The Invisible Hand of Stability

Or, Learnable Mixed Nash Equilibria are Collectively Rational

Geelon So (UC San Diego, CSE) | April 9, 2026

This talk: a surprising connection

Dynamical Stability $\longleftrightarrow$ Collective Rationality

arXiv > cs > arXiv:2510.14907

Learnable Mixed Nash Equilibria are Collectively Rational

Geelon So, Yi-An Ma

We extend the study of learning in games to dynamics that exhibit non-asymptotic stability. We do so through the notion of uniform stability, which is concerned with equilibria of individually utility-seeking dynamics. Perhaps surprisingly, it turns out to be closely connected to economic properties of collective rationality. Under mild non-degeneracy conditions and up to strategic equivalence, if a mixed equilibrium is not uniformly stable, then it is not weakly Pareto optimal: there is a way for all players to improve by jointly deviating from the equilibrium. On the other hand, if it is locally uniformly stable, then the equilibrium must be weakly Pareto optimal. Moreover, we show that uniform stability determines the last-iterate convergence behavior for the family of incremental smoothed best-response dynamics, used to model individual and corporate behaviors in the markets. Unlike dynamics around strict equilibria, which can stabilize to socially-inefficient solutions, individually utility-seeking behaviors near mixed Nash equilibria lead to collective rationality.

Classic problem: learning in games

- General sum, N-player normal-form games
- Infinite repeated setting
- Uncoupled players
- Simple learning rules

Which Nash equilibria can emerge?

Background

An impossibility result on finding Nash equilibria

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Main result: there is a family of games $\mathscr{U}$ such that:

An impossibility result on finding Nash equilibria

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Main result: there is a family of games $\mathscr{U}$ such that:

1. Each game has a unique Nash equilibrium.

An impossibility result on finding Nash equilibria

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Main result: there is a family of games $\mathscr{U}$ such that:

1. Each game has a unique Nash equilibrium.
2. For every uncoupled dynamics, the dynamics will fail to converge to the Nash equilibrium with <u>asymptotic stability</u> for at least one game in this family.

Asymptotic stability

A notion of “strong learnability”

- Dynamics are robust to perturbations
- It will eventually converge to the Nash equilibrium
- Guarantees last-iterate convergence

Significance of this result

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Raises questions on the theoretical viability of the Nash equilibrium

- Is the Nash equilibrium meaningful if players cannot find it?

Some foreshadowing

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Asymptotic stability is too strong of a criterion of learnability

Some foreshadowing

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Asymptotic stability is too strong of a criterion of learnability

- We will introduce a non-asymptotic notion called uniform stability.

Some foreshadowing

Uncoupled dynamics do not lead to Nash equilibrium, Hart and Mas-Colell 2003

Asymptotic stability is too strong of a criterion of learnability

- We will introduce a non-asymptotic notion called <u>uniform stability</u>.
- Perhaps it is also not so bad we cannot find certain Nash equilibria.

Which Nash equilibria can we find?

Convergent dynamics exist for certain important subclasses of equilibria:

Which Nash equilibria can we find?

Convergent dynamics exist for certain important subclasses of equilibria:

- For many standard learning dynamics, there are result of the form:

An equilibrium admits asymptotic convergence
if and only if it is a <u>strict</u> Nash equilibrium.

Which Nash equilibria can we find?

Convergent dynamics exist for certain important subclasses of equilibria:

- For many standard learning dynamics, there are result of the form:

An equilibrium admits asymptotic convergence
if and only if it is a strict Nash equilibrium.

- Strictness: each player has a deterministic strategy that is locally optimal.

Which Nash equilibria can we find?

Convergent dynamics exist for certain important subclasses of equilibria:

- For many standard learning dynamics, there are result of the form:

An equilibrium admits asymptotic convergence
if and only if it is a <u>strict</u> Nash equilibrium.

- Strictness: each player has a deterministic strategy that is locally optimal.
- Fictitious play, replicator dynamics, regularized learning, no-regret learning…

  Samuelson and Zhang (1992); Vlatakis-Gkaragkounis et. al. 2020; Giannou et. al. (2021) …

Which Nash equilibria can we find?

Mixed Nash equilibria do not usually to admit convergence,
but convergent dynamics have been discovered for zero-sum games:

Which Nash equilibria can we find?

Mixed Nash equilibria do not usually to admit convergence,
but convergent dynamics have been discovered for zero-sum games:

- Optimistic gradient ascent
- Extragradient method

Which Nash equilibria can we find?

Mixed Nash equilibria do not usually to admit convergence,
but convergent dynamics have been discovered for zero-sum games:

- Optimistic gradient ascent
- Extragradient method

Further foreshadowing: our work gives some explanation as to why.

Which Nash equilibria are good?

Price of Anarchy: the social cost of decentralized rationality.

Which Nash equilibria are good?

Price of Anarchy: the social cost of decentralized rationality.

- Situations where the inability to cooperate, coordinate, or communicate lead to everyone being worse off.
  - Prisoner’s Dilemma
  - Tragedy of the Commons

Which Nash equilibria are good?

Price of Anarchy: the social cost of decentralized rationality.

- Situations where the inability to cooperate, coordinate, or communicate lead to everyone being worse off.
  - Prisoner’s Dilemma
  - Tragedy of the Commons

This work: <u>weak Pareto optimality</u> as a minimal notion of collective rationality.

Nash equilibrium vs. Pareto optimum

Nash Equilibrium - Individual Rationality

Fixing everyone else, no player chooses to play a suboptimal strategy.

Nash equilibrium vs. Pareto optimum

Nash Equilibrium - Individual Rationality

Fixing everyone else, no player chooses to play a suboptimal strategy.

Weak Pareto Optimum - Collective Rationality

There is no alternative joint strategy that is strictly better for everyone.

Prior Work: Stability of Social Conventions

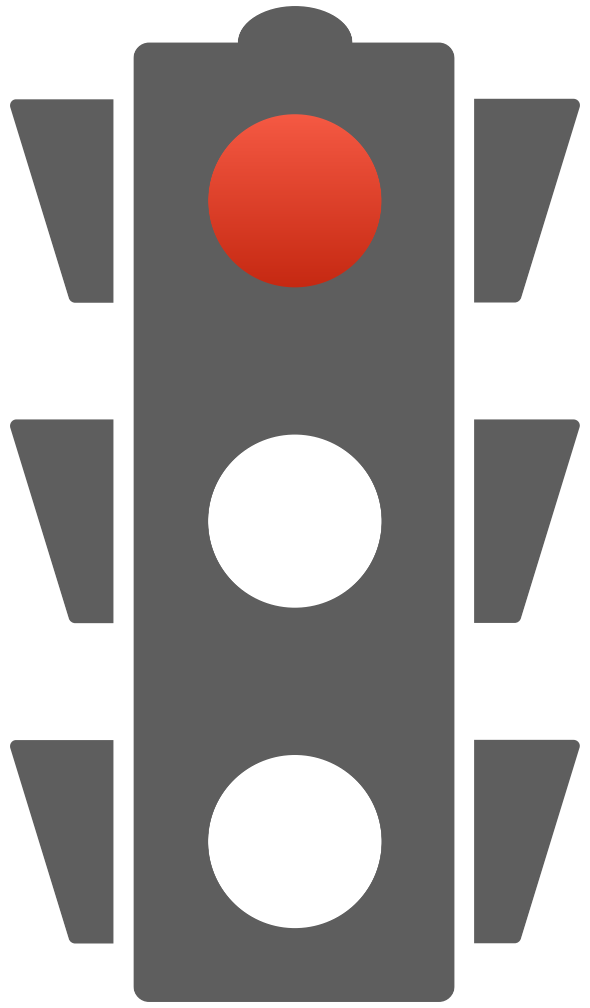

Dynamics of multi-agent systems around strict Nash equilibria are well-understood.

Example. Once enough people agree **RED/GREEN** means **STOP/GO**, everyone else is forced to adopt this convention.

Prior Work: Stability of Social Conventions

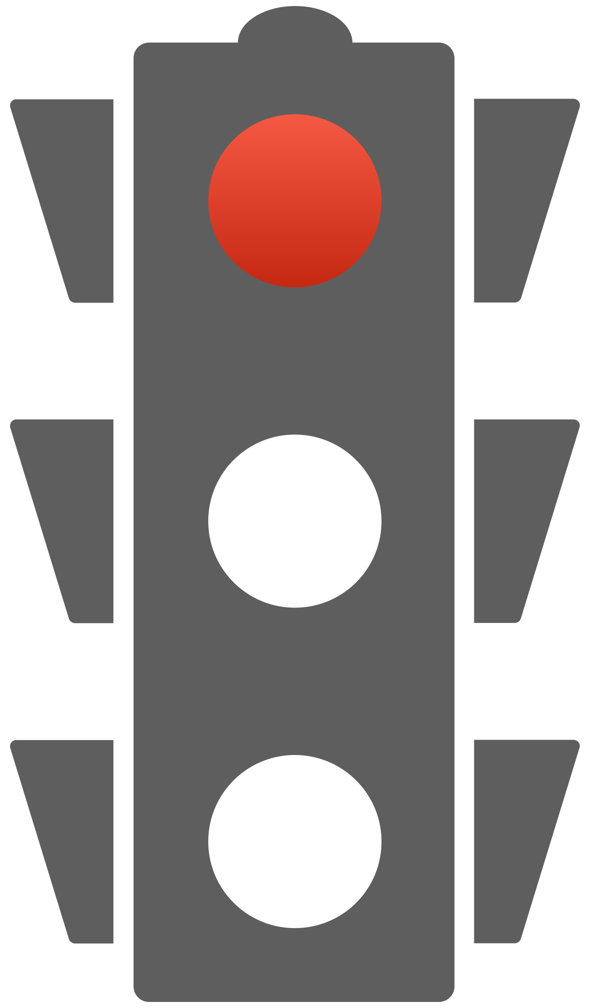

Dynamics of multi-agent systems around strict Nash equilibria are well-understood.

Example. Once enough people agree **RED/GREEN** means **STOP/GO**, everyone else is forced to adopt this <u>convention</u>.

Stability does <u>not</u> guarantee collective rationality.

Stability in Non-Strict Equilibria

What about dynamics around non-strict Nash equilibria?

Example. When the job market is at equilibrium, many career options are viable, leading to a mixed population.

Stability in Non-Strict Equilibria

What about dynamics around **non-strict** Nash equilibria?

Example. When the job market is at equilibrium, many career options are viable, leading to a mixed population.

Stability guarantees collective rationality.

This work

Our results

1. **Connection between stability and collective rationality.**

local uniform stability $\implies$ Pareto optimality * $\implies$ pointwise uniform stability

* Technically, strategic Pareto optimality, which holds up to strategic equivalence.

Our results

1. **Connection between stability and collective rationality.**

   local uniform stability $\Longrightarrow$ Pareto optimality * $\Longrightarrow$ pointwise uniform stability

2. **Dynamics of a specific class of learning rules.**
   - Incremental smooth best-response dynamics
   - Local uniform stability $\Longrightarrow$ all dynamics can be stabilized
   - Not uniformly stable $\Longrightarrow$ certain dynamics can never be stabilized

* Technically, strategic Pareto optimality, which holds up to strategic equivalence.

