 0$, we trivially have $s > s_0$. Hence, by applying Lemma G.4 and Proposition 5.1, it is possible to conclude that \mathbf{z}^* is a strict linearly stable point.

Notice that θ_j are the eigenvalues of the matrix

$$\tau \mathbf{H}_\tau^* = \begin{bmatrix} \mathbf{A} & \mathbf{C} \\ -\tau \mathbf{C}^\top & -\tau \mathbf{B} \end{bmatrix}.$$

As in Section 4.2, by applying a similarity transformation, we may assume without loss of generality that \mathbf{B} is a diagonal matrix. Following the notation therein, let $-\mathbf{B} = \operatorname{diag}\{\nu_1, \dots, \nu_r, 0, \dots, 0\}$.

Fix any $j \in \mathcal{I}$. Recall that we have a convergent Puiseux series expansion $\lambda_j(\epsilon) = \sum_{k=1}^{\infty} \alpha_j^{(k)} \epsilon^{k/p_j}$, and notice that it suffices to assume that $\epsilon \in \mathbb{R}$, as we are studying the limit as $\epsilon \rightarrow 0^+$. Then we see that, as $\epsilon \rightarrow 0^+$, either $\operatorname{Re}(\frac{1}{\epsilon} \lambda_j(\epsilon)) = \operatorname{Re}(\theta_j(\epsilon))$ converges to a finite number, diverges to $+\infty$, or diverges to $-\infty$. For the sake of contradiction, suppose that $\lim_{\tau \rightarrow \infty} \operatorname{Re}(\theta_j(\epsilon)) < 0$. By the Schur determinantal formula, the eigenvalue θ_j of $\tau \mathbf{H}_\tau^*$ must be a root of the equation

$$\begin{aligned} \det(\tau \mathbf{H}_\tau^* - \theta \mathbf{I}) &= \det \left(\begin{bmatrix} \mathbf{A} - \theta \mathbf{I} & \mathbf{C} \\ -\tau \mathbf{C}^\top & -\tau \mathbf{B} - \theta \mathbf{I} \end{bmatrix} \right) \\ &= \det(-\tau \mathbf{B} - \theta \mathbf{I}) \det(\mathbf{A} - \theta \mathbf{I} - \tau \mathbf{C}(\tau \mathbf{B} + \theta \mathbf{I})^{-1} \mathbf{C}^\top) \\ &= \det(-\tau \mathbf{B} - \theta \mathbf{I}) \det \left(\mathbf{A} - \theta \mathbf{I} - \mathbf{C} \left(\mathbf{B} + \frac{\theta}{\tau} \mathbf{I} \right)^{-1} \mathbf{C}^\top \right). \end{aligned}$$

Since $\mathbf{B} \preceq 0$, and as for sufficiently large τ we would have $\operatorname{Re}(\theta_j(\epsilon)) < 0$, it must be the case where $\det(-\tau \mathbf{B} - \theta_j \mathbf{I}) \neq 0$. With this in mind, observe that

$$\begin{aligned} &\mathbf{A} - \theta \mathbf{I} - \mathbf{C} \left(\mathbf{B} + \frac{\theta}{\tau} \mathbf{I} \right)^{-1} \mathbf{C}^\top \\ &= \mathbf{A} - \theta \mathbf{I} - \mathbf{C} \left(\operatorname{diag} \left\{ -\nu_1 + \frac{\theta}{\tau}, \dots, -\nu_r + \frac{\theta}{\tau}, \frac{\theta}{\tau}, \dots, \frac{\theta}{\tau} \right\} \right)^{-1} \mathbf{C}^\top \\ &= \mathbf{A} - \theta \mathbf{I} + \mathbf{C} \operatorname{diag} \left\{ \frac{\tau}{\tau \nu_1 - \theta}, \dots, \frac{\tau}{\tau \nu_r - \theta}, -\frac{\tau}{\theta}, \dots, -\frac{\tau}{\theta} \right\} \mathbf{C}^\top \\ &= \mathbf{A} - \mathbf{C} \mathbf{B}^\dagger \mathbf{C}^\top - \theta \mathbf{I} + \mathbf{C} \operatorname{diag} \left\{ \frac{1}{\nu_1} \cdot \frac{1}{\frac{\tau \nu_1}{\theta} - 1}, \dots, \frac{1}{\nu_r} \cdot \frac{1}{\frac{\tau \nu_r}{\theta} - 1}, -\frac{\tau}{\theta}, \dots, -\frac{\tau}{\theta} \right\} \mathbf{C}^\top. \end{aligned}$$

To see why the last line holds, observe that $\frac{\tau}{\tau \nu - \theta} = \frac{1}{\nu} + \frac{1}{\nu(\frac{\tau \nu}{\theta} - 1)}$. Let us denote

$$\mathbf{E} := \operatorname{diag} \left\{ \frac{1}{\nu_1} \cdot \frac{1}{\frac{\tau \nu_1}{\theta} - 1}, \dots, \frac{1}{\nu_r} \cdot \frac{1}{\frac{\tau \nu_r}{\theta} - 1}, -\frac{\tau}{\theta}, \dots, -\frac{\tau}{\theta} \right\}$$

so that $\theta_j(\epsilon)$ becomes a root of the equation

$$\det(\mathbf{A} - \mathbf{C} \mathbf{B}^\dagger \mathbf{C}^\top - \theta \mathbf{I} + \mathbf{C} \mathbf{E} \mathbf{C}^\top) = 0. \quad (33)$$

Now suppose that $\lim_{\tau \rightarrow \infty} \operatorname{Re}(\theta_j(\epsilon)) = -\infty$. For sufficiently large τ , having $\operatorname{Re}(\theta_j(\epsilon)) < 0$ implies $\operatorname{Re}(\theta_j(\epsilon)/\tau) < 0$. So by Lemma G.10,

$$-1 \leq \operatorname{Re} \left(\frac{1}{\frac{\tau \nu_k}{\theta} - 1} \right) \leq 0 \quad (34)$$

for all $k = 1, \dots, r$. Here, note that the Hermitian part of $\mathbf{C}\mathbf{E}\mathbf{C}^\top$ is

$$\begin{aligned} \frac{1}{2} (\mathbf{C}\mathbf{E}\mathbf{C}^\top + (\mathbf{C}\mathbf{E}\mathbf{C}^\top)^\mathbf{H}) &= \mathbf{C} \left(\frac{\mathbf{E} + \mathbf{E}^\mathbf{H}}{2} \right) \mathbf{C}^\top \\ &= \mathbf{C} \operatorname{Re}(\mathbf{E}) \mathbf{C}^\top \end{aligned}$$

where $\operatorname{Re}(\mathbf{E})$ denotes the *elementwise* real part of \mathbf{E} , by the fact that \mathbf{E} is diagonal. Hence, using Corollary G.9 with (34), we get that the Hermitian part of $\mathbf{C}\mathbf{E}\mathbf{C}^\top$ has a spectrum that is bounded below by a constant, uniformly on τ . As we assume that $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top \succeq \mathbf{0}$, by Weyl's theorem, the spectrum of the Hermitian part of $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top + \mathbf{C}\mathbf{E}\mathbf{C}^\top$ is also bounded below by a constant, uniformly on τ . Finally, noticing that the Hermitian part of $\theta\mathbf{I}$ is $\operatorname{Re}(\theta)\mathbf{I}$, and that we assume $\lim_{\tau \rightarrow \infty} \operatorname{Re}(\theta_j(\epsilon)) = -\infty$, Bendixson's theorem tells us that τ being sufficiently large implies every eigenvalue of $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top - \theta_j\mathbf{I} + \mathbf{C}\mathbf{E}\mathbf{C}^\top$ having a strictly positive real part. However, this means that $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top - \theta_j\mathbf{I} + \mathbf{C}\mathbf{E}\mathbf{C}^\top$ is invertible, hence (33) cannot hold.

