Learning in Games and the Topology of Dynamical Systems

Georgios Piliouras

SINGAPORE UNIVERSITY OF TECHNOLOGY AND DESIGN

Christos Papadimitriou

Berkeley
UNIVERSITY OF CALIFORNIA

Ioannis Panageas

Georgia Tech

Game Theory Basics

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

Encodes economic interactions between many agents.
n number of agents.
Each agent i has a set of actions S_i to choose from.
The utility of each agent u_i is a function of all agent actions.

$$u_i : \times_i S_i \rightarrow \mathcal{R}$$

The main solution concept is that of Nash equilibrium:
A profile of strategies for each agent so that no agent can deviate unilaterally and increase his own payoff.

Game Theory Basics

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

	A	B
A	1,-1	-1,1
B	-1,1	1,-1

Encodes economic interactions between many agents.
n number of agents.
Each agent i has a set of actions S_i to choose from.
The utility of each agent u_i is a function of all agent actions.

$$u_i : \times_i S_i \to \mathcal{R}$$

The main solution concept is that of Nash equilibrium:
A profile of strategies for each agent so that no agent can deviate unilaterally and increase his own payoff.
Randomization is necessary for NE to always exist (Nash '50).

Nash Equilibria: The Nash equilibrium of solution concepts

Several aspects of Nash equilibrium as a solution concept are undesirable

Non-uniqueness

Computational hardness of computing a sample (even approximate) equilibrium in large games

Learning dynamics typically do not converge to equilibria even in simple small games

Nash Equilibria: The Nash equilibrium of solution concepts

A ***lot of work*** has been invested in reconciling these problems to the extent possible:

Non-uniqueness: **Axiomatic refinements of Nash equilibria** by economists.

Computational hardness of computing a sample (even approximate) equilibrium in large games: Exploring the tractability of computing **ϵ-approximate equilibria** by computer scientists.

Learning dynamics typically do not converge to equilibria even in simple small games: Focusing on games where dynamics behave well (e.g. **potential)** by evolutionary game theorists.

The "unspoken" thesis

Any solution concept *must* satisfy stationarity. Since 1920s. Unchallenged hypothesis across different communities.

Why should one be skeptical of this assumption:

Actual behavior of multi-body systems is seldom equilibrium:

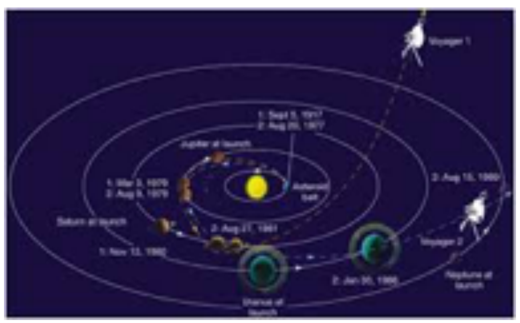

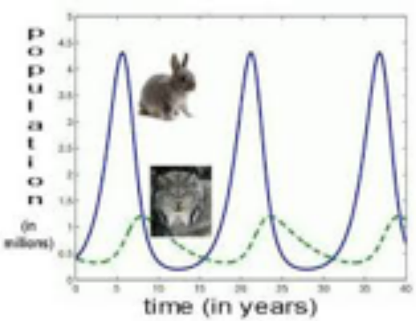

Epicycle Effects in Game Theory: When studying the implications of a false but deeply engrained assumption, a specific pattern of results emerges according to which more and more relaxations are necessary to accommodate the theory. Eventually it becomes unsustainable.

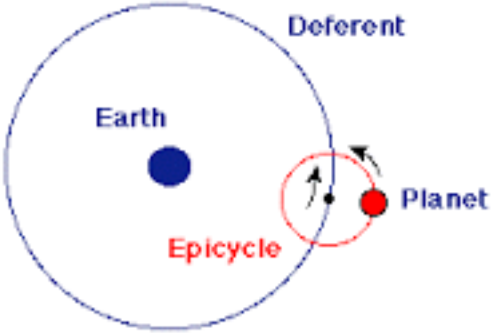

Epicycles

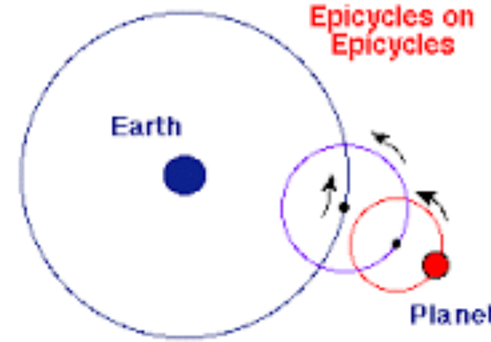

Stationary Earth Hypothesis: (a.k.a. Geocentric Theory)
Thesis: Earth does not move.
As computational/observational tools became more refined it became clear that in order to support the theory increasingly unrealistic explanations were needed. (e.g. Planets move along cycles whose centers move along cycles, e.t.c.) Number of epicycles needed grew from 6,8,80,… Then ***theory collapsed***.