Comparison of Dynamics Around Non-Strict NE

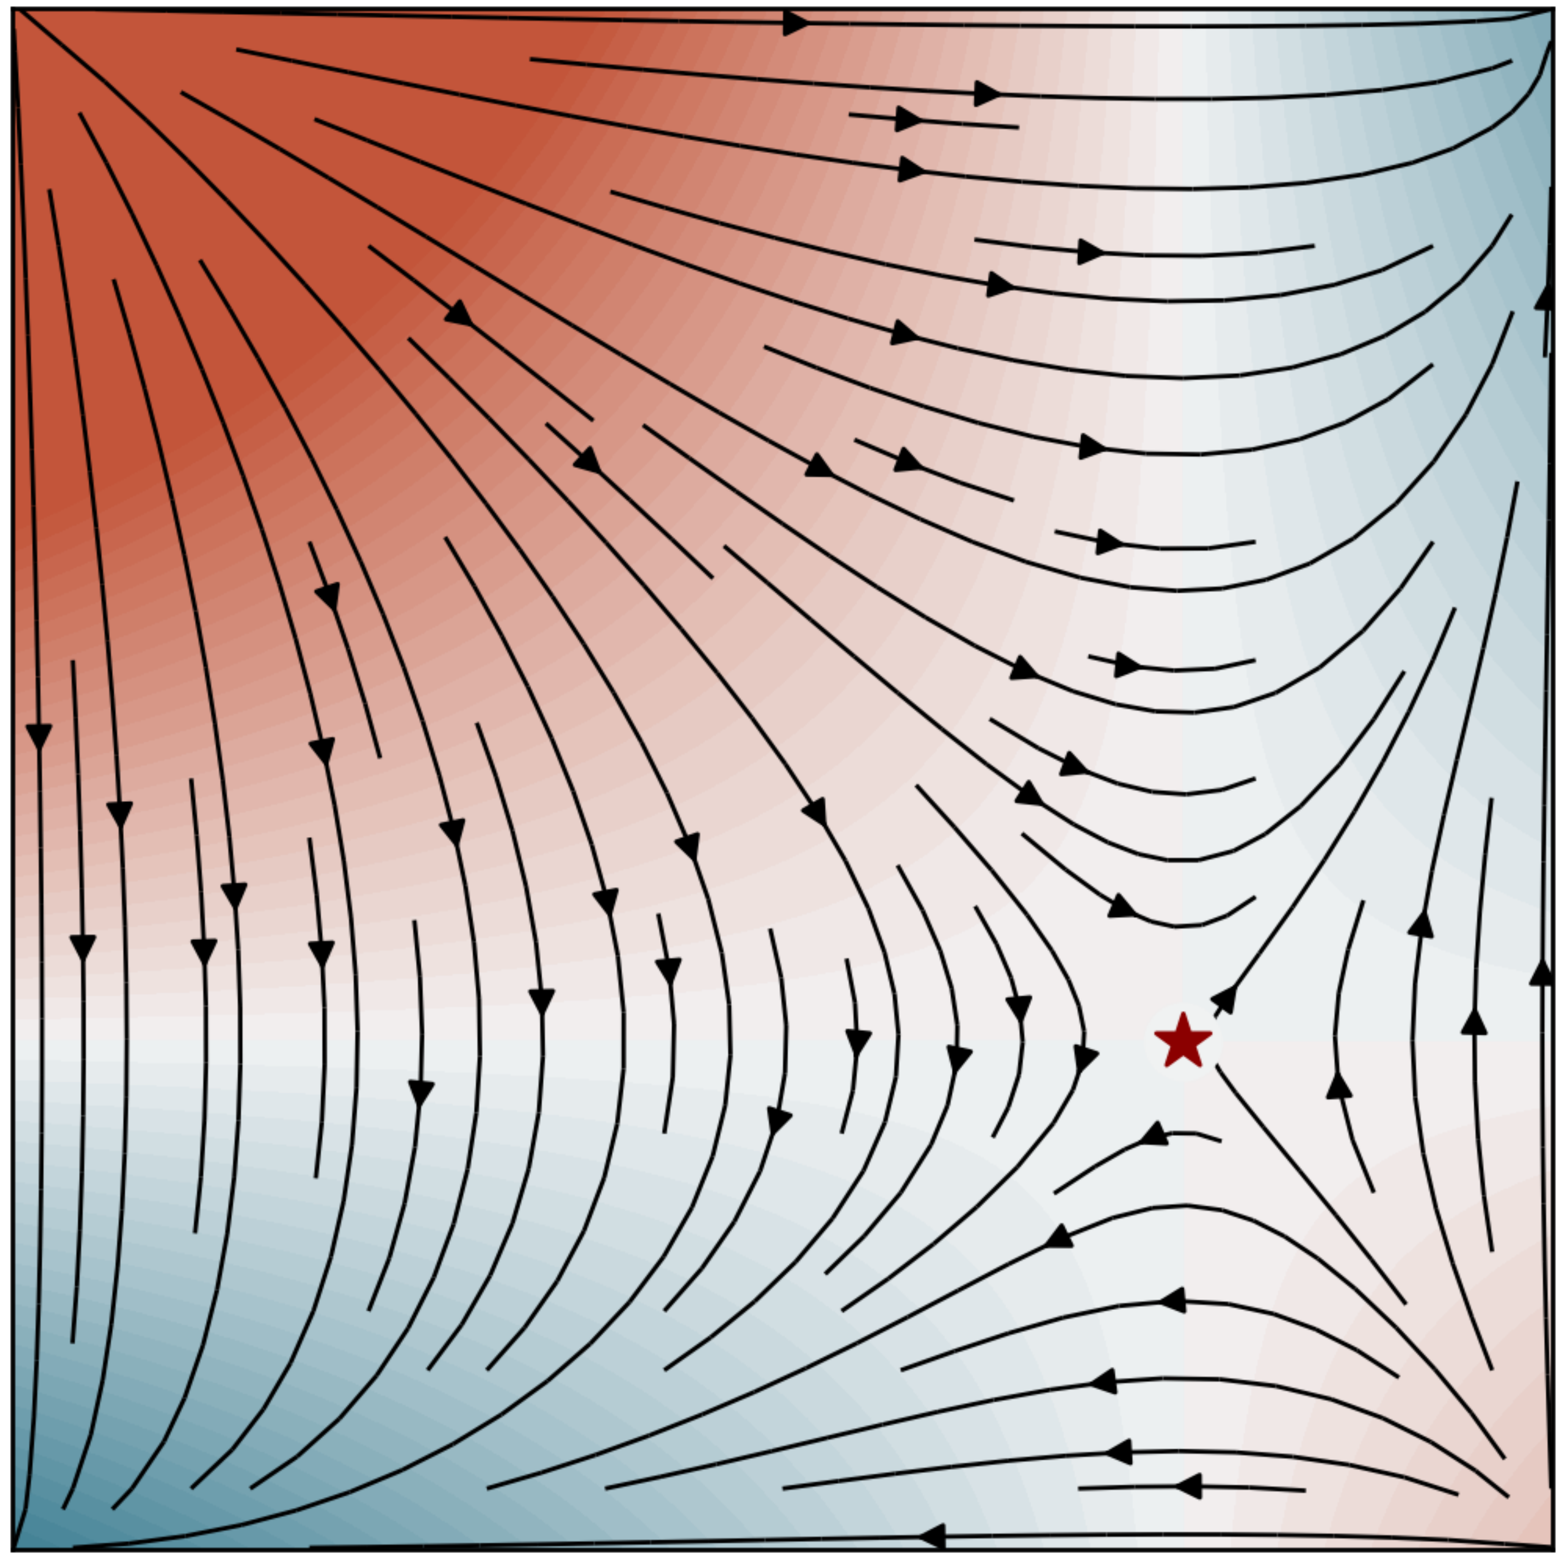

Unstable Nash equilibria are <u>not</u> strategically Pareto optimal.

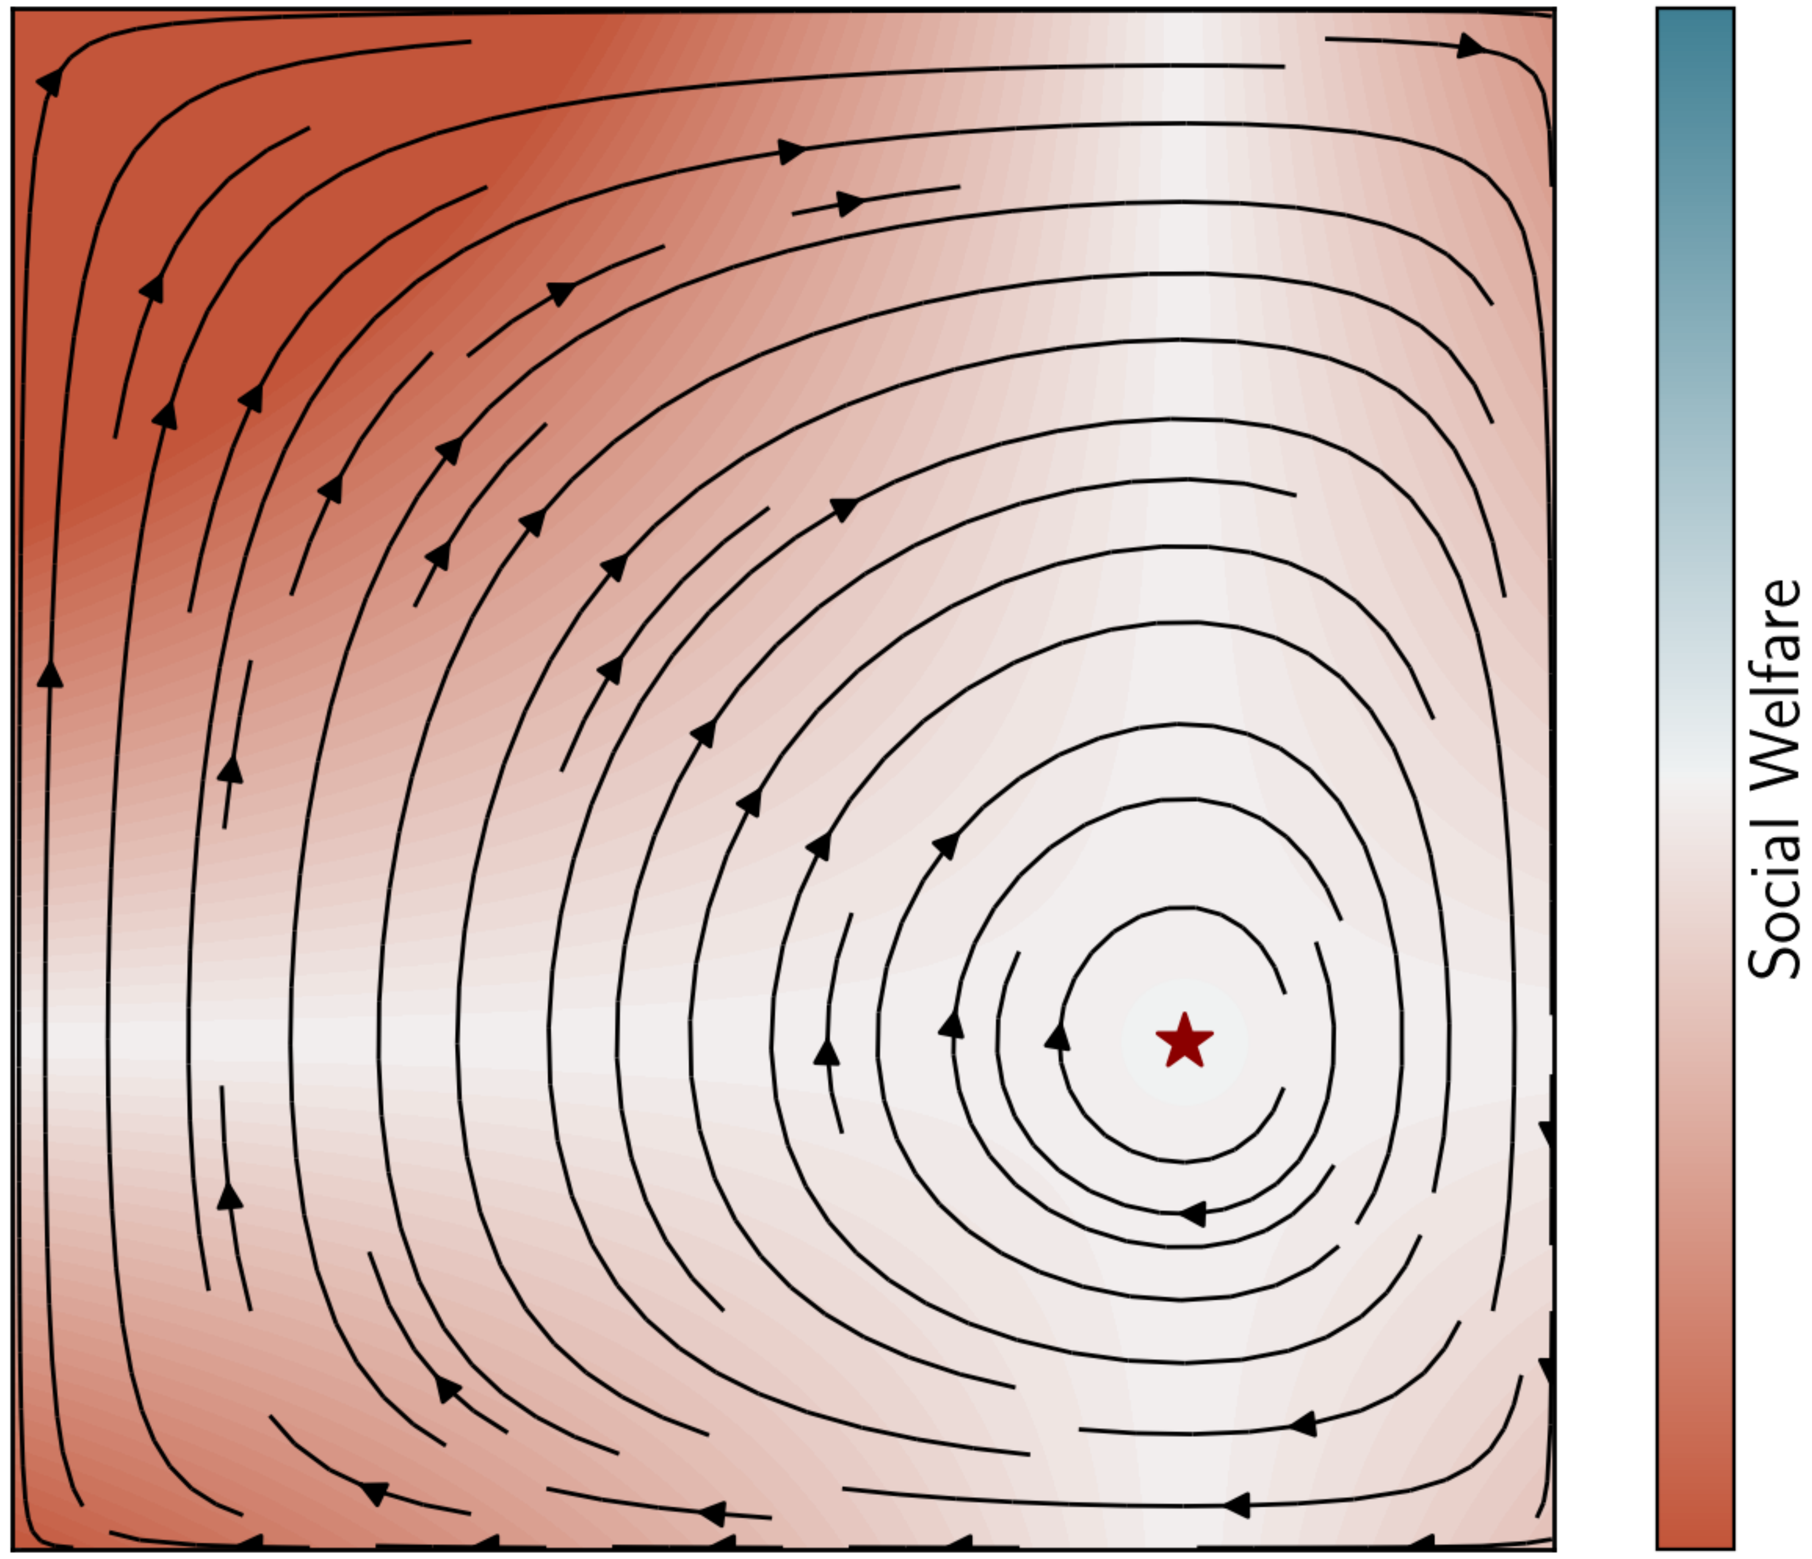

Uniformly stable Nash equilibria are strategically Pareto optimal.

Main takeaway

The mixed Nash equilibria that are robustly learnable are the ones that are strategically Pareto optimal.

- The notion of learnability is a weaker, non-asymptotic criterion.
- Non-convergence is not necessarily a bad thing if the equilibrium is Pareto dominated.

Basic Concepts

Normal-form game

- General sum, N-player normal-form games
- Player $n \in [N]$ has utility

$$u_n(\mathbf{x}) \equiv u_n(x_1, \ldots, x_N) \equiv u_n(x_n, \mathbf{x}_{-n}) \,.$$

  - The strategy space Ω_n is the probability simplex over some k_n actions.
  - The utility u_n is a multilinear polynomial in $(x_1, \ldots, x_N)$:

$$u_n(\mathbf{x}) = \sum_{i_m \in [k_m]} T_{n;i_1,\ldots,i_N} x_{1,i_1} \cdots x_{N,i_N} \,.$$

Normal-form game

- General sum, N-player normal-form games
- Player $n \in [N]$ has utility

$$u_n(\mathbf{x}) \equiv u_n(x_1, \ldots, x_N) \equiv u_n(x_n, \mathbf{x}_{-n}) .$$

  - The strategy space Ω_n is the probability simplex over some k_n actions.
  - In the two-player setting, u_n is a bilinear polynomial (matrix game):

$$u_1(x_1, x_2) = x_1^\top J_{12} x_2 \qquad \text{and} \qquad u_2(x_1, x_2) = x_2^\top J_{21} x_1 .$$

Repeated Games

Players enter a game with limited knowledge:

- They know what set of actions they/others can take.
- They know their own utilities.