We are left with the case where $\operatorname{Re}(\theta_j(\epsilon))$ converges to a negative finite value, say $-\psi_j$, as $\tau \rightarrow \infty$. In that case, notice that

$$\lim_{\tau \rightarrow \infty} \frac{1}{\frac{\tau\nu_k}{\theta_j} - 1} = 0$$

for all $k = 1, \dots, r$. As we assume that $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top \succeq \mathbf{0}$, following the same logic used in the previous paragraph, this time we can find τ_0 such that for any $\tau > \tau_0$ the spectrum of the Hermitian part of $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top + \mathbf{C}\mathbf{E}\mathbf{C}^\top$ becomes bounded below by $-\psi_j/2$. Meanwhile, there also exists τ_1 such that $\tau > \tau_1$ implies $-\operatorname{Re}(\theta_j(\epsilon)) > \frac{2}{3}\psi_j$. Therefore, by Bendixson's theorem, whenever $\tau > \max\{\tau_0, \tau_1\}$ the eigenvalues of $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top - \theta_j\mathbf{I} + \mathbf{C}\mathbf{E}\mathbf{C}^\top$ will have a positive real part that is greater than $\psi_j/6$. But then, this again implies that (33) does not hold, which is absurd.

Therefore, for any j it holds that $\lim_{\tau \rightarrow \infty} \operatorname{Re}(\theta_j(\epsilon)) \geq 0$. As claimed in the above, this completes the proof. \square

Proof of Theorem 5.3(ii). Suppose that $\mathbf{S}_{\text{res}} \not\preceq \mathbf{0}$, or equivalently, there exists some k such that $\mu_k < 0$. Then, by Theorem 4.3, we have $\lambda_j(\epsilon) = \epsilon\mu_k + o(\epsilon)$ for some j . In other words, for some j it holds that $\lambda_j(\epsilon)/\epsilon \rightarrow \mu_k$ as $\epsilon \rightarrow 0+$. So by choosing $s = \frac{1}{2|\mu_k|} > 0$ we must have $\lambda_j(\epsilon)/\epsilon \in \overline{\mathcal{D}}_s$ whenever τ is sufficiently large. Since $\epsilon\overline{\mathcal{D}}_s \subset \overline{\mathcal{D}}_s$ whenever $0 < \epsilon \leq 1$, it follows that $\lambda_j(\epsilon) \in \overline{\mathcal{D}}_s$ whenever τ is sufficiently large. However by Proposition 5.1 this implies that \mathbf{z}^* is not strict linearly stable, which is absurd. Therefore, it is necessary that $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$.

Next, suppose that $\mathbf{B} \not\preceq \mathbf{0}$, or equivalently, there exists some k such that $\nu_k < 0$. Then, by Theorem 4.3, as $\lambda_j(\epsilon)$ converges to ν_k for some j , if we choose $s = \frac{1}{2|\nu_k|} > 0$ we must have $\lambda_j(\epsilon) \in \overline{\mathcal{D}}_s$ whenever τ is sufficiently large. However again by Proposition 5.1 this implies that \mathbf{z}^* is not strict linearly stable, which is absurd. Therefore, it is necessary that $\mathbf{B} \preceq \mathbf{0}$. \square

G.6 PROOF OF THEOREM 5.4

Let us first identify what happens if σ_j are all distinct. In doing so, the following proposition, which summarizes the discussion made in (Kato, 1995, Section II.1.2), will be useful. For further details and the proof of the statements, see the original reference.

Proposition G.11. *Consider a convergent perturbation series $\mathbf{T}(\mathcal{X}) = \mathbf{T}_0 + \mathcal{X}\mathbf{T}_1 + \mathcal{X}^2\mathbf{T}_2 + \dots$, and suppose that 0 is in the spectrum of the unperturbed operator \mathbf{T}_0 . Let $\lambda_1(\mathcal{X}), \dots, \lambda_s(\mathcal{X})$ be the eigenvalues of $\mathbf{T}(\mathcal{X})$ with $\lambda_j(0) = 0$ for all $j = 1, \dots, s$. We may assume that such functions are holomorphic on a deleted neighborhood of 0.*

Furthermore, we may assume that those functions can be grouped into $\{\lambda_1(\mathcal{X}), \dots, \lambda_{g_1}(\mathcal{X})\}, \{\lambda_{g_1+1}(\mathcal{X}), \dots, \lambda_{g_1+g_2}(\mathcal{X})\}, \dots$ so that within each group, by letting g to be the size of that group and relabelling the indices of the functions in that group as $\lambda_{j_1}, \dots, \lambda_{j_g}$, the Puiseux series expansion

$$\lambda_{j_h}(\epsilon) = \sum_{\rho=0}^{\infty} a_\rho \left(\omega^h \epsilon^{1/g} \right)^\rho$$

holds for all $h = 1, \dots, g$. Here, ω is a primitive g^{th} root of unity.

The fact that the eigenvalues can be grouped in a way that is described above is crucial, leading to the following observation.

Lemma G.12. *Suppose that all σ_j are distinct. Then, the Puiseux series $\lambda_j(\epsilon) = \sum_{k=1}^{\infty} \alpha_j^{(k)} \epsilon^{k/p_j}$ essentially becomes a power series of $\epsilon^{1/2}$. In particular, $\varrho_j \geq 1$.*

Proof. Fix any λ_j with $j \in \mathcal{I}$. By Proposition G.11, there exists some positive integer g such that we can find $\lambda_{j_1} = \lambda_j, \dots, \lambda_{j_g}$ whose Puiseux series expansions are

$$\lambda_{j_s}(\epsilon) = \sum_{\rho=0}^{\infty} a_{\rho} \left(\omega^s \epsilon^{1/g} \right)^{\rho}$$

for ω a primitive g^{th} root of unity. In particular, the absolute values of the coefficients corresponding to the terms of the same order in each of these g series must be all equal. Since all σ_j values are distinct by assumption, and since they are all real and positive, we conclude that $g = 2$. It is also now clear that the second lowest order term in the Puiseux series expansion is of order $O(\epsilon)$. \square

Now, observe that if $\varrho_j > 1$, then $\iota_j = 0$. Hence, in the case where $\varrho_j > 1$, there is no loss of generality when considering it as having $\varrho_j = 1$ and $\zeta_j = 0$. In other words, by allowing $\zeta_j = 0$, we may assume that the Puiseux expansion of λ_j is of the form

$$\lambda_j = \pm i \sigma_j \sqrt{\epsilon} + \zeta_j \epsilon + o(\epsilon).$$

This suggests a reparametrization, say for example letting $\sqrt{\epsilon} = \delta$, so that the Puiseux series can be dealt as if it is an ordinary power series.

In case where an eigenvalue of a perturbed linear operator (on a finite dimensional space) admits a power series expansion, much more is known on quantitative results regarding the coefficients of that power series. For instance, the following proposition summarizes the discussions made in Sections I.5.3, II.2.2, and II.3.1 in (Kato, 1995). For the details and the proof of the statements, see the original reference.

Proposition G.13. *Consider a convergent perturbation series $\mathbf{T}(\varkappa) = \mathbf{T}_0 + \varkappa \mathbf{T}_1 + \varkappa^2 \mathbf{T}_2 + \dots$. Suppose that λ_0 is a simple eigenvalue of \mathbf{T}_0 , and define*

$$\mathbf{P}_{\lambda_0} := \frac{1}{2\pi i} \int_{\Gamma} (\zeta \mathbf{I} - \mathbf{T}_0)^{-1} d\zeta \quad (35)$$

where Γ is a positively oriented simple closed contour around λ_0 which does not contain any other eigenvalues of \mathbf{T}_0 . Then, there exists an eigenvalue $\lambda(\varkappa)$ of a perturbed operator $\mathbf{T}(\varkappa)$ such that $\lambda(0) = \lambda_0$, which admits a power series expansion

$$\lambda(\varkappa) = \lambda_0 + \varkappa \lambda^{(1)} + \varkappa^2 \lambda^{(2)} + \dots \quad (36)$$

Here, it holds that

$$\lambda^{(1)} = \text{tr}(\mathbf{T}_1 \mathbf{P}_{\lambda_0}).$$

Moreover, let R be the radius of convergence of the power series (36). Then, the fact that the perturbation series is convergent guarantees the existence of nonnegative constants a and c satisfying $\|\mathbf{T}_n\| \leq a c^{n-1}$ for all $n = 1, 2, \dots$. For such a and c it holds that

$$R \geq r_0 := \min_{\zeta \in \Gamma} \frac{1}{a \|\zeta \mathbf{I} - \mathbf{T}_0\| + c}, \quad (37)$$

and there exists a positive constant ρ such that

$$\left| \lambda(\varkappa) - \lambda_0 - \varkappa \lambda^{(1)} \right| \leq \frac{\rho |\varkappa|^2}{r_0 (r_0 - |\varkappa|)}. \quad (38)$$

The operator \mathbf{P}_{λ_0} defined as in (35) is called the *eigenprojection* of \mathbf{T}_0 for λ_0 . If \mathbf{T}_0 is additionally assumed to be normal, we have the following characterization of the eigenprojection for λ_0 .