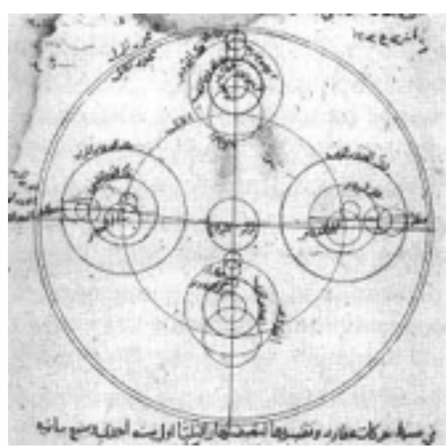

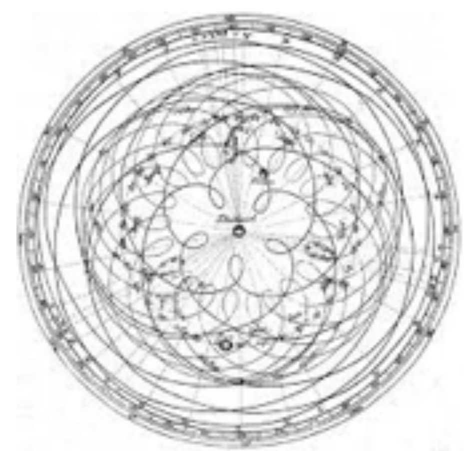

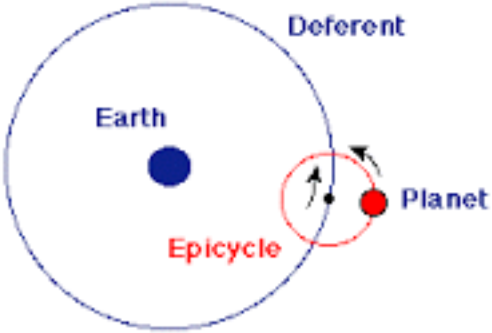

Epicycles

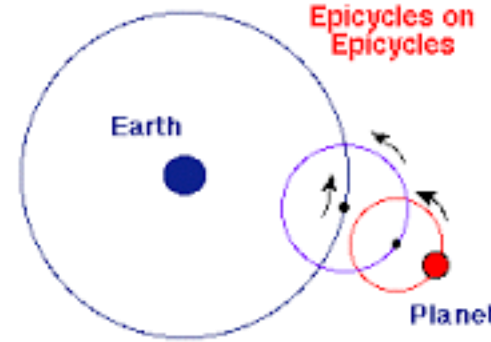

Stationary Earth Hypothesis: (a.k.a. Geocentric Theory)
Thesis: Earth does not move.
As computational/observational tools became more refined it became clear that in order to support the theory increasingly complex trajectories were needed. (e.g. Planets move along cycles whose centers move along cycles). Number of epicycles needed grew from 6,8,80,… Then ***theory collapsed***.

Similar phenomena in game theory:
Route 1: Inapproximability of Nash equilibria. [No efficient computation Daskalakis-Goldberg-Papadimitriou '06, No FPTAS Chen-Deng-Teng '06, No PTAS Rubinstein '15]
Route 2: Learning. Increasingly universal non-convergence results [Shapley'64, Books: Hofbauer-Sigmund'98, Sandholm'10]

Our thesis: A computational theory of games must be *generative*

There is *no equilibrium without a disequilibrium process.*

We need **a new solution concept**:

Applicable to any game and any dynamic!

That does not share the weaknesses of Nash equilibrium.

Effectively a solution concept that captures all dynamical systems. If such a result existed, it would be a fundamental result of pure mathematics.

Our thesis: A computational theory of games must be *generative*

There is *no equilibrium without a disequilibrium process.*

We need **a new solution concept**:

Applicable to any game and any dynamic!

That does not share the weaknesses of Nash equilibrium.

Effectively a solution concept that captures all dynamical systems. If such a result existed, it would be a fundamental result of pure mathematics.=>

Conley '78 Fundamental theorem of dynamical systems.