They do not know:

- What other players want.
- How other players learn.

The players are **uncoupled**.

Repeated Games

Players enter a game with limited knowledge:

- They know what set of actions they/others can take.
- They know their own utilities.

They do not know:

- What other players want.
- How other players learn.

Players will repeatedly play the game, observe what other do, and adjust their own behaviors for future rounds.

Learning in games

- **Infinitely repeated game:** For $t = 1,2,\ldots$:
  - Each player $n \in [N]$ plays a strategy $x_n(t)$
  - Everyone observes the joint strategy played $\mathbf{x}(t)$

Learning in games

- **Infinitely repeated game:** For $t = 1,2,\ldots$:*
  - Each player $n \in [N]$ plays a strategy $x_n(t)$
  - Everyone observes the joint strategy played $\mathbf{x}(t)$

*The continuous time version is also sensible.

Learning in games

- **Infinitely repeated game:** For $t = 1,2,\ldots$:*
  - Each player $n \in [N]$ plays a strategy $x_n(t)$
  - Everyone observes the joint strategy played $\mathbf{x}(t)$
- **Last-iterate convergence:** as $t \to \infty$, does $\mathbf{x}(t) \to \mathbf{x}^*$ converge to Nash?

*The continuous time version is also sensible.

Simple Learning Rules

- **Best-response dynamics**

$$x_n(t+1) = \arg\max_{z_n \in \Omega_n} u_n\big(z_n; \mathbf{x}_{-n}(t)\big)$$

- **Smoothed best-response dynamics**
- **Gradient ascent**
- **Positive-definite dynamics**

Simple Learning Rules

- **Best-response dynamics**
- **Smoothed best-response dynamics**

$$x_n(t+1) = \arg\max_{z_n \in \Omega_n} u_n\Big(z_n; \mathbf{x}_{-n}(t)\Big) + \text{regularizer}(z_n)$$

- **Gradient ascent**
- **Positive-definite dynamics**

Simple Learning Rules

- **Best-response dynamics**
- **Smoothed best-response dynamics**
- **Gradient ascent-ascent**

$$\dot{x}_n(t) = \eta \nabla_n u_n\Big(x_n(t); \mathbf{x}_{-n}(t)\Big)$$

- **Positive-definite dynamics**

Simple Learning Rules

- **Best-response dynamics**
- **Smoothed best-response dynamics**
- **Gradient ascent**
- **Positive-definite dynamics**

$$\dot{x}(t) = H_n^{-1} \nabla_n u_n\big(\mathbf{x}(t)\big)$$

where H_n is positive-definite (also, preconditioned gradient ascent, mirror ascent,…).

Dynamical Equilibria

The **dynamics** of a (Markov) learning rule is governed by a transition map:

$$T : \mathbf{\Omega} \to \mathbf{\Omega},$$

where $\mathbf{x}(t+1) = T\big(\mathbf{x}(t)\big)$.

Dynamical Equilibria

The **dynamics** of a (Markov) learning rule is governed by a transition map:

$$T : \mathbf{\Omega} \to \mathbf{\Omega},$$

where $\mathbf{x}(t+1) = T\big(\mathbf{x}(t)\big)$.

- A joint strategy $\mathbf{x}^\dagger$ is an **equilibrium** or a **fixed point** if:

$$\mathbf{x}^\dagger = T(\mathbf{x}^\dagger).$$

Classic Notions of Dynamical Stability

1. Asymptotic stability
2. Non-asymptotic/Lyapunov stability
3. Instability

1. Asymptotic Stability

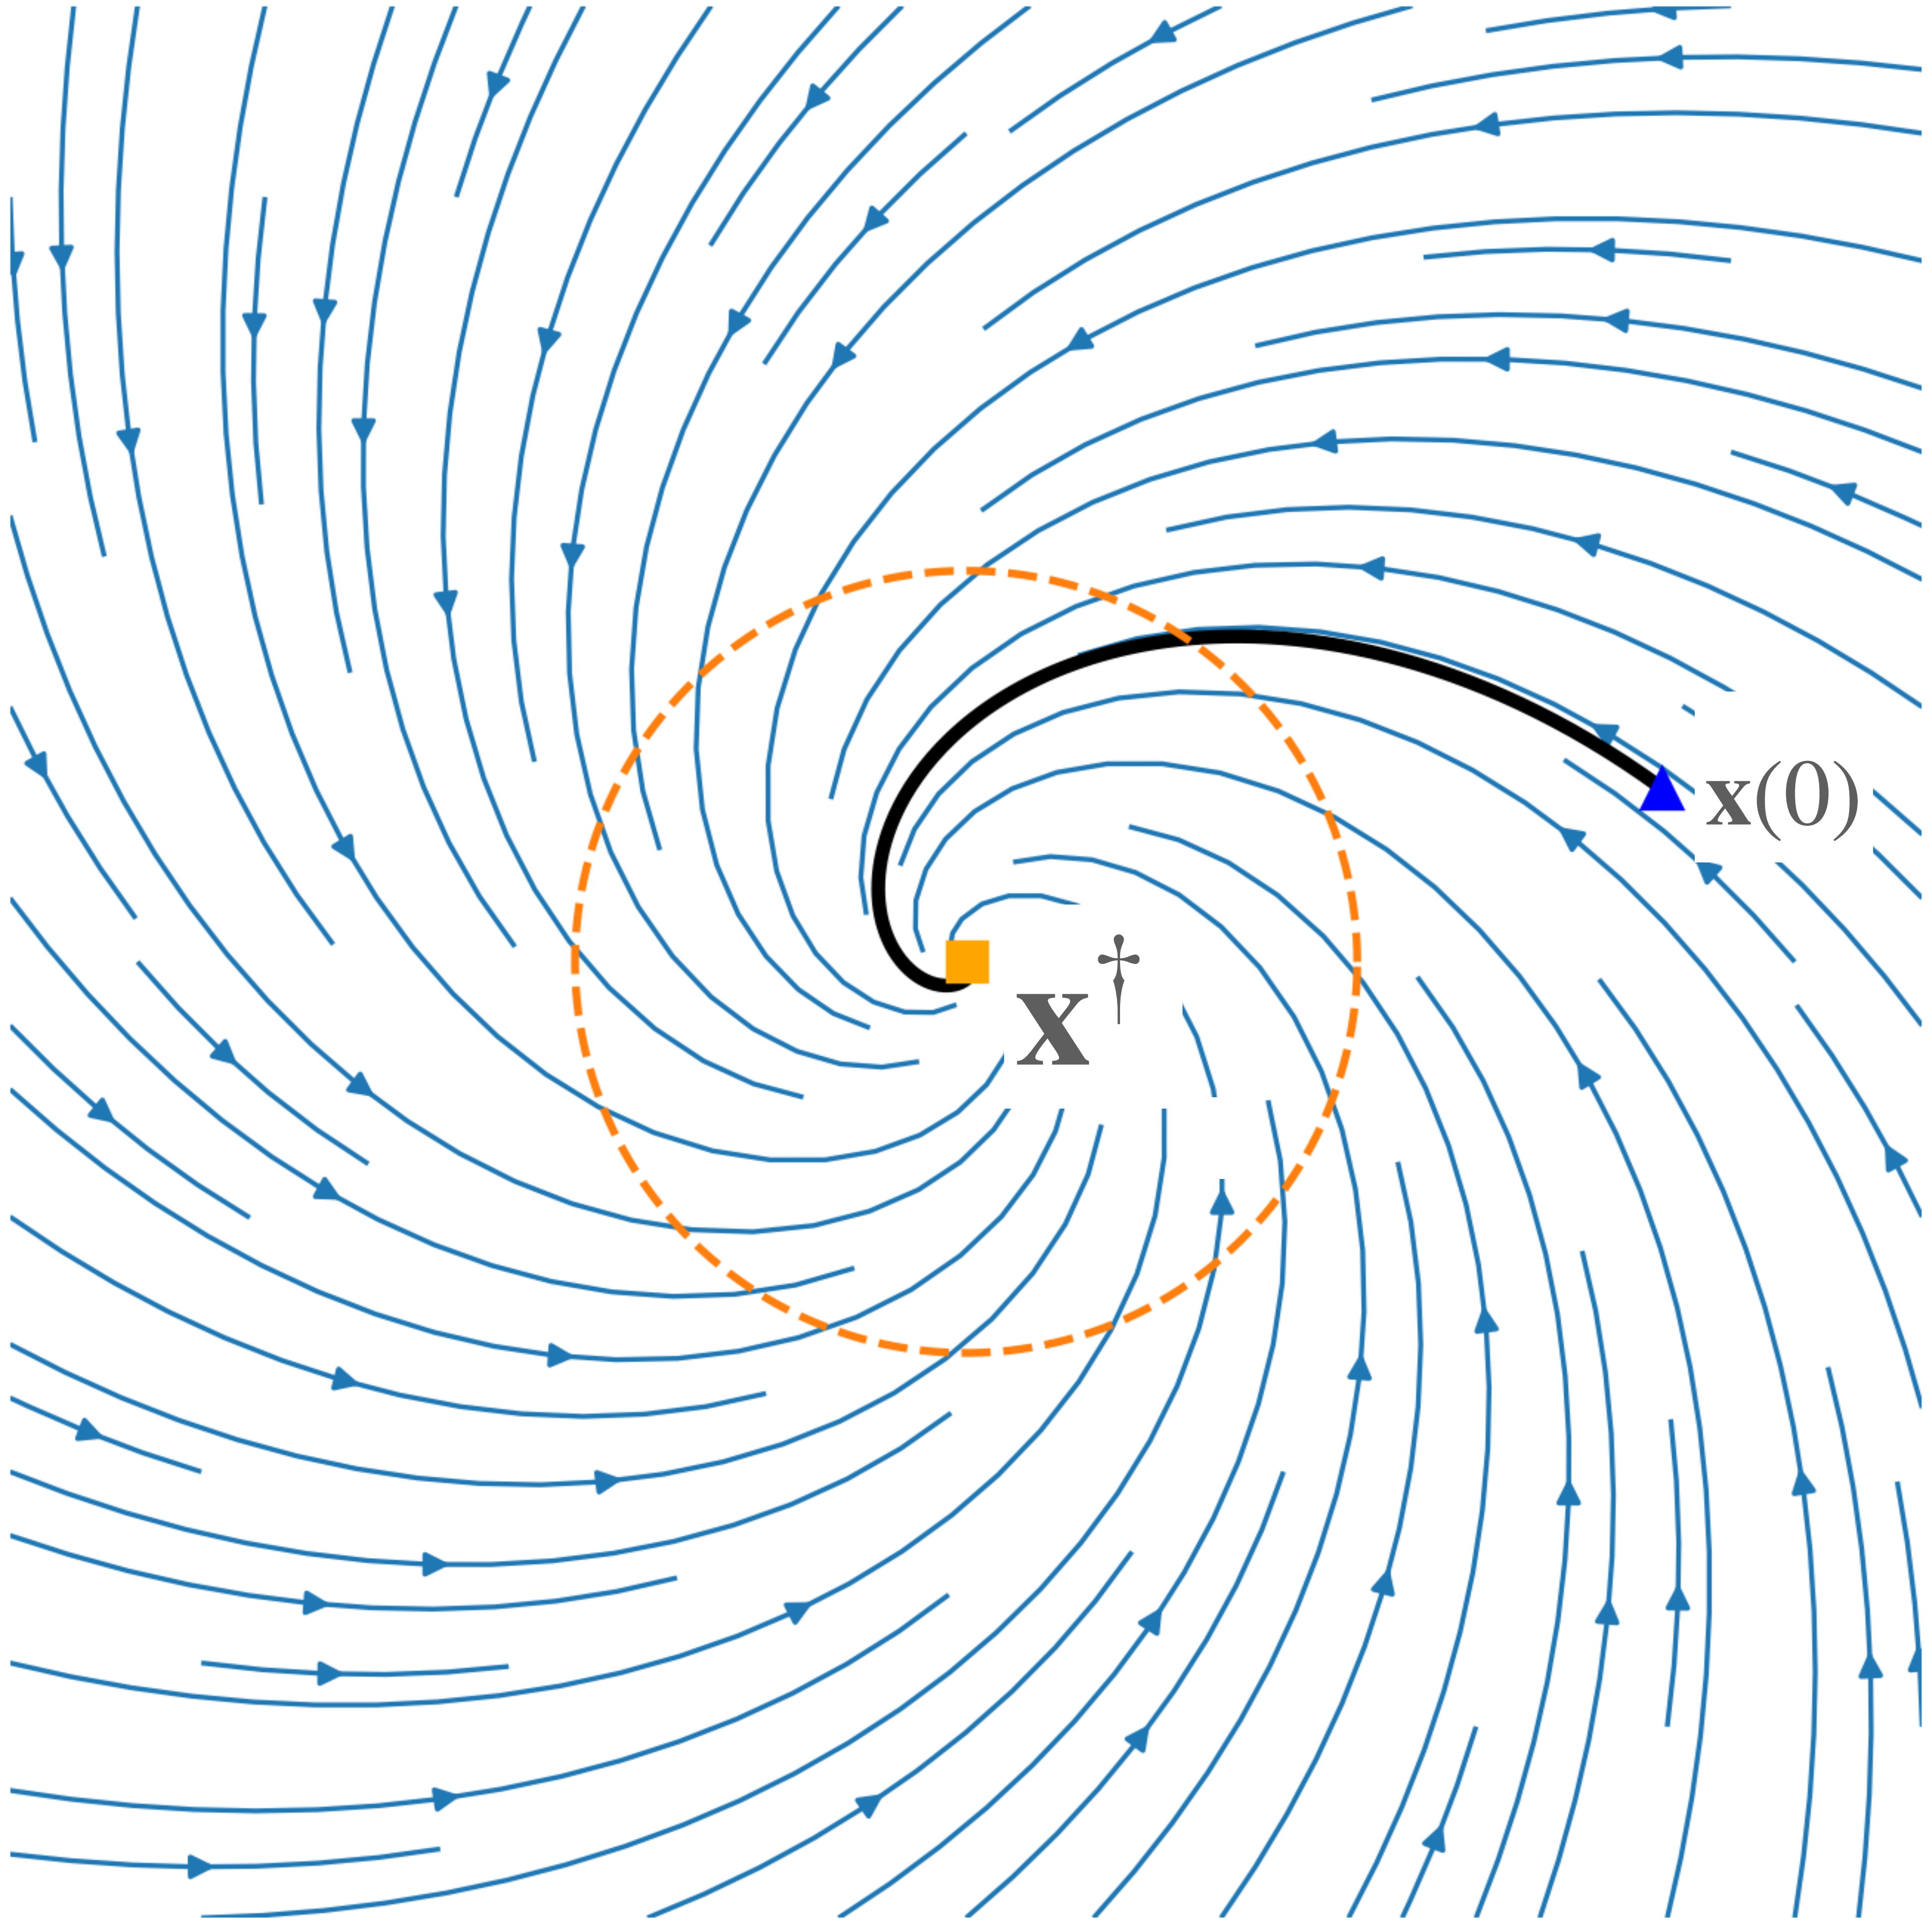

An equilibrium $\mathbf{x}^\dagger$ is asymptotically stable if there is some $\delta > 0$ such that:

$$\|\mathbf{x}(0) - \mathbf{x}^\dagger\| < \delta \quad \Longrightarrow \quad \lim_{t\to\infty} \mathbf{x}(t) = \mathbf{x}^\dagger$$

A strong notion of stability: convergence toward equilibrium is stable against small perturbations.