Lemma G.14. *Suppose that λ_0 is a simple eigenvalue of a normal operator \mathbf{T}_0 . Let \mathbf{w} be the unit norm eigenvector of \mathbf{T}_0 associated with λ_0 . Then, $\mathbf{P}_{\lambda_0} = \mathbf{w}\mathbf{w}^\top$.*

Proof. Because \mathbf{T}_0 is normal, it admits a diagonalization $\mathbf{T}_0 = \mathbf{W}\mathbf{\Theta}\mathbf{W}^\top$, where we may assume without loss of generality that $\mathbf{\Theta} = \text{diag}\{\theta_1, \theta_2, \dots\}$ with $\theta_1 = \lambda_0$ and \mathbf{W} is a unitary matrix with its first column being \mathbf{w} . Then, by definition (35), it holds that

$$\begin{aligned} \mathbf{P}_{\lambda_0} &= \frac{1}{2\pi i} \int_{\Gamma} (\zeta \mathbf{I} - \mathbf{T}_0)^{-1} d\zeta \\ &= \mathbf{W} \left(\frac{1}{2\pi i} \int_{\Gamma} (\zeta \mathbf{I} - \mathbf{\Theta})^{-1} d\zeta \right) \mathbf{W}^\top \\ &= \mathbf{W} \left(\text{diag} \left\{ \frac{1}{2\pi i} \int_{\Gamma} \frac{d\zeta}{\zeta - \theta_1}, \frac{1}{2\pi i} \int_{\Gamma} \frac{d\zeta}{\zeta - \theta_2}, \dots \right\} \right) \mathbf{W}^\top \\ &= \mathbf{W} \text{diag}\{1, 0, \dots, 0\} \mathbf{W}^\top \\ &= \mathbf{w}\mathbf{w}^\top \end{aligned}$$

where in the fourth line we use the fact that Γ encloses $\lambda_0 = \theta_1$ but no other eigenvalues of \mathbf{T}_0 . \square

Proof of Theorem 5.4. Recall that, in the proof of Theorem 4.3, we have shown that any eigenvalue such that $\lambda_j \rightarrow 0$ as $\epsilon \rightarrow 0$ should be a solution of the equation (24), namely the equation

$$0 = \det \begin{bmatrix} \kappa \mathbf{I} - \sqrt{\epsilon} (\mathbf{A} - \mathbf{C}_1 (\lambda \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\mathbf{C}_2 \\ \mathbf{C}_2^\top & \kappa \mathbf{I} \end{bmatrix}.$$

Reparametrizing the above by $\delta = \sqrt{\epsilon}$ gives us

$$0 = \det \begin{bmatrix} \kappa \mathbf{I} - \delta (\mathbf{A} - \mathbf{C}_1 (\kappa \delta \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\mathbf{C}_2 \\ \mathbf{C}_2^\top & \kappa \mathbf{I} \end{bmatrix}.$$

Now fix any index $j \in \{1, 2, \dots, d_2 - r\}$. The cases where $j \in \{d_1 + 1, \dots, d_1 + d_2 - r\}$ can be dealt in the exact same way that is presented hereinafter, by simply changing every instance of σ_j to $-\sigma_j$. We a priori know that $\kappa_j(\delta) := \lambda_j(\delta)/\delta = i\sigma_j + \zeta_j\delta + o(\delta)$. Fix $\delta_0 > 0$ arbitrarily, but small enough so that $\lambda_j(\delta_0) \|\mathbf{D}\| < 1$. Setting $\hat{\kappa} = \lambda_j(\delta_0)/\delta_0$, we have that κ_j is a solution of the equation

$$0 = \det \begin{bmatrix} \kappa \mathbf{I} - \delta_0 (\mathbf{A} - \mathbf{C}_1 (\hat{\kappa} \delta_0 \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top) & -\mathbf{C}_2 \\ \mathbf{C}_2^\top & \kappa \mathbf{I} \end{bmatrix}.$$

In other words, $\kappa_j(\delta_0)$ is an eigenvalue of

$$\begin{aligned} &\begin{bmatrix} \mathbf{0} & \mathbf{C}_2 \\ -\mathbf{C}_2^\top & \mathbf{0} \end{bmatrix} + \delta_0 \begin{bmatrix} \mathbf{A} - \mathbf{C}_1 (\hat{\kappa} \delta_0 \mathbf{I} - \mathbf{D})^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \\ &= \begin{bmatrix} \mathbf{0} & \mathbf{C}_2 \\ -\mathbf{C}_2^\top & \mathbf{0} \end{bmatrix} + \delta_0 \begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} + \sum_{k=2}^{\infty} \delta_0^k \begin{bmatrix} \hat{\kappa}^{k-1} \mathbf{C}_1 \mathbf{D}^{k-2} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}. \end{aligned}$$

Notice that the infinite series expansion above is valid, as δ_0 is chosen to satisfy the condition $\hat{\kappa} \delta_0 \|\mathbf{D}\| = \lambda_j(\delta_0) \|\mathbf{D}\| < 1$. Now, because $\hat{\kappa} \rightarrow i\sigma_j$ as $\delta_0 \rightarrow 0$, by restricting the range of δ_0 further if necessary, we may assume that $\hat{\kappa}$ is uniformly bounded on δ_0 . In particular, we can set the range of δ_0 so that $|\hat{\kappa}| \leq \sigma_j + 1$ uniformly on δ_0 . Then, the coefficients of the perturbation series are bounded by

$$\left\| \begin{bmatrix} \hat{\kappa}^{k-1} \mathbf{C}_1 \mathbf{D}^{k-2} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \right\| \leq \|\mathbf{C}_1\|^2 (\sigma_j + 1)^{k-1} \|\mathbf{D}\|^{k-2}.$$

Notice that these bounds are independent of δ_0 . Invoking Proposition G.13, with restricting the range of δ_0 further if necessary in order to guarantee the convergence of the power series (36), we obtain

$$\kappa_j(\delta_0) = i\sigma_j + \delta_0 \text{tr} \left(\begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{P} \right) + \delta_0^2 \kappa^{(2)} + \dots$$

Moreover, by further restricting the range of δ_0 if necessary so that $\delta_0 \leq r_0/2$, where r_0 is the quantity introduced in (37), we can apply the bound (38) to get a constant C that is independent of δ_0 and satisfies

$$\left| \kappa_j(\delta_0) - i\sigma_j - \delta_0 \operatorname{tr} \left(\begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{P} \right) \right| \leq C\delta_0^2.$$

Note that this estimate holds for all sufficiently small δ_0 . Therefore, we can conclude that

$$\zeta_j = \operatorname{tr} \left(\begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{P} \right).$$