Crash Course in Dynamical Systems

Define (in the context of dynamical systems) and **motivate** (in the context of game theory) the following notions:

Lyapunov function:

Invariant function:

Chain equivalence/recurrence:

Fundamental Theorem of Dynamical Systems:

Dynamical System

Let (X,d) be a compact metric space X with metric d.
(e.g [0,1]^2, Euclidean distance)

Let T^+ denote either $Z^+ = \{0,1,2,...\}$ or $R^+ = [0,\infty)$.
A dynamical system on X is a continuous map $\phi : T^+ \times X \to X$
that satisfies the following two properties:
(i) $\phi(0,x) = x$
(ii) $\phi(s,\phi(t,x)) = \phi(s+t,x)\ \forall s,t \in T^+$

Typically, we study systems of
ODEs of the form:

$$\frac{dx}{dt} = f(x)$$

which under smoothness assumptions
on f admit a unique solution/flow.

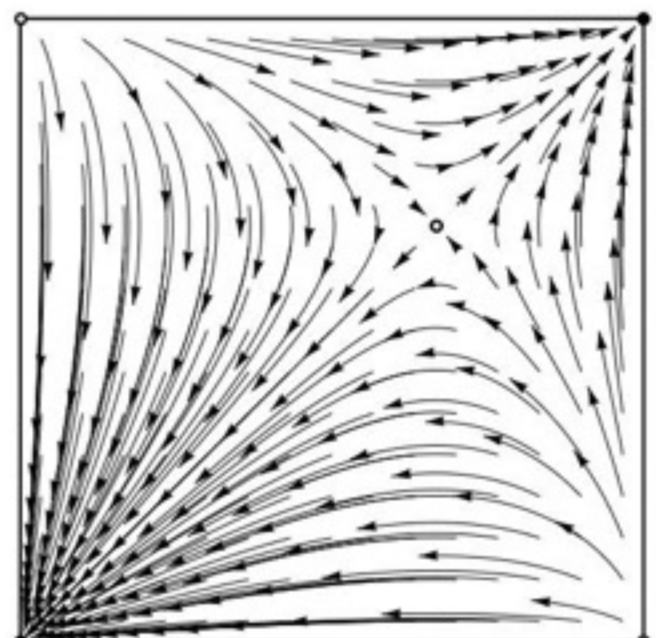

Lyapunov/potential function

A (strict) Lyapunov function of $\frac{dx}{dt} = f(x)$ is a scalar function L(x) which is (strictly) decreasing on non-equilibrium solutions.

Under strict Lyapunov fns, for each non-equilibrium solution x(t) L(x(t))<L(x(0)) for all t>0. I.e., whenever f(x(t)) $\neq \mathbf{0}$

$$\frac{d}{dt}L(x(t)) = \nabla L(x(t)) \cdot \frac{dx}{dt} = \nabla L(x(t)) \cdot f(x(t)) < 0$$

The ω-limit set of an orbit with initial condition *x* is defined as:

$y \in \omega(x) \leftrightarrow$ for some sequence $t_n \to \infty$, $x(t_k) \to y$.

Every ω(*x*) is a nonempty, compact, connected set consisting entirely of equilibria and upon which the Lyapunov function is constant.

First integral/Invariant function

A first integral/invariant of $\frac{dx}{dt} = f(x)$ is a scalar function I(x) which is constant on solutions.

In other words, for each solution x(t) to the differential equation, I(x(t)) = I(x(0)) for all t. I.e.,

$$0 = \frac{d}{dt}I(x(t)) = \nabla I(x(t)) \cdot \frac{dx}{dt} = \nabla I(x(t)) \cdot f(x(t))$$

Every solution to the dynamical system is constrained to move along a single level set {I(u) = c} of the first integral, namely the level set that contains the initial data.

(e.g., energy conservation in ideal pendulum)

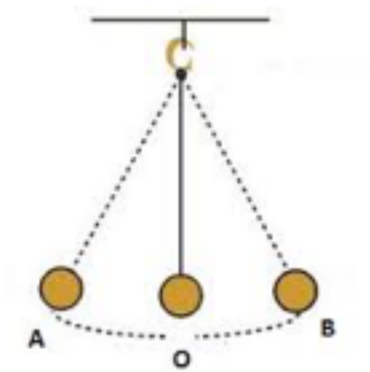

Chain recurrence

Periodicity for computer scientists

Suppose Alice wants to check that point x_0 is a periodic point of function g:X->X, however she can only check the accuracy of her computations up to accuracy $\varepsilon>0$.

Every time she computes another iteration $x_{i+1} = g(x_i)$, Bob can corrupt her computation by a small amount $|x_{i+1} - g(x_i)| < \varepsilon$.