2. Lyapunov (Non-Asymptotic) Stability

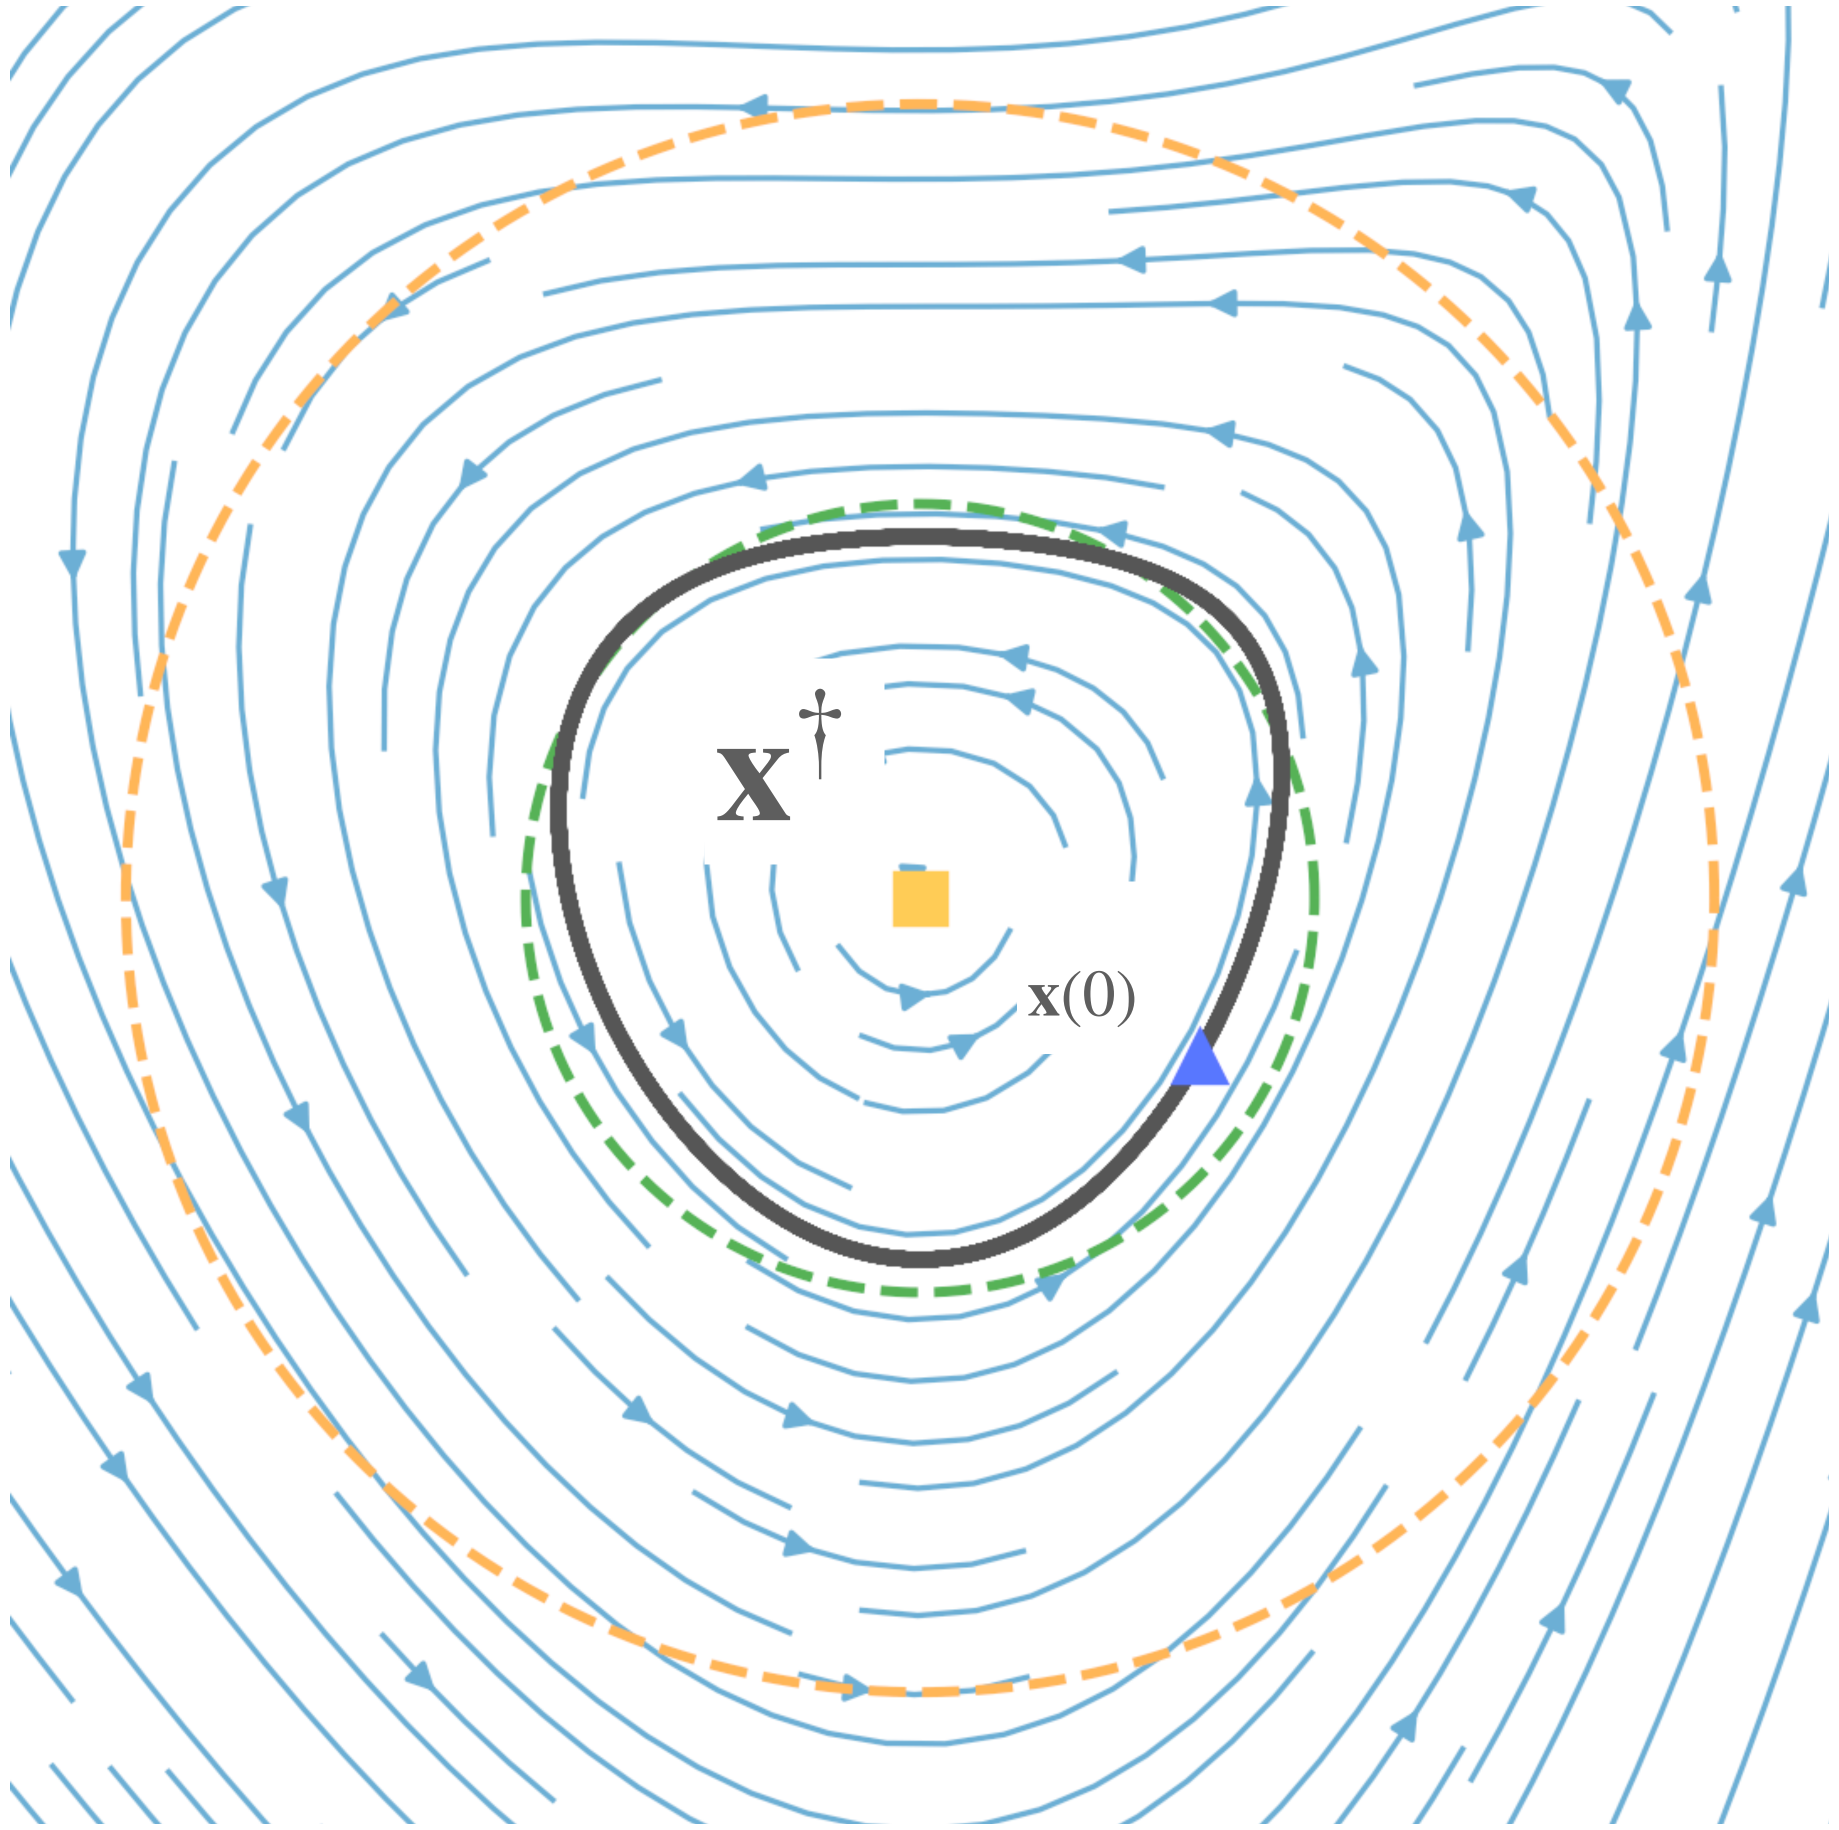

An equilibrium $\mathbf{x}^\dagger$ is **Lyapunov stable** if for all $\varepsilon > 0$, there is some $\delta > 0$ such that:

$$\|\mathbf{x}(0) - \mathbf{x}^\dagger\| < \delta \quad \Longrightarrow \quad \sup_{t>0} \|\mathbf{x}(t) - \mathbf{x}^\dagger\| < \varepsilon$$

A weaker notion of stability: dynamics near equilibrium remains close to it.