Observe that, for \mathbf{u}_j and \mathbf{v}_j the left and right singular vectors of \mathbf{C}_2 associated with the singular value σ_j respectively, the vector $\begin{bmatrix} \mathbf{u}_j \\ i\mathbf{v}_j \end{bmatrix}$ is the eigenvector of the unperturbed operator $\begin{bmatrix} \mathbf{0} & \mathbf{C}_2 \\ -\mathbf{C}_2^\top & \mathbf{0} \end{bmatrix}$ that is associated to the eigenvalue $i\sigma_j$. Normalizing the eigenvector gives us

$$\hat{\mathbf{w}} = \frac{1}{\sqrt{\|\mathbf{u}_j\|^2 + \|\mathbf{v}_j\|^2}} \begin{bmatrix} \mathbf{u}_j \\ i\mathbf{v}_j \end{bmatrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} \mathbf{u}_j \\ i\mathbf{v}_j \end{bmatrix},$$

and Lemma G.14, we moreover have $\mathbf{P} = \hat{\mathbf{w}}\hat{\mathbf{w}}^\top$. Then, by the cyclic invariance of the trace, we get

$$\begin{aligned} \zeta_j &= \operatorname{tr} \left(\begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \mathbf{P} \right) \\ &= \hat{\mathbf{w}}_j^\top \begin{bmatrix} \mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \hat{\mathbf{w}}_j \\ &= \frac{1}{2} \mathbf{u}_j^\top (\mathbf{A} + \mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top) \mathbf{u}_j. \end{aligned}$$

By the structure of \mathbf{B} , we have $\mathbf{C}_1 \mathbf{D}^{-1} \mathbf{C}_1^\top = -\mathbf{C} \mathbf{B}^\dagger \mathbf{C}^\top$, so we are done. \square

H PROOFS AND MISSING DETAILS FOR SECTION 5.3

H.1 ON THE PEANUT-SHAPED REGION \mathcal{P}_η

The peanut-shaped region \mathcal{P}_η defined in (8) is illustrated in Figure 3. This \mathcal{P}_η is a dilation of \mathcal{P}_1 by a factor of $\frac{1}{\eta}$, and contains a subset of the imaginary axis $\left\{ it : t \in \left(-\frac{1}{\eta}, 0\right) \cup \left(0, \frac{1}{\eta}\right) \right\} \subset i\mathbb{R}$.




Figure 3: The peanut-shaped region \mathcal{P}_η in the complex plane

Similar to the analysis in Appendix E.2, the tangential information alone may not be sufficient for discrete-time case stability analysis: see Figure 4. Consider two peanut-shaped regions $\mathcal{P}_{\eta/16}$ and \mathcal{P}_η on the complex plane. Both are tangent to the imaginary axis at 0, but have different dilation parameter; $\frac{1}{\eta}$ and $\frac{16}{\eta}$ respectively. Suppose that some $\lambda_j(\epsilon)$ satisfies (27) with $\varrho_j = 1$ and $\frac{\xi_j}{\sigma_j^2} = -\frac{\eta}{9}$. Then we can visually see that $\lambda_j(\epsilon) \in \mathcal{P}_\eta$ and $\lambda_j(\epsilon) \notin \mathcal{P}_{\eta/16}$ becomes true as $\epsilon \rightarrow 0+$.

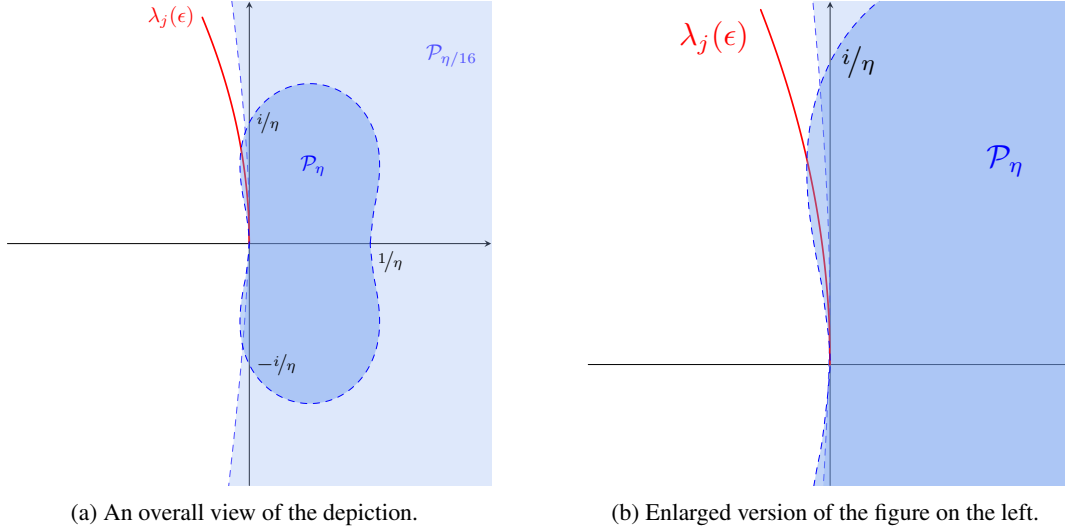



Figure 4: Two peanut-shaped region and a curve, all tangent to the same line at the same point.

Figure 5: Target sets of discrete-time τ -EG and τ -GDA, both using η as their step sizes. The (blue) peanut-shaped region is the target set of τ -EG, and the (orange) disk is the target set of τ -GDA.

H.2 TARGET SETS OF DISCRETE-TIME METHODS

Similar to the continuous-time case, let us compare this to the target set of τ -GDA. As studied by [Fiez & Ratliff \(2021\)](#), the Jacobian of (discrete) τ -GDA is $\mathbf{I} - \eta \mathbf{H}_\tau$, so the target set of τ -GDA becomes the open disk with radius $1/\eta$ centered at $1/\eta$. In [Figure 5](#) we have overlaid the depictions of the target set of τ -EG and the target set of τ -EG, when both methods use η as their step size.

Notice that it is now visually apparent why $\mathbf{0}$ in [Example 2](#) is not a strict linearly stable point for τ -GDA. Clearly, the target set of τ -GDA—the disk—will never contain $\pm i\sqrt{\epsilon}$ for any $\epsilon > 0$.

H.3 PROOF OF PROPOSITION 5.5

For convenience, let us define a subset of the complex plane

$$\mathcal{D}'_{\eta/2} := \left\{ z \in \mathbb{C} : \left| z - \frac{1}{\eta} \right| < \frac{1}{\eta} \right\}$$

which is an open disk centered at $\frac{1}{\eta}$ with radius $\frac{1}{\eta}$.

At an equilibrium point z^* , it holds that

$$\mathbf{J}_\tau(z^*) = \mathbf{I} - \eta \mathbf{H}_\tau^* (\mathbf{I} - \eta \mathbf{H}_\tau^*),$$

and one can easily show that $\rho(\mathbf{J}_\tau^*) < 1$ if and only if $\text{spec}(\mathbf{H}_\tau^*(\mathbf{I} - \eta\mathbf{H}_\tau^*)) \in \mathcal{D}'_{\eta/2}$. The only remaining part to be proven is that $\text{spec}(\mathbf{H}_\tau^*(\mathbf{I} - \eta\mathbf{H}_\tau^*)) \subset \mathcal{D}'_{\eta/2}$ is equivalent to $\text{spec}(\mathbf{H}_\tau^*) \subset \mathcal{P}_\eta$.

Let λ be an eigenvalue of \mathbf{H}_τ^* . Then $\text{spec}(\mathbf{H}_\tau^*(\mathbf{I} - \eta\mathbf{H}_\tau^*)) \subset \mathcal{D}'_{\eta/2}$ implies $\lambda - \eta\lambda^2 \in \mathcal{D}'_{\eta/2}$. By letting $\lambda = x + iy$, where x denotes the real part and y denotes the imaginary part of λ , we have $\lambda - \eta\lambda^2 = x - \eta x^2 + \eta y^2 + i(y - 2\eta xy)$. Then, the condition $\lambda - \eta\lambda^2 \in \mathcal{D}'_{\eta/2}$ can be equivalently written as

$$(\eta x - \eta^2 x^2 + \eta^2 y^2 - 1)^2 + (\eta y - 2\eta^2 xy)^2 < 1.$$

This can be further simplified as

$$\left(\eta x - \frac{1}{2}\right)^2 + \eta^2 y^2 + \frac{3}{4} < \sqrt{1 + 3\eta^2 y^2},$$

and the assertion then follows from the definition of \mathcal{P}_η .

H.4 PROOF OF THEOREM 5.6

In proving Theorem 5.6, the following two results will be used.

Proposition H.1. For $\lambda_j(\epsilon)$ as in (27), it holds that

$$\lim_{\epsilon \rightarrow 0^+} \text{Re} \left(\frac{1}{\lambda_j(1 - \eta\lambda_j)} \right) = \iota_j + \eta. \quad (39)$$

Proof. Using (29) and substituting (27) into the above leads to

$$\begin{aligned} \lim_{\epsilon \rightarrow 0^+} \text{Re} \left(\frac{1}{\lambda_j(1 - \eta\lambda_j)} \right) &= \lim_{\epsilon \rightarrow 0^+} \frac{\text{Re}(\lambda_j(1 - \eta\lambda_j))}{\text{Re}(\lambda_j(1 - \eta\lambda_j))^2 + \text{Im}(\lambda_j(1 - \eta\lambda_j))^2} \\ &= \lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^{2j-1} + o(\epsilon^{2j-1}) + \eta \sigma_j^2 + o(1)}{\sigma_j^2 + o(1)} \\ &= \iota_j + \eta \end{aligned}$$

where the second line uses the fact that $\varrho > 1/2$ implies $\epsilon^{2\varrho} = o(\epsilon)$. \square

Lemma H.2. Let $s_0 := \max_{j \in \mathcal{I}} \iota_j$. If $\frac{\eta}{2} > s_0$, then there exists some $\tau^* > 0$ such that whenever $\tau > \tau^*$, all the eigenvalues λ_j with $j \in \mathcal{I}$ lie in \mathcal{P}_η . Conversely, if $\frac{\eta}{2} < s_0$, then for any sufficiently large τ , there exists some $j \in \mathcal{I}$ such that $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$.

Proof of Lemma H.2. Recall that, in the proof of Proposition 5.5, we have shown that the peanut-shaped region can be represented as follows

$$\mathcal{P}_\eta = \left\{ z \in \mathbb{C} : \text{Re} \left(\frac{1}{z(1 - \eta z)} \right) > \frac{\eta}{2} \right\}.$$

Therefore, if $-s_0 = \min_{j \in \mathcal{I}} \iota_j > -\frac{\eta}{2}$ then for any sufficiently large τ and for all j we will have $\text{Re} \left(\frac{1}{\lambda_j(1 - \eta\lambda_j)} \right) > \frac{\eta}{2}$. Hence, all eigenvalues will lie in \mathcal{P}_η for a sufficiently large τ . On the other hand, if $-s_0 = \min_{j \in \mathcal{I}} \iota_j < -\frac{\eta}{2}$ then for some j we will have $\text{Re} \left(\frac{1}{\lambda_j(1 - \eta\lambda_j)} \right) < \frac{\eta}{2}$ whenever τ is sufficiently large. Hence, there exists some j such that $\lambda_j \notin \mathcal{P}_\eta$ as $\tau \rightarrow \infty$. \square

Proof of Theorem 5.6. Similar to the proof of the Theorem 5.2, we analyze the necessary and sufficient condition for each set of eigenvalues λ_j of \mathbf{H}_τ^* categorized in Theorem 4.3 to lie in \mathcal{P}_η below, and combining them leads to the claimed statement.

- (i) Consider λ_j with $j \in \mathcal{I}$. Suppose that $s_0 < \frac{1}{2L}$. Then, by Lemma H.2, for any step size η such that $s_0 < \frac{\eta}{2} < \frac{1}{2L}$, it holds that $\lambda_j(\epsilon) \in \mathcal{P}_\eta$ whenever τ is sufficiently large. Conversely, suppose that $s_0 \geq \frac{1}{2L}$. Then, for any η such that $0 < \eta < \frac{1}{L}$, we have

$\frac{\eta}{2} < s_0$. In that case, by Lemma H.2, for any sufficiently large timescale $\tau > 0$ there exists some $j \in \mathcal{I}$ such that $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$. Therefore, $s_0 < \frac{1}{2L}$ if and only if there exists some $0 < \eta^* < \frac{1}{L}$ such that for any η satisfying $\eta^* < \eta < \frac{1}{L}$, τ being sufficiently large implies $\lambda_j(\epsilon) \in \mathcal{P}_\eta$ for all $j \in \mathcal{I}$.

- (ii) Consider $\lambda_j = \epsilon\mu_k + o(\epsilon)$ for some k . Note that they converge to zero asymptotically along the real axis. In particular, whenever ϵ is sufficiently small, if $\mu_k > 0$ then $\lambda_j(\epsilon) \in \mathcal{P}_\eta$, and if $\mu_k < 0$ then $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$. Recall that in Theorem 4.3 we have established that μ_k are the eigenvalues of the restricted Schur complement \mathbf{S}_{res} . One can then deduce that $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$ if and only if there exists some τ^* where $\tau > \tau^*$ implies that every λ_j of order $\Theta(\epsilon)$ satisfies $\lambda_j(\epsilon) \in \mathcal{P}_\eta$.
- (iii) Consider $\lambda_j = \nu_k + o(1)$ for some k . By the proof of Theorem 5.2, we have $\|\lambda_j(\epsilon)\| \leq L$ for eigenvalues of \mathbf{H}_τ^* and $\|\nu_k\| \leq L$ for the eigenvalues of $-\mathbf{B}$. Now, let $0 < \eta < \frac{1}{L}$, and suppose that $\mathbf{B} \not\preceq \mathbf{0}$, so that there exists some k such that $-L \leq \nu_k < 0$. Then, since the half-open interval $[-L, 0)$ is contained in the interior of complement of \mathcal{P}_η , for sufficiently large τ we will have $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$. Conversely, if $\mathbf{B} \preceq \mathbf{0}$, then $\nu_k > 0$ for all k , so as a consequence, whenever $\eta > 0$ and τ is sufficiently large, we will have $\lambda_j(\epsilon) \in \mathcal{P}_\eta$ for all j such that $\lambda_j(\epsilon) = \nu_k + o(1)$. Therefore, $\mathbf{B} \preceq \mathbf{0}$ if and only if there exists some $0 < \eta < \frac{1}{L}$ such that τ being sufficiently large implies every λ_j of order $\Theta(1)$ satisfies $\lambda_j(\epsilon) \in \mathcal{P}_\eta$.

Combining all the discussions made, we conclude that $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $s_0 < \frac{1}{2L}$ if and only if there exists some $0 < \eta^* < \frac{1}{L}$ such that, for any η satisfying $\eta^* < \eta < \frac{1}{L}$, τ being sufficiently large implies $\lambda_j(\tau) \in \mathcal{P}_\eta$ for all j . The conclusion then follows from Proposition 5.5. \square

H.5 ADDITIONAL STABILITY CONDITIONS OF TWO-TIMESCALE EG IN DISCRETE-TIME

Theorem H.3. *Let \mathbf{z}^* be an equilibrium point. Suppose Assumptions 1 and 2' hold.*

- (i) *(Sufficient condition) Suppose that \mathbf{z}^* satisfies $\mathbf{S} \succeq \mathbf{0}$ and $\mathbf{B} \preceq \mathbf{0}$. Then for any step size $0 < \eta < 1/L$, the point \mathbf{z}^* is a strict linearly stable point of ∞ -EG.*
- (ii) *(Necessary condition) Suppose that \mathbf{z}^* is a strict linearly stable equilibrium point of ∞ -EG for any step size $0 < \eta < 1/L$. Then, it satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$ and $\mathbf{B} \preceq \mathbf{0}$.*

Theorem H.4. *Under Assumptions 1 and 2', suppose that all the values of σ_j are distinct. Then, an equilibrium point \mathbf{z}^* satisfies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, $\mathbf{B} \preceq \mathbf{0}$, and $\mathbf{u}_j^\top \mathbf{S} \mathbf{u}_j \geq 0$ for every left singular vector \mathbf{u}_j of \mathbf{C}_2 if and only if \mathbf{z}^* is a strict linearly stable point of ∞ -EG for any $0 < \eta < 1/L$.*

H.6 PROOF OF THEOREM H.3

Proof of Theorem H.3(i). For any j such that $\lambda_j(\epsilon) = \nu_k + o(1)$, $\mathbf{B} \preceq \mathbf{0}$ implies that $\nu_k > 0$. Since $|\nu_k| \leq L$ and $0 < \eta < \frac{1}{L}$ implies that \mathcal{P}_η contains a half-open interval $(0, L]$, for any sufficiently large τ we have $\lambda_j(\epsilon) \in \mathcal{P}_\eta$.