3. Unstable

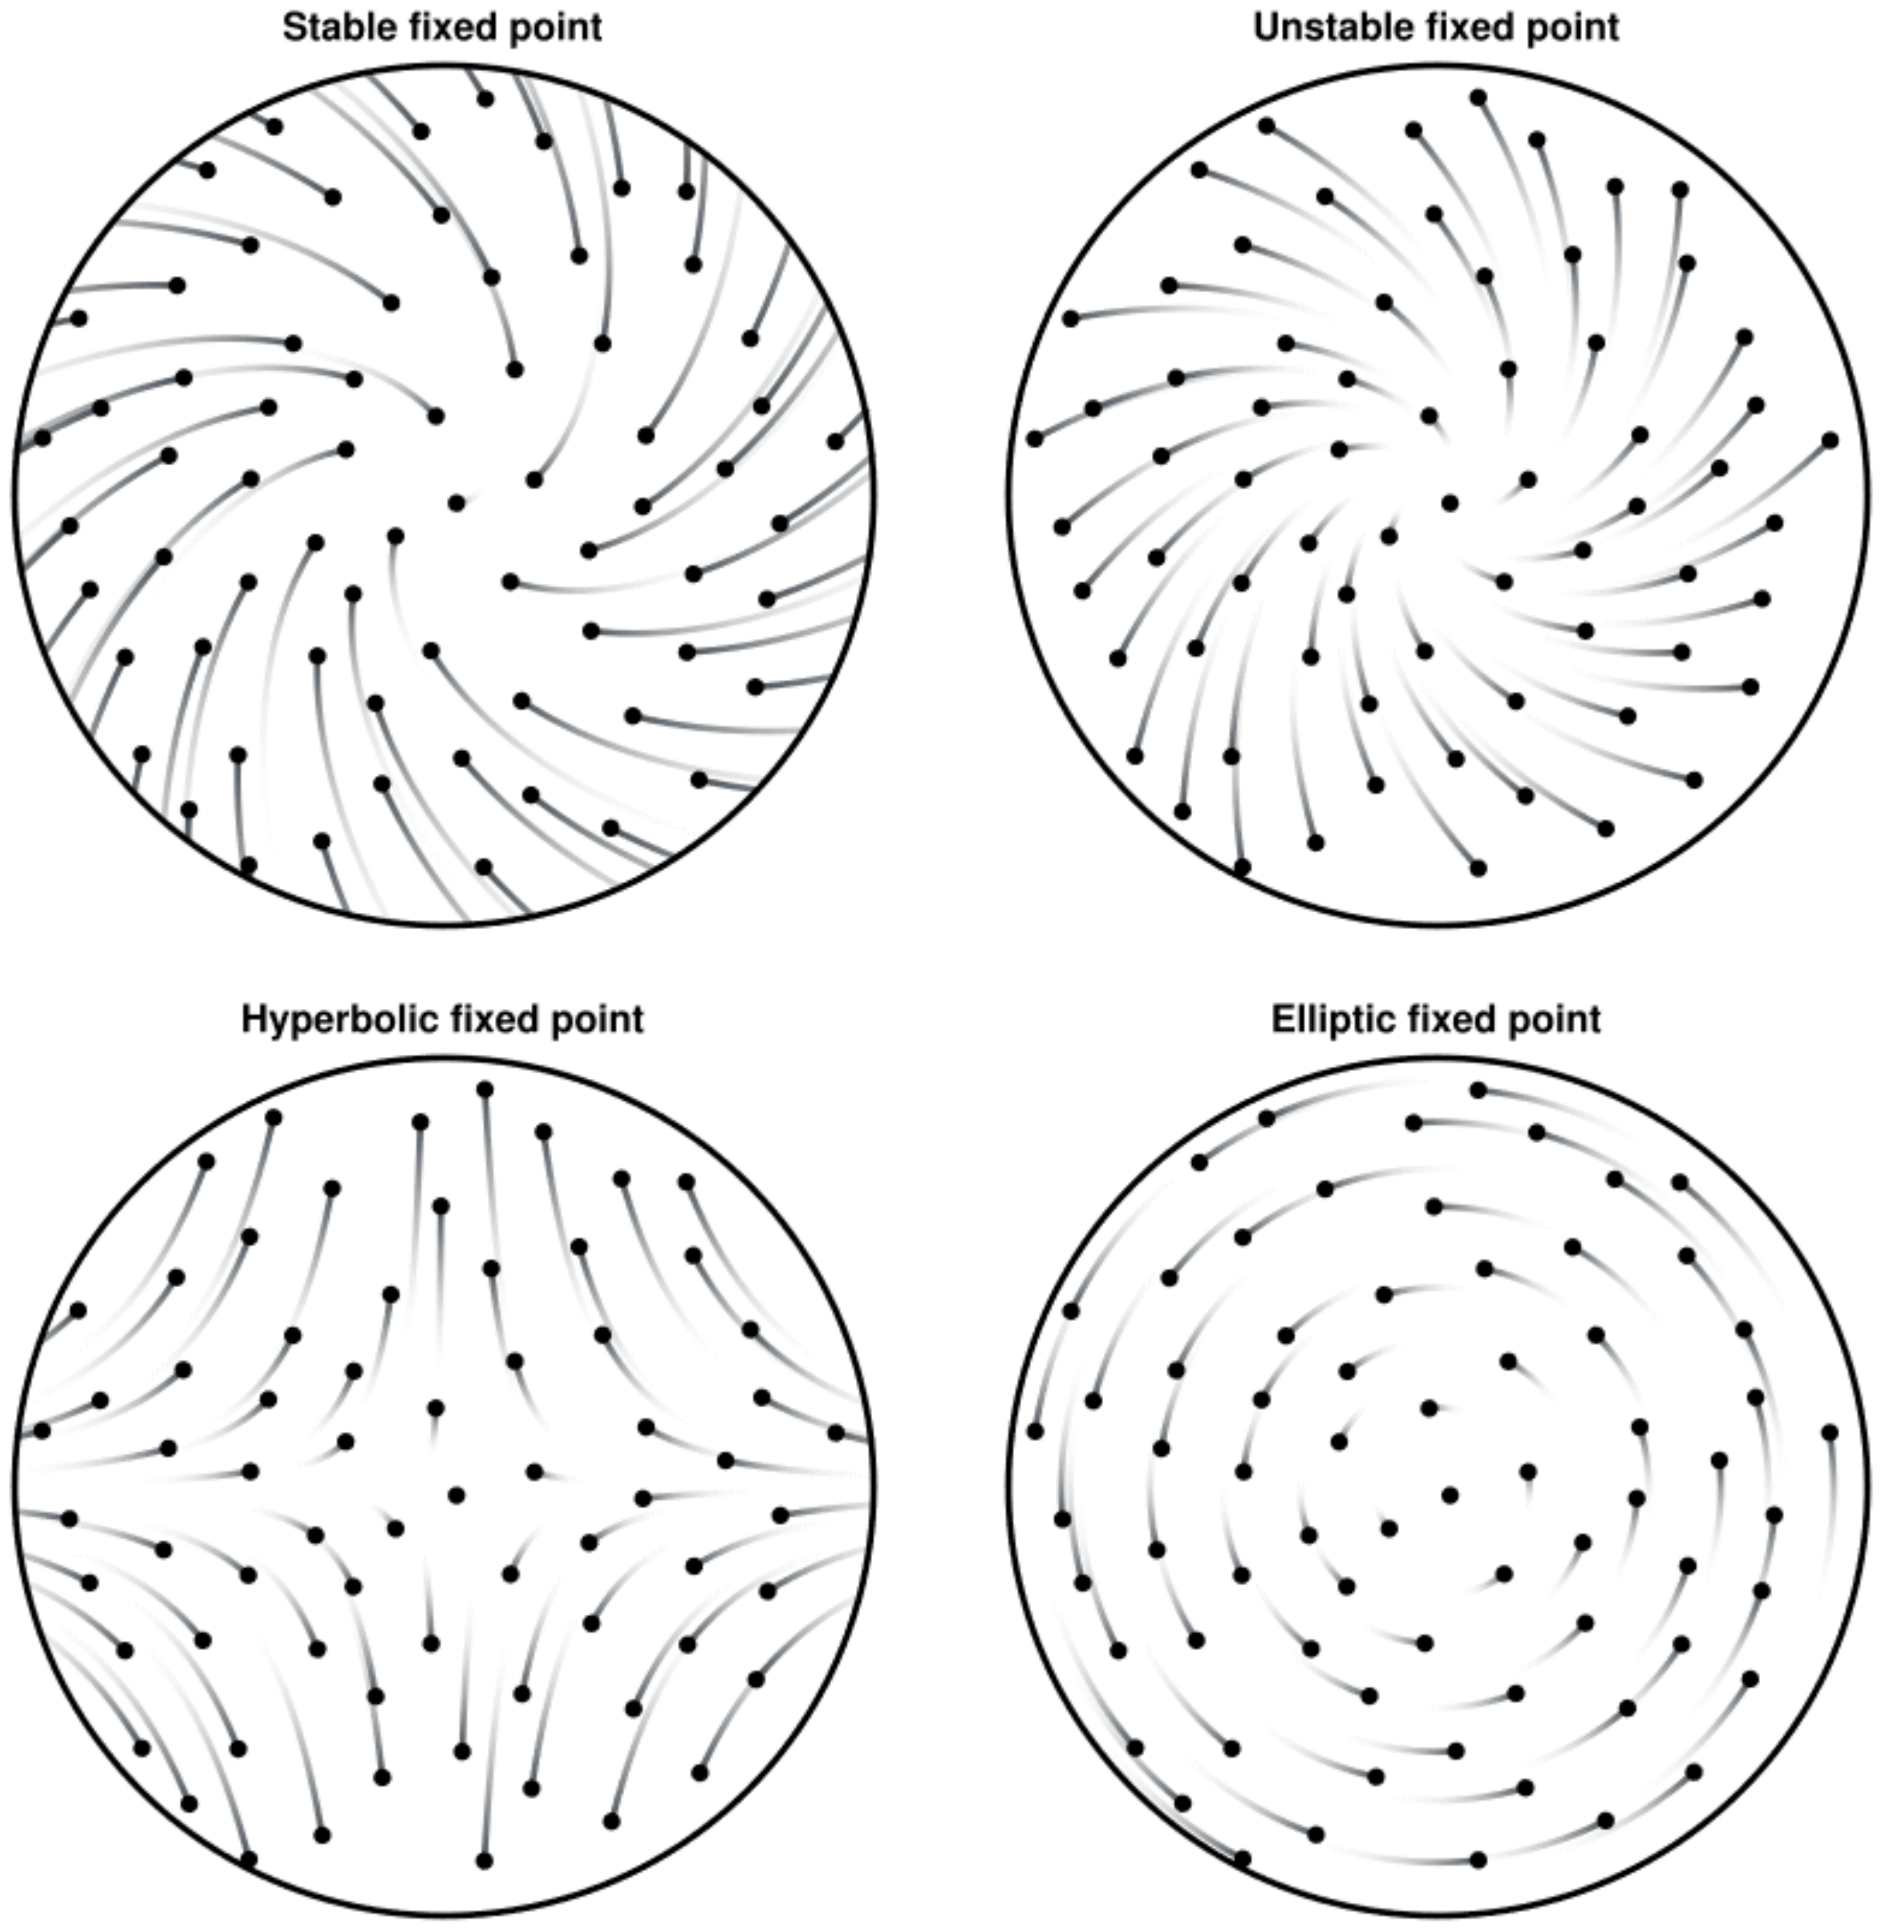

An equilibrium $\mathbf{x}^\dagger$ is unstable if it is not Lyapunov stable.

Linear Stability Analysis

Consider the following continuous-time, smooth dynamical system:

$$\dot{\mathbf{x}}(t) = T\big(\mathbf{x}(t)\big).$$

- We can study the **stability** of the dynamics around a fixed point $\mathbf{x}^\dagger$ by looking at the **linearized dynamics** around $\mathbf{x}^\dagger$.

Linearized Dynamics

For simplicity, let $\mathbf{x} \in \mathbb{R}^N$. The **Jacobian** $\nabla T(\mathbf{x}) \in \mathbb{R}^{N \times N}$ is given by:

$$\nabla T = \begin{bmatrix} \nabla_1 T_1 & \nabla_2 T_1 & \cdots & \nabla_N T_1 \\ \nabla_1 T_2 & \nabla_2 T_2 & & \\ \vdots & & \ddots & \\ \nabla_1 T_N & & & \nabla_N T_N \end{bmatrix}$$

Linearized Dynamics

For simplicity, let $\mathbf{x} \in \mathbb{R}^N$. The **Jacobian** $\nabla T(\mathbf{x}) \in \mathbb{R}^{N \times N}$ is given by:

$$\nabla T = \begin{bmatrix} \nabla_1 T_1 & \nabla_2 T_1 & \cdots & \nabla_N T_1 \\ \nabla_1 T_2 & \nabla_2 T_2 & & \\ \vdots & & \ddots & \\ \nabla_1 T_N & & & \nabla_N T_N \end{bmatrix}$$

Around the fixed point, the linearized dynamics are defined by:

$$\dot{\mathbf{z}}(t) = \nabla T(\mathbf{x}^\dagger)\, \mathbf{z}(t)$$

where $\mathbf{z} \equiv \mathbf{x} - \mathbf{x}^\dagger$.

Stability from Spectral Analysis

Stability of **linear** dynamical systems well-understood:

Stability from Spectral Analysis

Stability of **linear** dynamical systems well-understood:

- If all eigenvalues of ∇T have **negative real parts** $\implies$ asymptotic stability

Stability from Spectral Analysis

Stability of **linear** dynamical systems well-understood:

- If all eigenvalues of ∇T have **negative real parts** $\Longrightarrow$ asymptotic stability
- If there is an eigenvalue with **positive real part** $\Longrightarrow$ instability

Stability from Spectral Analysis

Stability of **linear** dynamical systems well-understood:

- If all eigenvalues of ∇T have **negative real parts** $\implies$ asymptotic stability
- If there is an eigenvalue with **positive real part** $\implies$ instability
- If all eigenvalues are **purely imaginary** $\implies$ neutral stability

Stability from Spectral Analysis

Stability of **linear** dynamical systems well-understood:

- If all eigenvalues of ∇T have **negative real parts** $\implies$ asymptotic stability
- If there is an eigenvalue with **positive real part** $\implies$ instability
- If all eigenvalues are **purely imaginary** $\implies$ neutral stability
  - This case is much harder to characterize for **non-linear** systems

Linearizing Gradient Ascent Dynamics

The gradient ascent dynamics:

$$\dot{x}_n(t) = \nabla_n u_n\big(\mathbf{x}(t)\big)$$

Linearizing Gradient Ascent Dynamics

The gradient ascent dynamics:

$$\dot{x}_n(t) = \nabla_n u_n\big(\mathbf{x}(t)\big)$$

- The Jacobian of this dynamics is given by:

$$\mathbf{J} = \begin{bmatrix} \nabla^2_{11} u_1 & \nabla^2_{12} u_1 & \cdots & \nabla^2_{1N} u_1 \\ \nabla^2_{21} u_2 & \nabla^2_{22} u_2 & & \\ \vdots & & \ddots & \\ \nabla^2_{N1} u_N & & & \nabla^2_{NN} u_N \end{bmatrix}$$

Linearizing Gradient Ascent Dynamics

The gradient ascent dynamics:

$$\dot{x}_n(t) = \nabla_n u_n\big(\mathbf{x}(t)\big)$$

- The Jacobian of this dynamics is given by:

$$\mathbf{J} = \begin{bmatrix} 0 & \nabla^2_{12} u_1 & \cdots & \nabla^2_{1N} u_1 \\ \nabla^2_{21} u_2 & 0 & & \\ \vdots & & \ddots & \\ \nabla^2_{N1} u_N & & & 0 \end{bmatrix}$$

  - The diagonal is zero by multilinearity.