And for any j such that $\lambda_j(\epsilon) = \epsilon\mu_k + o(\epsilon)$, since $\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top \succeq \mathbf{0}$ implies $\mathbf{S}_{\text{res}} \succeq \mathbf{0}$, for any sufficiently large τ we have $\lambda_j(\epsilon) \in \mathcal{P}_\eta$.

On the other hand, for any $j \in \mathcal{I}$ proof of Theorem 5.3(i) implies that $\lim_{\epsilon \rightarrow 0^+} \frac{\xi_j \epsilon^{\eta_j - 1}}{\sigma_j^2} \geq 0$ hence $s_0 = \max_{j \in \mathcal{I}}(-\iota_j) \leq 0$. The conclusion then follows from Lemma H.2. \square

To prove Theorem H.3(ii), we use the following technical lemma. For convenience, let us define a subset of the complex plane, for a positive constant $a > 0$,

$$\mathcal{O}_a := \left\{ z \in \mathbb{C} : |z + a| < \frac{a}{2} \right\},$$

which an open disk centered at $-a$ with radius $\frac{a}{2}$.

Lemma H.5. $\mathcal{O}_a \cap \mathcal{P}_\eta = \emptyset$ for any positive constant $a > 0$.

Proof. Noticing that \mathcal{O}_a lies in the left half plane, the only region to care about is \mathbb{C}_- . Thus, if the disk \mathcal{O}_a and the peanut-shaped \mathcal{P}_η do not intersect on that region, then the assertion follows immediately.

Consider a circle centered at origin with radius R , denoted by \mathcal{O}^* . Then, if the circle \mathcal{O}^* and the boundary of \mathcal{O}_a intersects, then it intersects at a point z_1 with a real part $\operatorname{Re} z_1 = -(\frac{R^2}{2a} + \frac{3a}{8})$. Similarly, if the circle \mathcal{O}^* and the boundary of \mathcal{P}_η intersects, then it intersects at a point z_2 with a real part $\operatorname{Re} z_2 = \frac{1}{4\eta} + \frac{\eta R^2}{4} - \sqrt{\frac{1}{16\eta^2} - \frac{3\eta^2 R^4}{16} + \frac{3R^2}{8}}$. For \mathcal{O}_a and \mathcal{P}_η to have an overlap, there must exist some R such that $\operatorname{Re} z_1 = \operatorname{Re} z_2$. We show that such R does not exist for any positive a and η , by proving the following statement

$$\operatorname{Re} z_2 - \operatorname{Re} z_1 = \frac{1}{4\eta} + \frac{\eta R^2}{4} + \frac{R^2}{2a} + \frac{3a}{8} - \sqrt{\frac{1}{16\eta^2} - \frac{3\eta^2 R^4}{16} + \frac{3R^2}{8}} > 0$$

for any positive values a , η and R . This is done by showing that the following

$$\begin{aligned} & \left(\frac{1}{4\eta} + \frac{\eta R^2}{4} + \frac{R^2}{2a} + \frac{3a}{8} \right)^2 - \left(\frac{1}{16\eta^2} - \frac{3\eta^2 R^4}{16} + \frac{3R^2}{8} \right) \\ &= \frac{9a^2}{64} + \frac{3a}{16\eta} + R^2 \left(\frac{3a\eta}{16} + \frac{1}{4a\eta} + \frac{1}{8} \right) + R^4 \left(\frac{1}{4a^2} + \frac{\eta}{4a} + \frac{\eta^2}{4} \right) > 0 \end{aligned}$$

holds for any positive a , η and R , and this concludes the proof. \square

Proof of Theorem H.3(ii). Suppose that $\mathbf{S}_{\text{res}} \not\leq \mathbf{0}$, or equivalently, there exists some k such that $\mu_k < 0$. Then, by Theorem 4.3, we have $\lambda_j(\epsilon) = \epsilon\mu_k + o(\epsilon)$ for some j . So by choosing sufficiently large τ , Lemma H.5 implies that $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$. However by Proposition 5.5 this implies that \mathbf{z}^* is not strict linearly stable, which is absurd. Therefore, it is necessary that $\mathbf{S}_{\text{res}} \leq \mathbf{0}$.

Next, suppose that $\mathbf{B} \not\leq \mathbf{0}$, or equivalently, there exists some k such that $\nu_k < 0$. Then, by Theorem 4.3, as $\lambda_j(\epsilon)$ converges to ν_k for some j , if we choose sufficiently large τ , we must have $\lambda_j(\epsilon) \notin \mathcal{P}_\eta$. However again by Proposition 5.5 this implies that \mathbf{z}^* is not strict linearly stable, which is absurd. Therefore, it is necessary that $\mathbf{B} \leq \mathbf{0}$. \square

H.7 PROOF OF THEOREM H.4

Note that $\zeta_j = \frac{1}{2}\mathbf{U}_j^\top(\mathbf{A} - \mathbf{C}\mathbf{B}^\dagger\mathbf{C}^\top)\mathbf{U}_j$ by the proof of Theorem 5.4. Then the fact that $s_0 \leq 0$ directly follows from the assumption. Lemma H.2 implies that the eigenvalues of the form $\lambda_j(\epsilon) = \pm i\sigma_j\sqrt{\epsilon} + o(\sqrt{\epsilon})$ lie in \mathcal{P}_η for any positive η . And the remaining eigenvalues lie in \mathcal{P}_η if and only if $\mathbf{S}_{\text{res}} \geq \mathbf{0}$ and $\mathbf{B} \leq \mathbf{0}$. The assertion then follows from the Proposition 5.5.

Conversely, suppose that $\mathbf{S}_{\text{res}} \not\geq \mathbf{0}$ or $\mathbf{B} \not\leq \mathbf{0}$. Then instability of the \mathbf{z}^* directly follows from the proof of Theorem H.3(ii).

Next, suppose that there exists some j such that $\zeta_j < 0$. Then the fact that $s_0 > 0$ directly follows from Lemma G.12. Hence for step size $\frac{\eta}{2} < s_0$, the \mathbf{z}^* is not strict linearly stable point by Lemma H.2 and Proposition 5.5.

I PROOF FOR SECTION 5.4

I.1 PROOF OF THEOREM 5.7

To prove Theorem 5.7 we need the following results.