Linearizing Gradient Ascent Dynamics

The gradient ascent dynamics:

$$\dot{x}_n(t) = \nabla_n u_n\big(\mathbf{x}(t)\big)$$

- The Jacobian of this dynamics is given by:

$$\mathbf{J} = \begin{bmatrix} 0 & \nabla^2_{12} u_1 & \cdots & \nabla^2_{1N} u_1 \\ \nabla^2_{21} u_2 & 0 & & \\ \vdots & & \ddots & \\ \nabla^2_{N1} u_N & & & 0 \end{bmatrix}$$

  - This object is called the **game Jacobian.**

Jacobian of learning dynamics

It turns out that the Jacobian of many learning dynamics is of the form:

$$\nabla T = \mathbf{H}^{-1}\mathbf{J}$$

where $\mathbf{J}$ is the game Jacobian and $\mathbf{H}$ is block-diagonal and positive definite:

$$\mathbf{H} = \begin{bmatrix} H_1 & & & \\ & H_2 & & \\ & & \ddots & \\ & & & H_N \end{bmatrix}$$

Jacobian of learning dynamics

- **Preconditioned gradient ascent**

$$\dot{x}_n(t) = H_n^{-1}\,\nabla_n u_n\big(\mathbf{x}(t)\big)$$

  - The preconditioned H_n also captures some local geometry.
    - Player-dependent notion of “steepest ascent”.

Jacobian of learning dynamics

- **Smoothed best-response dynamics**

$$x_n(t+1) = \arg\max_{z_n \in \Omega_n} \; u_n\big(z_n; \mathbf{x}_{-n}(t)\big) - \phi_n(z_n)$$

  - Usually ϕ_n is a smooth, strictly convex regularizer
    - Entropy map, negative-definite quadratic.
    - It defines the “local geometry” on Ω_n.

Jacobian of learning dynamics

- **Smoothed best-response dynamics**

$$x_n(t+1) = \arg\max_{z_n \in \Omega_n} \; u_n\big(z_n; \mathbf{x}_{-n}(t)\big) - \phi_n(z_n)$$

- Usually ϕ_n is a smooth, strictly convex regularizer

$$\nabla T(\mathbf{x}) = \mathbf{H}^{-1}\mathbf{J} \qquad \text{where} \quad \mathbf{H} = \begin{bmatrix} \nabla^2 \phi_1 & & & \\ & \nabla^2 \phi_2 & & \\ & & \ddots & \\ & & & \nabla^2 \phi_N \end{bmatrix}$$

The Jacobian of the dynamics are traceless

For learning dynamics whose Jacobians are of the form:

$$\nabla T = \mathbf{H}^{-1}\mathbf{J} = \begin{bmatrix} 0 & H_1^{-1}J_{12} & \cdots & H_1^{-1}J_{1N} \\ H_2^{-1}J_{21} & 0 & & \\ \vdots & & \ddots & \\ H_N^{-1}J_{N1} & & & 0 \end{bmatrix}$$

the trace is zero, $\mathrm{tr}\left(\mathbf{H}^{-1}\mathbf{J}\right) = 0$.

Barrier to Asymptotic Stability

- Suppose the dynamics around a fixed point is **traceless:** $\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = 0.$

Barrier to Asymptotic Stability

- Suppose the dynamics around a fixed point is **traceless:** $\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = 0.$
    - Recall that the trace is equal to the sum of the eigenvalues:

$$\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = \lambda_1 + \cdots + \lambda_N \,.$$

Barrier to Asymptotic Stability

- Suppose the dynamics around a fixed point is **traceless:** $\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = 0.$
    - Recall that the trace is equal to the sum of the eigenvalues:

$$\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = \lambda_1 + \cdots + \lambda_N\,.$$

- Asymptotic stability requires that the real part $\Re(\lambda_n) < 0$ for all $n \in [N]$.

Barrier to Asymptotic Stability

- Suppose the dynamics around a fixed point is **traceless:** $\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = 0$.
  - Recall that the trace is equal to the sum of the eigenvalues:

$$\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = \lambda_1 + \cdots + \lambda_N \,.$$

- Asymptotic stability requires that the real part $\Re(\lambda_n) < 0$ for all $n \in [N]$.
  - If some λ_n has a negative real part, then another λ_m must have positive real part. Then, the dynamics are **unstable**.

Barrier to Asymptotic Stability

- Suppose the dynamics around a fixed point is **traceless:** $\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = 0$.
  - Recall that the trace is equal to the sum of the eigenvalues:

$$\mathrm{tr}\Big(\nabla T(\mathbf{x}^\dagger)\Big) = \lambda_1 + \cdots + \lambda_N .$$

- Asymptotic stability requires that the real part $\Re(\lambda_n) < 0$ for all $n \in [N]$.
  - If some λ_n has a negative real part, then another λ_m must have positive real part. Then, the dynamics are **unstable**.

This is the very high-level argument used across many works showing that certain dynamics fail to converge stably to interior Nash equilibria.

Stability from Spectral Analysis

- ~~If all eigenvalues of ∇T have **negative real parts** $\implies$ asymptotic stability~~
- If there is an eigenvalue with **positive real part** $\implies$ instability
- If all eigenvalues are **purely imaginary** $\implies$ neutral stability

The remaining possibility

Neutral stability $\implies$ purely imaginary eigenvalues

Uniform Stability: Non-Asymptotic Stability

Definition. The game Jacobian $\mathbf{J}(\mathbf{x})$ is **uniformly stable** at $\mathbf{x}$ if the eigenvalues of $\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})$ are purely imaginary for all positive-definite block-diagonal matrix $\mathbf{H}$,

$$\mathrm{spec}\left(\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})\right) \subset i\mathbb{R}.$$

Uniform Stability: Non-Asymptotic Stability

Definition. The game Jacobian $\mathbf{J}(\mathbf{x})$ is **uniformly stable** at $\mathbf{x}$ if the eigenvalues of $\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})$ are purely imaginary for all positive-definite block-diagonal matrix $\mathbf{H}$,

$$\mathrm{spec}\left(\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})\right) \subset i\mathbb{R}.$$

- A Nash equilibrium $\mathbf{x}^\star$ is pointwise uniformly stable if $\mathbf{J}(\mathbf{x}^\star)$ is uniformly stable.

Uniform Stability: Non-Asymptotic Stability

Definition. The game Jacobian $\mathbf{J}(\mathbf{x})$ is **uniformly stable** at $\mathbf{x}$ if the eigenvalues of $\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})$ are purely imaginary for all positive-definite block-diagonal matrix $\mathbf{H}$,

$$\mathrm{spec}\left(\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})\right) \subset i\mathbb{R}.$$

- A Nash equilibrium $\mathbf{x}^\star$ is pointwise uniformly stable if $\mathbf{J}(\mathbf{x}^\star)$ is uniformly stable.
- It is locally uniformly stable if there is an open set U containing $\mathbf{x}^\star$ such that $\mathbf{J}(\mathbf{x})$ is uniformly stable for all $\mathbf{x} \in U$.

Uniform Stability: Motivation

Definition. The game Jacobian $\mathbf{J}(\mathbf{x})$ is **uniformly stable** at $\mathbf{x}$ if the eigenvalues of $\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})$ are purely imaginary for all positive-definite block-diagonal matrix $\mathbf{H}$,

$$\mathrm{spec}\left(\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})\right) \subset i\mathbb{R}.$$

- The imaginary spectrum captures the only chance for the dynamics to exhibit neutral stability (instead of instability).

Uniform Stability: Motivation

Definition. The game Jacobian $\mathbf{J}(\mathbf{x})$ is **uniformly stable** at $\mathbf{x}$ if the eigenvalues of $\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})$ are purely imaginary for all positive-definite block-diagonal matrix $\mathbf{H}$,

$$\operatorname{spec}\left(\mathbf{H}^{-1}\mathbf{J}(\mathbf{x})\right) \subset i\mathbb{R}.$$

- The imaginary spectrum captures the only chance for the dynamics to exhibit neutral stability (instead of instability).
- The quantifier "for all $\mathbf{H}$" expands the notion of *uncoupledness*: players do not know how others learn from their experiences.

Technical Detail: Strategic Equivalence

All of the dynamics mentioned factor through *strategic equivalence*:

- Take two utilities u_n and $u_n' = u_n + v_n$ where $v_n(\mathbf{x}_{-n})$ does not depend on x_n. The dynamics for player n are the same for both. For example:

$$\nabla_n u_n \equiv \nabla_n \left(u_n + v_n\right).$$

- Parts of the utility can be "external" to the strategic component of the game.
  - We focus on **purely strategic** utilities (or, define strategic Pareto optimality to depend only on the strategic component of the utility).

Technical Detail: Non-Degenerate Game

Our results hold under somewhat mild degeneracy assumptions.

- Consider the Taylor expansion of $u_n(\mathbf{x}) - u_n(\mathbf{x}^\star)$ about the equilibrium.
  - At the equilibrium, the linear terms are zero; lowest-order terms are *bilinear*.
  - The game Jacobian $\mathbf{J}$ precisely captures this bilinear information.
  - Our assumptions roughly make sure that the lower-order terms dominate.

The Meaning of Uniform Stability

An Algebraic Characterization

Theorem. Under non-degeneracy, the game Jacobian $\mathbf{J}$ is uniformly stable if and only if there are weights $\lambda_1, \ldots, \lambda_N > 0$ so that $\mathbf{\Lambda J}$ is skew-symmetric where:

$$\mathbf{\Lambda} = \begin{bmatrix} \lambda_1 I & & & \\ & \lambda_2 I & & \\ & & \ddots & \\ & & & \lambda_N I \end{bmatrix}$$

- Thus, we can rescale the utilities $\tilde{u}_n = \lambda_n u_n$ and the corresponding Jacobian $\tilde{\mathbf{J}}$ is skew-symmetric.

Local Uniform Stability implies Pareto Optimality

Theorem. Assuming non-degeneracy. If $\mathbf{x}^\star$ be a mixed Nash equilibrium, then:

$$\text{local uniform stability} \quad \Longrightarrow \quad \text{strategic Pareto optimality.}$$

A simple class: polymatrix games

In a polymatrix game, each utility u_n decomposes as:

$$u_n(\mathbf{x}) = \sum_{m \neq n} (x_n - x_n^\star)^\top J_{nm}(x_m - x_m^\star).$$

- Players are simply playing multiple two-player games at once.
- There are only bilinear interactions, and no higher-order interactions.
- The game Jacobian fully characterizes (the strategic component of) the game.

Uniform Stability implies Pareto Optimality

Theorem (polymatrix version). Consider a non-degenerate, purely strategic, polymatrix game. Let $\mathbf{x}^\star$ be a mixed Nash equilibrium. Then:

$$\text{uniform stability} \qquad \Longrightarrow \qquad \text{Pareto optimality.}$$

Proof

1. **By the algebraic characterization, there are $\lambda_1, \ldots, \lambda_N > 0$ such that:**

$$\lambda_n J_{nm} + \lambda_m J_{mn}^\top = 0.$$

2. **It follows that the sum of rescaled utilities is zero:**

$$\sum_{n \in [N]} \lambda_n u_n(\mathbf{x}) = \sum_{n \in [N]} \sum_{m \neq n} \lambda_n (x_n - x_n^\star)^\top J_{nm} (x_m - x_m^\star) = 0.$$

   - The (n, m)-term in $\lambda_n J_{nm}$ cancels with the (m, n)-term of $\lambda_m J_{mn}$.

3. **Thus, there is no way to choose $\mathbf{x}$ so that $u_n(\mathbf{x}) > 0$ for all $n \in [N]$.**

   - It follows that $\mathbf{x}^\star$ is Pareto optimal.

A comment for general games

In general-sum games, there may be higher-order interactions.

- It turns out that local uniform stability imposes a lot of structure on the utilities.
- Proof is quite non-trivial, but the general scheme is to construct functions $\lambda_1(\mathbf{x}), \ldots, \lambda_N(\mathbf{x}) > 0$ such that in a neighborhood around $\mathbf{x}^\star$, we also have:

$$\sum_{n \in [N]} \lambda_n(\mathbf{x})\, u_n(\mathbf{x}) = 0.$$

Pareto Optimality implies Uniform Stability

Theorem. Assume non-degeneracy. If $\mathbf{x}^\star$ is a mixed Nash equilibrium. Then:

$$\text{Pareto optimality} \qquad \Longrightarrow \qquad \text{Uniform Stability.}$$

Proof (Two-Player Case)

1. Recenter the utilities $u_1 - u_1^\star$ and $u_2 - u_2^\star$ so that $u_n(\mathbf{x}^\star) = 0$.
2. Pareto optimality implies that $\operatorname{sign}(u_1) = -\operatorname{sign}(u_2)$.
3. The zero set of u_1 and u_2 coincide.
4. These utilities are polynomials, so by Hilbert's nullstellensatz, they are scalar multiples of each other (up to scaling, the game is zero-sum).
5. Their Jacobians are also scalar multiples of each other.
6. Apply algebraic characterization of uniform stability.

* Technically, the nullstellensatz does not apply. The actual proof uses SVD and is only slightly longer.

A comment for general games

For N-player games, we need to show the following result:

Theorem. Assume non-degeneracy. A mixed Nash equilibrium $\mathbf{x}^\star$ is weakly Pareto optimal if and only if it is a strong Nash equilibrium.

- A strong Nash equilibrium is one where coalitions of players cannot improve their utilities.
- The result proceeds by induction on coalitions.

Summary

1. **Connection between stability and collective rationality.**

   local uniform stability $\Longrightarrow$ Pareto optimality $\Longrightarrow$ pointwise uniform stability

Comparison of Dynamics Around Non-Strict NE

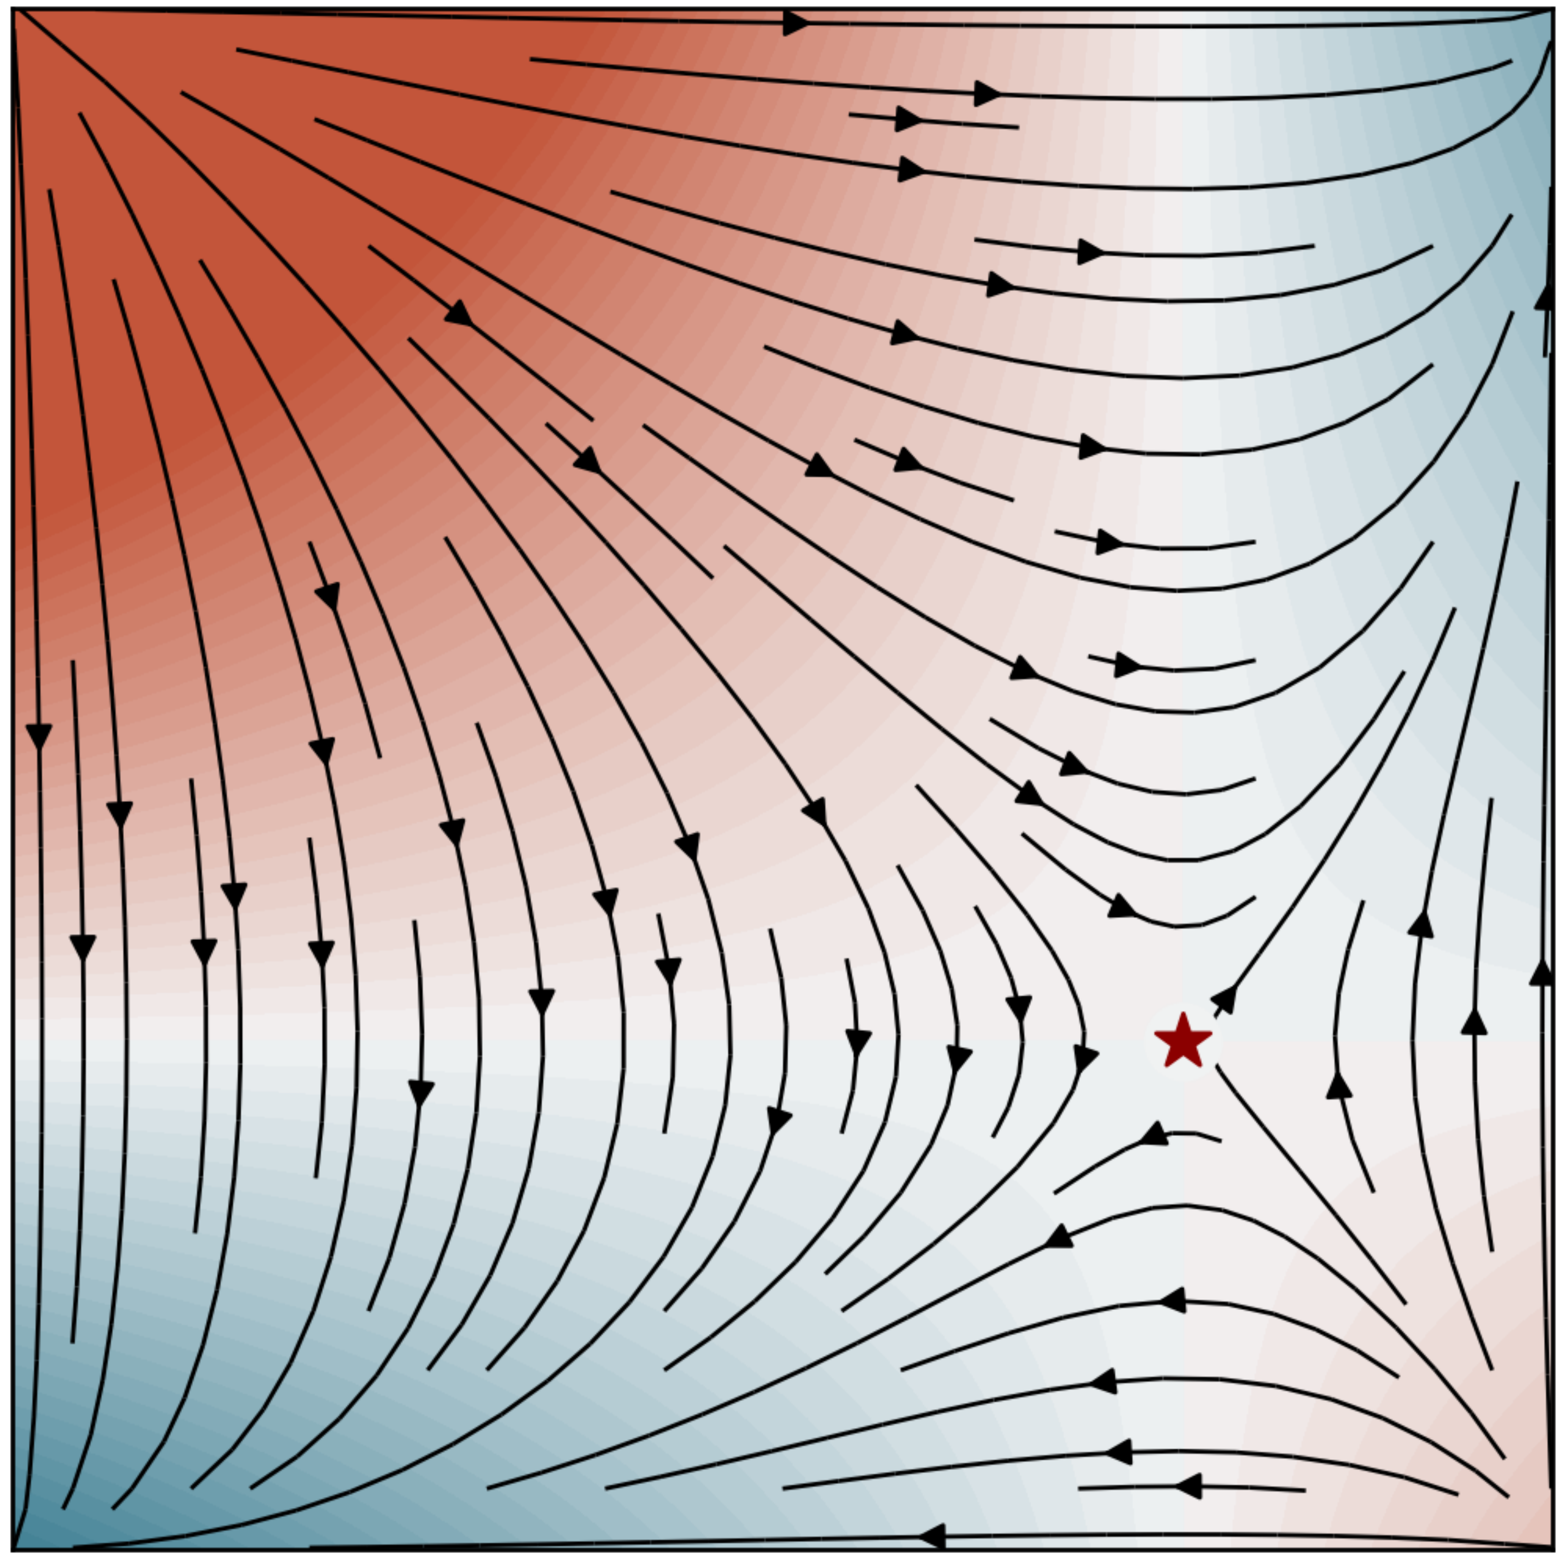

Unstable Nash equilibria are <u>not</u> strategically Pareto optimal.

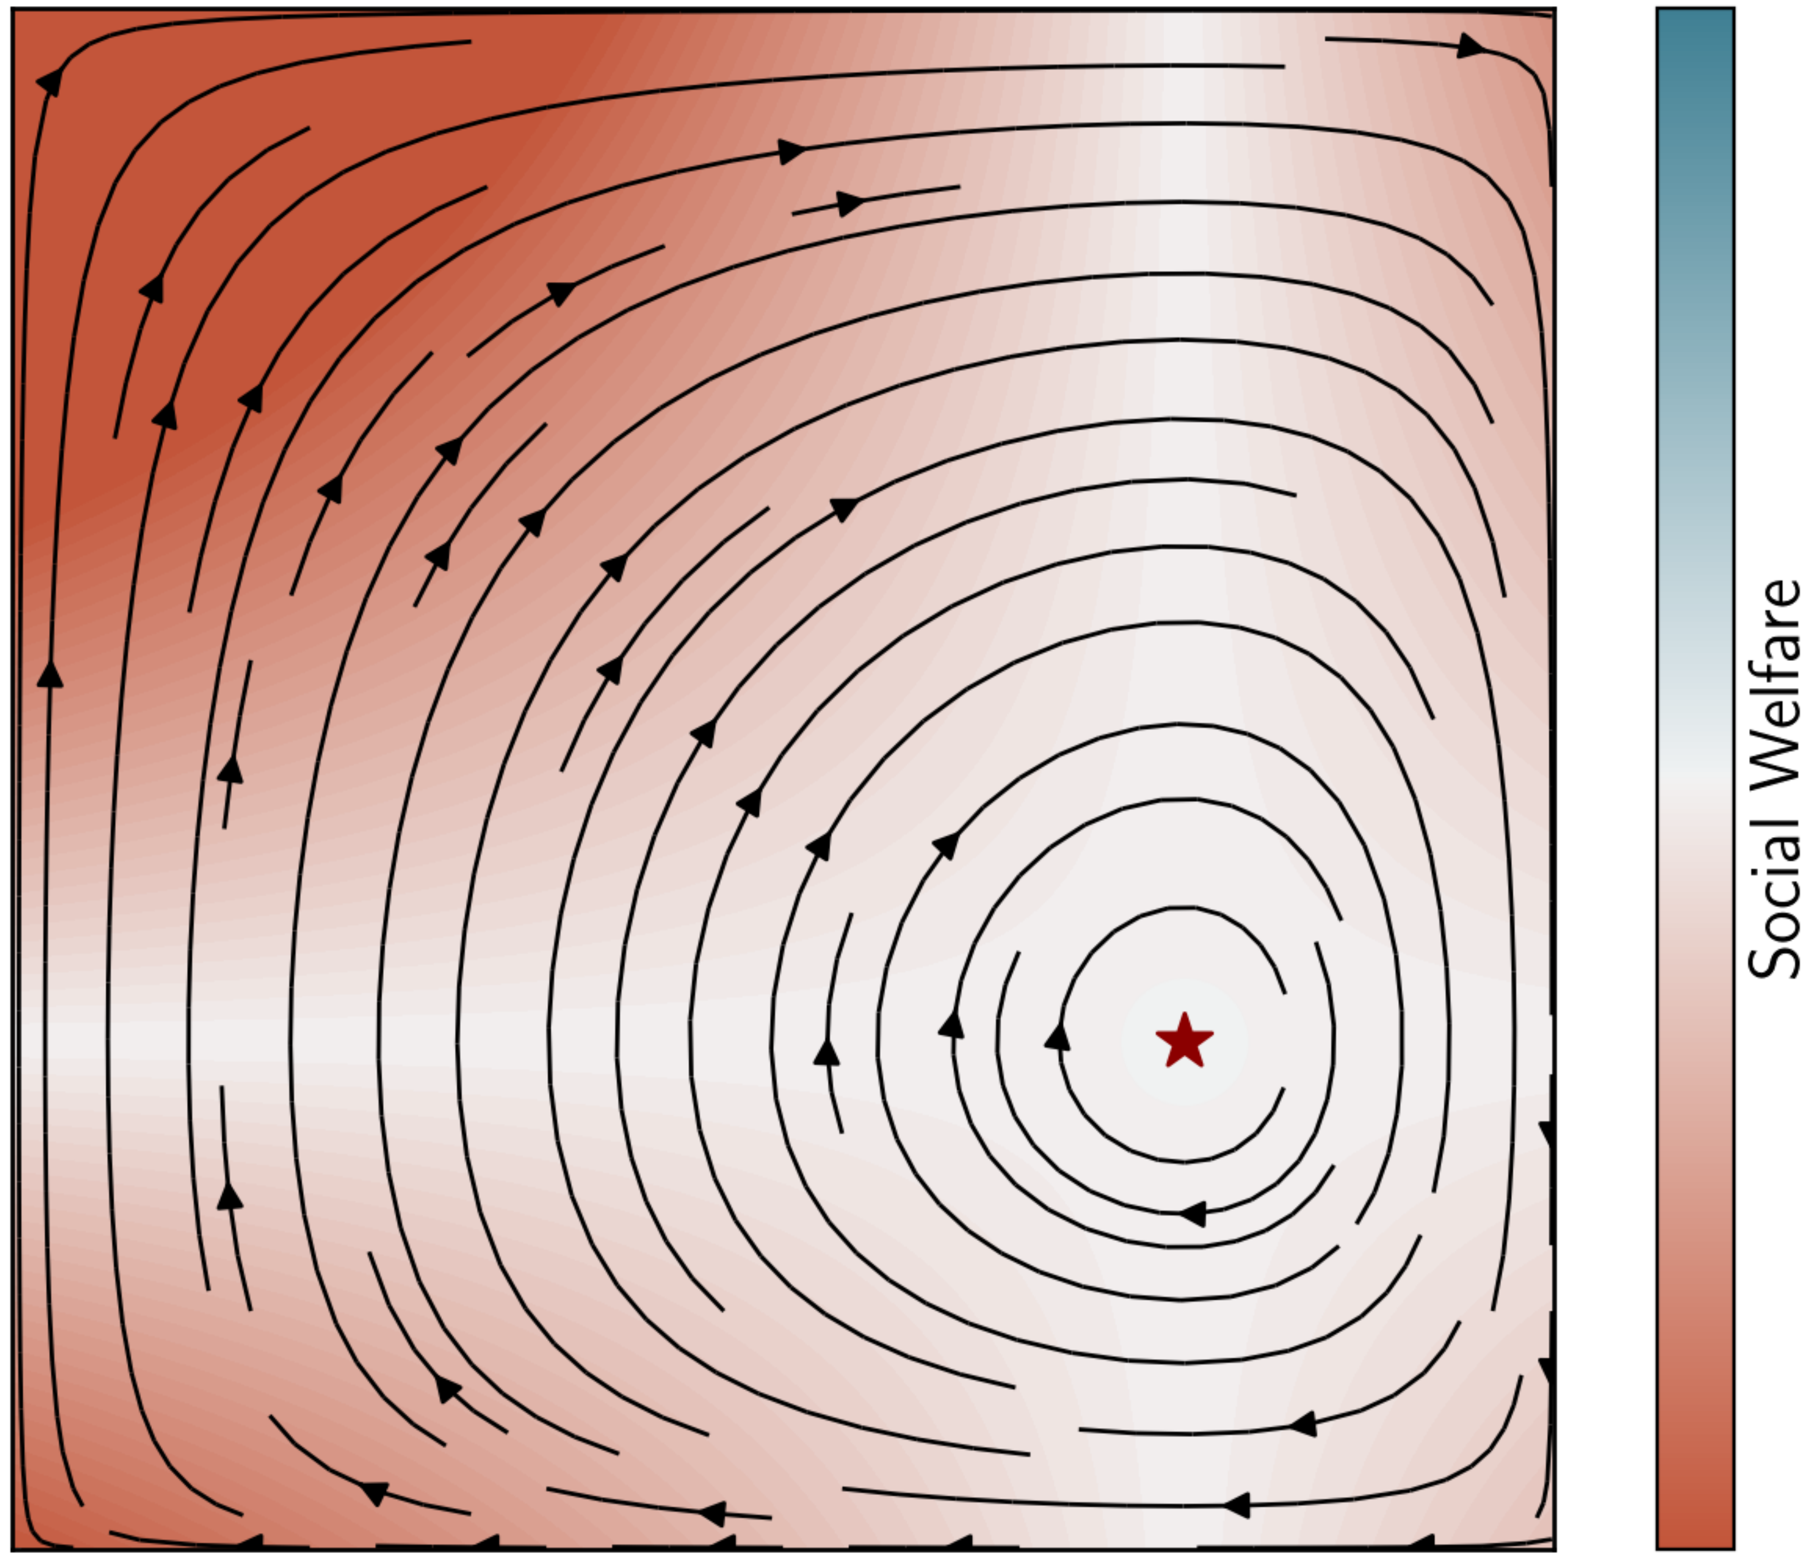

Uniformly stable Nash equilibria are strategically Pareto optimal.