Proposition I.1. *Under Assumption 1 and $0 < \eta < \frac{\sqrt{5}-1}{2L}$, we have $\det(D\mathbf{w}_\tau(\mathbf{z})) \neq 0$ for all \mathbf{z} .*

Proof. We begin by observing that

$$D\mathbf{w}_\tau(\mathbf{z}) = \mathbf{I} - \eta\mathbf{\Lambda}_\tau D\mathbf{F}(\mathbf{z} - \eta\mathbf{\Lambda}_\tau \mathbf{F}(\mathbf{z})) (\mathbf{I} - \eta\mathbf{\Lambda}_\tau D\mathbf{F}(\mathbf{z})).$$

Since $\|\Lambda_\tau\| \leq 1$, Assumption 1 implies that $\|\Lambda_\tau D\mathbf{F}\| \leq \|\Lambda_\tau\| \|D\mathbf{F}\| \leq L$. Therefore, whenever $0 < \eta < \frac{\sqrt{5}-1}{2L}$, for clarity by letting $\tilde{z} = z - \eta\Lambda_\tau\mathbf{F}(z)$, we obtain the bound

$$\begin{aligned} \|\eta\Lambda_\tau D\mathbf{F}(z - \eta\Lambda_\tau\mathbf{F}(z))(I - \eta\Lambda_\tau D\mathbf{F}(z))\| &\leq \|\eta\Lambda_\tau D\mathbf{F}(z - \eta\Lambda_\tau\mathbf{F}(z))\| \|I - \eta\Lambda_\tau D\mathbf{F}(z)\| \\ &\leq \|\eta\Lambda_\tau D\mathbf{F}(\tilde{z})\| (1 + \|\eta\Lambda_\tau D\mathbf{F}(z)\|) \\ &\leq \eta L (1 + \eta L) \\ &< 1. \end{aligned}$$

It follows that any eigenvalue of $\eta\Lambda_\tau D\mathbf{F}(z - \eta\Lambda_\tau\mathbf{F}(z))(I - \eta\Lambda_\tau D\mathbf{F}(z))$ has an absolute value strictly less than 1, and hence, 0 cannot be an eigenvalue of $D\mathbf{w}_\tau(z)$. The conclusion is now immediate. \square

Proposition I.2. *For any $z^* \in \mathcal{T}^*$, under Assumption 2, there exists a positive constant $\tau^* > 0$ such that $z^* \in \mathcal{A}^*(\mathbf{w}_\tau)$ for any $\tau > \tau^*$.*

Proof. By Proposition 5.5, we have

$$\begin{aligned} \mathcal{A}^*(\mathbf{w}_\tau) &= \{z^* : z^* = \mathbf{w}_\tau(z^*), \rho(D\mathbf{w}_\tau(z^*)) > 1\} \\ &= \{z^* : z^* = \mathbf{w}_\tau(z^*), \exists \lambda \in \text{spec}(\mathbf{H}_\tau(z^*)) \text{ s.t. } \lambda \notin \mathcal{P}_\eta\}. \end{aligned}$$

For any strict non-minimax point $z^* \in \mathcal{T}^*$, either $S_{\text{res}}(z^*)$ or $-\mathbf{B}(z^*)$ has at least one strictly negative eigenvalue. First, suppose that $S_{\text{res}}(z^*)$ has a strictly negative eigenvalue $\mu < 0$. By Theorem 4.3, there exists a constant τ^* such that at least one eigenvalue of \mathbf{H}_τ^* lies in a disk $\mathcal{D}_{\mu\epsilon}^\sharp$ for any $\tau > \tau^*$. So by Lemma H.5, we would have $\mathcal{D}_{\mu\epsilon}^\sharp \cap \mathcal{P}_\eta = \emptyset$. On the other hand, suppose that $-\mathbf{B}(z^*)$ has a strictly negative eigenvalue $\nu < 0$. Similarly, by Theorem 4.3, there exists a constant τ^* such that at least one eigenvalue of \mathbf{H}_τ^* lies in a disk \mathcal{D}_ν^\sharp for any $\tau > \tau^*$. So by Lemma H.5, we would have $\mathcal{D}_\nu^\sharp \cap \mathcal{P}_\eta = \emptyset$. Therefore, we can conclude that for any $z^* \in \mathcal{T}^*$, there exists a constant τ^* such that $z^* \in \mathcal{A}^*(\mathbf{w}_\tau)$ for any $\tau > \tau^*$. \square

Proof of Theorem 5.7. Because $z^* \in \mathcal{A}^*(\mathbf{w}_\tau)$ implies

$$\left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) = z^*\right\} \subset \left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) \in \mathcal{A}^*(\mathbf{w}_\tau)\right\},$$

by Theorem 2.2, there exists a positive constant $\tau^* > 0$ such that

$$\mu \left(\left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) = z^*\right\} \right) = 0$$

for any $\tau > \tau^*$.

Moreover, if \mathcal{T}^* is finite, then a maximum of τ^* for all $z^* \in \mathcal{T}^*$ is also finite. Let us denote such maximum by τ_{max}^* . Then, for any $\tau > \tau_{\text{max}}^*$ we have $\mathcal{T}^* \subset \mathcal{A}^*(\mathbf{w}_\tau)$. This implies that

$$\left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) \in \mathcal{T}^*\right\} \subset \left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) \in \mathcal{A}^*(\mathbf{w}_\tau)\right\},$$

for any $\tau > \tau_{\text{max}}^*$, and by Theorem 2.2, we can conclude that

$$\mu \left(\left\{z_0 : \lim_{k \rightarrow \infty} \mathbf{w}^k(z_0) \in \mathcal{T}^*\right\} \right) = 0 \quad \square.$$

J GLOBAL CONVERGENCE ANALYSIS OF TWO-TIMESCALE EG

Let us consider the *Minty variational inequality* (MVI) condition, defined as follows, as a requirement that holds for nonconvex-nonconcave setting.

Assumption 4 (Minty variational inequality). *For the saddle-gradient operator \mathbf{F} , a stationary point $z^* = (x^*, y^*)$ under consideration satisfies $\langle \mathbf{F}(z), z - z^* \rangle \geq 0$ for all z .*

Let us now show that the two-timescale EG

$$\begin{cases} \mathbf{u}_k = \mathbf{x}_k - \frac{\eta}{\tau} \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k) \\ \mathbf{v}_k = \mathbf{y}_k + \eta \nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k) \end{cases}, \quad \begin{cases} \mathbf{x}_{k+1} = \mathbf{x}_k - \frac{\eta}{\tau} \nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k) \\ \mathbf{y}_{k+1} = \mathbf{y}_k + \eta \nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k) \end{cases}.$$

finds a stationary point satisfying the MVI condition, built upon the proof for the plain EG (Diakonikolas et al., 2021). Note that Diakonikolas et al. (2021) consider a variant of EG, named EG+, and shows that it works for a even weaker condition, namely the *weak MVI* condition. The ideas behind our dynamical system analysis and the proof in the section are general enough to be further extended to EG+ and the weak MVI condition, but we leave a more detailed study in this direction as a future work.

The objective function f is said to be L -smooth if the inequality $\|\nabla f(\mathbf{z}) - \nabla f(\mathbf{w})\| \leq L\|\mathbf{z} - \mathbf{w}\|$ holds for any two points \mathbf{z} and \mathbf{w} in the domain of f . If $f \in C^2$, then it is known that f is L -smooth if and only if $\|\nabla^2 f\| \leq L$; see, e.g., (Beck, 2017, Theorem 5.12). As ∇f and \mathbf{F} are equal up to a sign difference, one can observe that Assumption 1 is equivalent to the assertion that f is L -smooth.

Theorem J.1. *Suppose that f is L -smooth, and Assumption 4 holds. Then, the iterates $\{\mathbf{x}_k\}_{k \geq 0}$ and $\{\mathbf{y}_k\}_{k \geq 0}$ of two-timescale EG with $\eta < \frac{1}{L}$ and $\tau \geq 1$ satisfies $\|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\| \rightarrow 0$ and $\|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\| \rightarrow 0$ as $k \rightarrow \infty$.*