Convergence under Uniform Stability

Smoothed best-response dynamics

Smoothed best-response map:

$$\Phi_n^{\beta}(\mathbf{x}) = \arg\max_{z_n \in \Omega_n} \; u_n(z_n; \mathbf{x}_{-n}) - \beta \phi_n(z_n)$$

- ϕ_n is a smooth, strictly convex regularizer.
- When ϕ_n is negative entropy, this yields the quantal response map:

$$\phi_n(x_n) = \sum_{i \in [k_n]} x_{n,i} \log x_{n,i} \qquad \text{and} \qquad \Phi_n^{\beta}(\mathbf{x}) \propto \exp\left(\frac{1}{\beta} \nabla_n u_n(\mathbf{x})\right)$$

- Smoothed best-response is a model of bounded rationality.
  - It approaches perfect rationality as $\beta \to 0$.

Incremental smoothed best-response dynamics

Incremental/partial updates:

$$\mathbf{x}(t+1) = (1-\eta)\,\mathbf{x}(t) + \eta\,\boldsymbol{\Phi}^{\beta}\big(\mathbf{x}(t)\big)$$

- This models the setting where each "player" corresponds to a population.
- In each round of the game, only an η-fraction of players from each population play the game.
  - η corresponds to a learning rate, governing how rapidly the dynamics change.

Incremental smoothed best-response dynamics

Smoothed equilibria: fixed points $\mathbf{x}^\beta$ of these dynamics.

- As the players become "more" rational ($\beta \to 0$), the equilibrium $\mathbf{x}^\beta \to \mathbf{x}^\star$ approaches the Nash equilibrium of the game.

- However, the dynamics also tend to become <u>less stable</u> as $\beta \to 0$.

- On the other hand, the dynamics tend to become <u>more stable</u> as $\eta \to 0$.

Stability of incremental dynamics

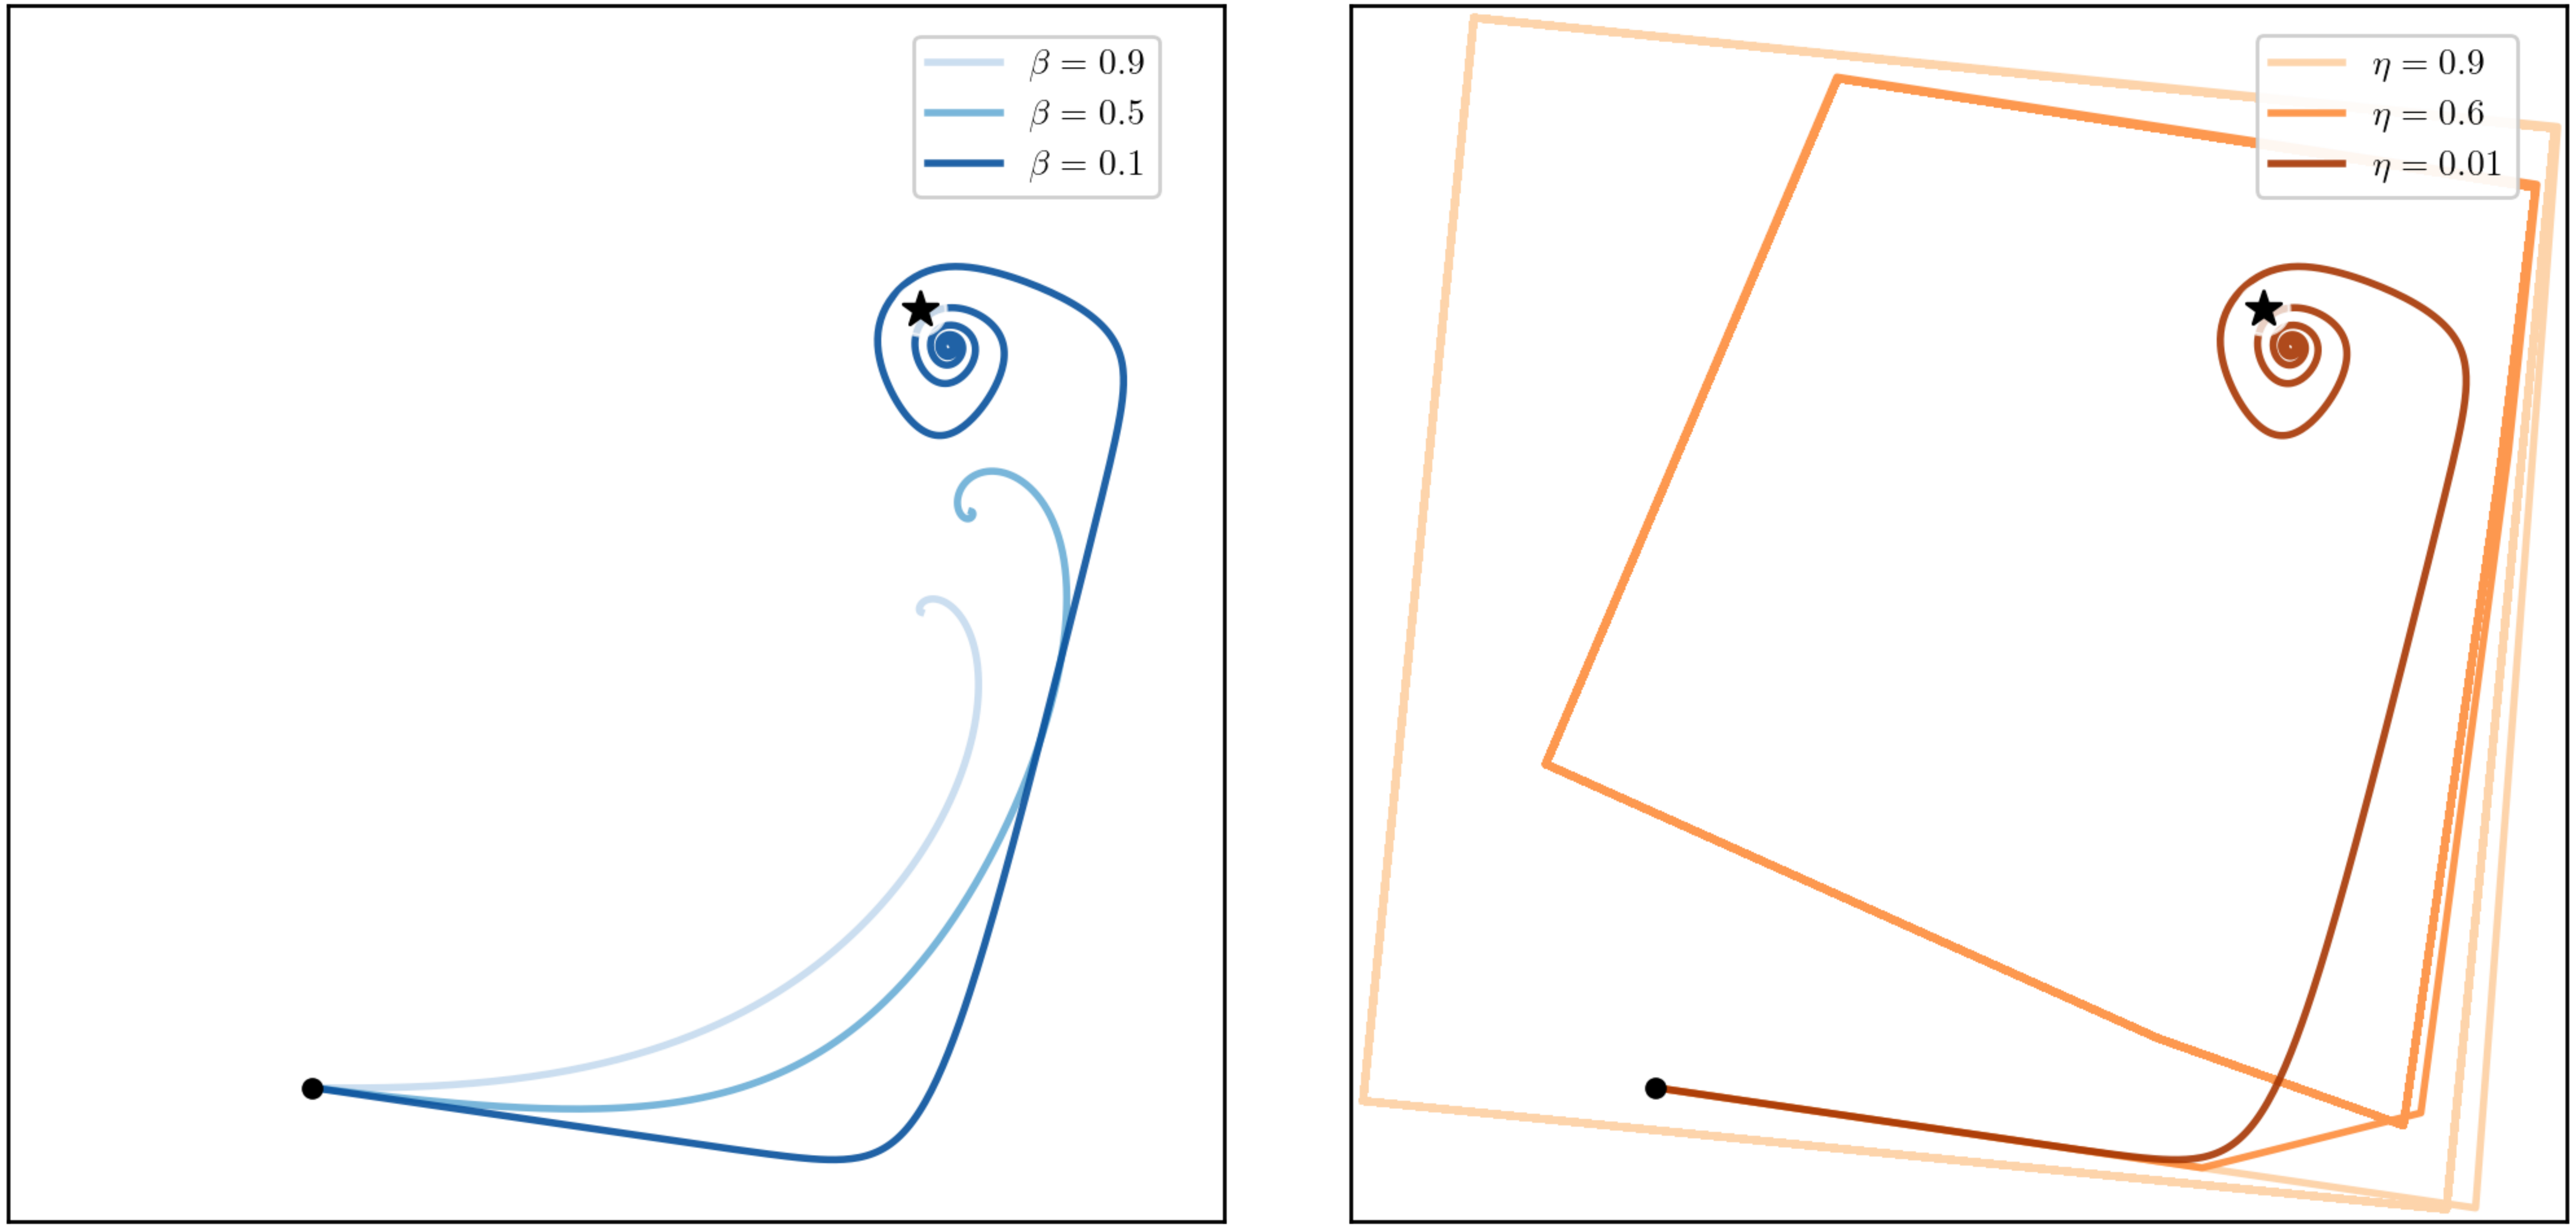

Stability of incremental dynamics

- The Jacobian of the dynamics is:

$$\nabla T(\mathbf{x}) = (1-\eta)\mathbf{I} + \frac{\eta}{\beta}\mathbf{H}(\mathbf{x})^{-1}\mathbf{J}(\mathbf{x})$$

- For discrete-time dynamics, the stability criterion is

$$\left|\lambda\left(\nabla T(\mathbf{x}^{\beta})\right)\right| < 1.$$

- If $\mathbf{J}$ is uniformly stable, then eigenvalues of the second term are imaginary, and choosing η sufficiently small stabilizes the dynamics.

- If $\mathbf{J}$ is not uniformly stable, then eigenvalues in the second term can have real part; once β becomes too small, there is no choice of η that can stabilize the dynamics.

Our results

1. **Connection between stability and collective rationality.**

   local uniform stability $\implies$ Pareto optimality $\implies$ pointwise uniform stability

2. **Dynamics of a specific class of learning rules.**
   - Incremental smooth best-response dynamics
   - Local uniform stability $\implies$ all dynamics can be stabilized
   - Not uniformly stable $\implies$ certain dynamics can never be stabilized

References

- Hart and Mas-Colell 2003, Uncoupled Dynamics Do Not Find Nash Equilibrium
- Samuelson and Zhang 1992, Evolutionary Stability in Asymmetric Games
- Vlatakis-Gkarangkounis et. al. 2020, No-regret learning and mixed Nash equilibria: They do not mix
- Giannou et. al. 2021, Survival of the strictest: Stable and unstable equilibria under regularized learning with partial information

Takeaways for CS Audience

- **Modern ML often implements multi-agent solutions.**
- **Decentralization introduces structural constraints.**
- **What are the ramifications and when are guardrails needed?**

Thank You!

Geelon So, geelon@ucsd.edu