Proof. We have the inequality

$$\begin{aligned} & \|\mathbf{x}_{k+1} - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_{k+1} - \mathbf{y}^*\|^2 \\ &= \left\| \mathbf{x}_k - \frac{\eta}{\tau} \nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k) - \mathbf{x}^* \right\|^2 + \frac{1}{\tau} \|\mathbf{y}_k + \eta \nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k) - \mathbf{y}^*\|^2 \\ &= \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 - \frac{2\eta}{\tau} (\langle \nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k), \mathbf{x}_k - \mathbf{x}^* \rangle - \langle \nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k), \mathbf{y}_k - \mathbf{y}^* \rangle) \\ &\quad + \frac{\eta^2}{\tau^2} \|\nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 + \frac{\eta^2}{\tau} \|\nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 \\ &\leq \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 - \frac{2\eta}{\tau} (\langle \nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k), \mathbf{x}_k - \mathbf{u}_k \rangle - \langle \nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k), \mathbf{y}_k - \mathbf{v}_k \rangle) \\ &\quad + \frac{\eta^2}{\tau^2} \|\nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 + \frac{\eta^2}{\tau} \|\nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 \\ &= \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 - \frac{2\eta}{\tau} \left(\left\langle \nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k), \frac{\eta}{\tau} \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k) \right\rangle - \langle \nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k), -\eta \nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k) \rangle \right) \\ &\quad + \frac{\eta^2}{\tau^2} \|\nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 + \frac{\eta^2}{\tau} \|\nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k)\|^2 \\ &= \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 + \frac{\eta^2}{\tau^2} (\|\nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k) - \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 - \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2) \\ &\quad + \frac{\eta^2}{\tau} (\|\nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k) - \nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 - \|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2) \\ &\leq \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 + \frac{\eta^2}{\tau} (\|\nabla_{\mathbf{x}} f(\mathbf{u}_k, \mathbf{v}_k) - \nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 + \|\nabla_{\mathbf{y}} f(\mathbf{u}_k, \mathbf{v}_k) - \nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2) \\ &\quad - \frac{\eta^2}{\tau^2} \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 - \frac{\eta^2}{\tau} \|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 \\ &\leq \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 + \frac{\eta^2 L^2}{\tau} (\|\mathbf{u}_k - \mathbf{x}_k\|^2 + \|\mathbf{v}_k - \mathbf{y}_k\|^2) \\ &\quad - \frac{\eta^2}{\tau^2} \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 - \frac{\eta^2}{\tau} \|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 \\ &\leq \|\mathbf{x}_k - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_k - \mathbf{y}^*\|^2 - \frac{\eta^2}{\tau^2} \left(1 - \frac{\eta^2 L^2}{\tau} \right) \|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 - \frac{\eta^2}{\tau} (1 - \eta^2 L^2) \|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 \end{aligned}$$

where the first inequality uses Assumption 4, the second inequality uses $\tau \geq 1$, the third inequality uses the L -smoothness of f , and the last inequality uses the update rules.

By taking a telescoping summation, we have

$$\begin{aligned} & \frac{\eta^2}{\tau^2} \left(1 - \frac{\eta^2 L^2}{\tau}\right) \sum_{i=0}^k \|\nabla_{\mathbf{x}} f(\mathbf{x}_i, \mathbf{y}_i)\|^2 + \frac{\eta^2}{\tau} (1 - \eta^2 L^2) \sum_{i=0}^k \|\nabla_{\mathbf{y}} f(\mathbf{x}_i, \mathbf{y}_i)\|^2 \\ & \leq \|\mathbf{x}_0 - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_0 - \mathbf{y}^*\|^2 - \|\mathbf{x}_{k+1} - \mathbf{x}^*\|^2 - \frac{1}{\tau} \|\mathbf{y}_{k+1} - \mathbf{y}^*\|^2 \\ & \leq \|\mathbf{x}_0 - \mathbf{x}^*\|^2 + \frac{1}{\tau} \|\mathbf{y}_0 - \mathbf{y}^*\|^2. \end{aligned}$$

Since both the series $\sum_{i=0}^k \|\nabla_{\mathbf{x}} f(\mathbf{x}_i, \mathbf{y}_i)\|^2$ and $\sum_{i=0}^k \|\nabla_{\mathbf{y}} f(\mathbf{x}_i, \mathbf{y}_i)\|^2$ are monotone and bounded, they converge. Therefore we have $\|\nabla_{\mathbf{x}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 \rightarrow 0$ and $\|\nabla_{\mathbf{y}} f(\mathbf{x}_k, \mathbf{y}_k)\|^2 \rightarrow 0$ as $k \rightarrow \infty$. \square

In general, even when a sequence $\{\mathbf{z}_k\}_{k \geq 0}$ satisfies $\lim_{k \rightarrow \infty} \|\mathbf{F}\mathbf{z}_k\| \rightarrow 0$, there is nothing we can say about the convergence of the sequence itself. In other words, the statement of Theorem J.1 alone does not guarantee the convergence of the iterates. However, what the condition $\lim_{k \rightarrow \infty} \|\mathbf{F}\mathbf{z}_k\| \rightarrow 0$ does guarantee, regardless of the MVI condition, is that any accumulation point of the sequence $\{\mathbf{z}_k\}_{k \geq 0}$ becomes a stationary point of the objective function f . Recall that a point \mathbf{z}' is called an accumulation point of a sequence $\{\mathbf{z}_k\}_{k \geq 0}$ if, for any given $\varepsilon > 0$, there exist infinitely many indices $k \geq 0$ such that $\|\mathbf{z}_k - \mathbf{z}'\| < \varepsilon$.

Lemma J.2. *Suppose that f is L -smooth, and let $\{\mathbf{z}_k\}_{k \geq 0}$ be a sequence such that $\|\mathbf{F}(\mathbf{z}_k)\| \rightarrow 0$ as $k \rightarrow \infty$. Then, any accumulation point of the sequence $\{\mathbf{z}_k\}_{k \geq 0}$ is a stationary point of the objective function f .*

Proof. Let \mathbf{z}' be an accumulation point of the sequence $\{\mathbf{z}_k\}_{k \geq 0}$. In that case, there exists a subsequence $\{\mathbf{z}_{k_j}\}_{j \geq 0}$ of $\{\mathbf{z}_k\}_{k \geq 0}$ that converges to \mathbf{z}' (Tao, 2016, Proposition 1.4.5). Then for any $j \geq 0$, by triangle inequality it holds that

$$\begin{aligned} 0 & \leq \|\mathbf{F}(\mathbf{z}')\| \leq \|\mathbf{F}(\mathbf{z}') - \mathbf{F}(\mathbf{z}_{k_j})\| + \|\mathbf{F}(\mathbf{z}_{k_j})\| \\ & \leq L\|\mathbf{z}' - \mathbf{z}_{k_j}\| + \|\mathbf{F}(\mathbf{z}_{k_j})\|. \end{aligned}$$

As both $\|\mathbf{z}' - \mathbf{z}_{k_j}\|$ and $\|\mathbf{F}(\mathbf{z}_{k_j})\|$ converge to 0 as $j \rightarrow \infty$, we must have $\|\mathbf{F}(\mathbf{z}')\| = 0$. \square

Finally, let us demonstrate that all local minimax points of the quadratic function f_q used in Example 4 satisfy the MVI condition.

Lemma J.3. *Consider the quadratic function $f(x_1, x_2, y_1, y_2, y_3) = \frac{1}{2}x_1^2 - \frac{1}{2}y_1^2 - \frac{1}{2}y_3^2 + x_2y_2 + y_1y_3$. Any stationary point of f is a local minimax point that satisfies the MVI condition.*

Proof. We already saw in Section D.3 that the stationary points of f are exactly the points of the form $(0, 0, t, 0, t)$ where $t \in \mathbb{R}$, and that every such point is a local minimax point.

Now fix any $t \in \mathbb{R}$ and let $\mathbf{z}^* = (0, 0, t, 0, t)$. One way to show that this point satisfies the MVI condition is to note that $\mathbf{F}\mathbf{z}^* = \mathbf{0}$, while f is in fact convex-concave hence its saddle gradient \mathbf{F} becomes a so-called *monotone* operator; see (Ryu & Yin, 2022, Section 2.2.3). Alternatively, one can also verify this by a direct computation: for any $\mathbf{z} = (x_1, x_2, y_1, y_2, y_3)$, it holds that

$$\begin{aligned} \langle \mathbf{F}\mathbf{z}, \mathbf{z} - \mathbf{z}^* \rangle & = \langle (x_1, y_2, y_1 - y_3, -x_2, -y_1 + y_3), (x_1, x_2, y_1 - t, y_2, y_3 - t) \rangle \\ & = x_1^2 + y_1^2 + y_3^2 - 2y_1y_3 \\ & = x_1^2 + (y_1 - y_3)^2 \\ & \geq 0 \end{aligned}$$

so the Minty variational inequality holds. \square