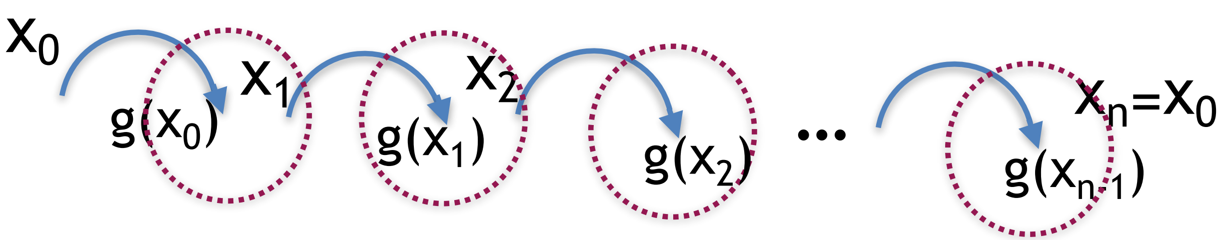

If Bob can always convince Alice that x_0 is periodic, no matter how small the allowable error is, then we say that x_0 is chain recurrent.

Chain recurrence (examples)

All equilibria of a dynamical system are chain recurrent.

All periodic points of a dynamical system are chain recurrent.

In fact, equilibria(f) $\subset$ periodic points(f) $\subset$ chain recurrent(f)

However, the set of chain recurrent points can be **strictly larger** than the set of periodic points.

E.g. Take a system defined by a point mass that traverses a circle of radius 1. On each discrete time step, the point mass takes makes a clockwise turn of 1 radians.

No configuration is periodic but all are chain recurrent.

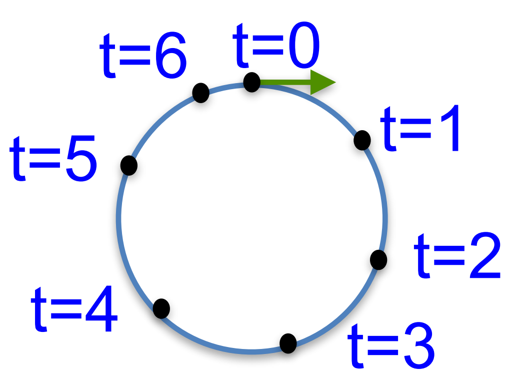

Fundamental theorem of dyn. systems

We know that in:

In *gradient(-like) dynamical systems*

There exists a strict Lyapunov function L such that

if x is not an equilibrium then

L(x(t)) will strictly decrease.

Conley (1976): Chain recurrence captures all limit behavior.

In **continuous dynamical systems**

There exists a function L such that

if x is not an **chain recurrence point** then

L(x(t)) will strictly decrease.

Recap

Define (in the context of dynamical systems) and **motivate** (in the context of game theory) the following notions:

Lyapunov function: decrease along trajectories

Invariant function: remain constant along trajectories

Chain equivalence/recurrence: lie on closed orbits modulo infinitesimal corrections.

Fundamental Theorem of Dynamical Systems: Chain recurrence captures all limit behavior.

Dynamical Systems in Games

Agent behavior is deterministic learning dynamic:
Input: Current mixed strategy of other agents
Output: Chosen (randomized) action

Game + learning dynamic -> dynamical system.

Even if we do not converge to equilibrium always, we can try to understand how the structure of the chain recurrence sets affects system performance.

How do CRS look like?

The answer **depends on the dynamic.**

How do chain recurrent sets (CRS) look for the most well studied combinations of games and dynamics.

Game class 1: Potential/Coordination games

Game “class” 2: 2x2 zero-sum games

These are easy instances. Hence, CRS=Equilibria, right?

How do CRS look like?

[Papadimitriou-P. ITCS ’16]

The answer **depends on the dynamic.**

How do chain recurrent sets (CRS) look for the most well studied combinations of games and dynamics.

Game class 1: Potential/Coordination games**: YES**

Game “class” 2: 2x2 zero-sum games: **NO, maximally different**

These are easy instances. Hence, CRS=Equilibria, right?

Replicator Dynamics

[Taylor-Jonker '78,Schuster-Sigmund '83]

We will assume that each agent updates their initial mixed strategy according to the replicator dynamics.

Basic tool in mathematical theory of selection and evolution.
Survival of the fittest: The probability of an action increases iff it outperforms the current (mixed) strategy.

$$\frac{dx_i(t)}{dt} = x_i[u_i(x) - \hat{u}(x)], \quad \hat{u}(x) = \sum_{i=1}^{n} x_i u_i(x)$$

rate of growth of action i

utility of action i at state x

expected utility at state x

Sanity check 1: Coordination games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

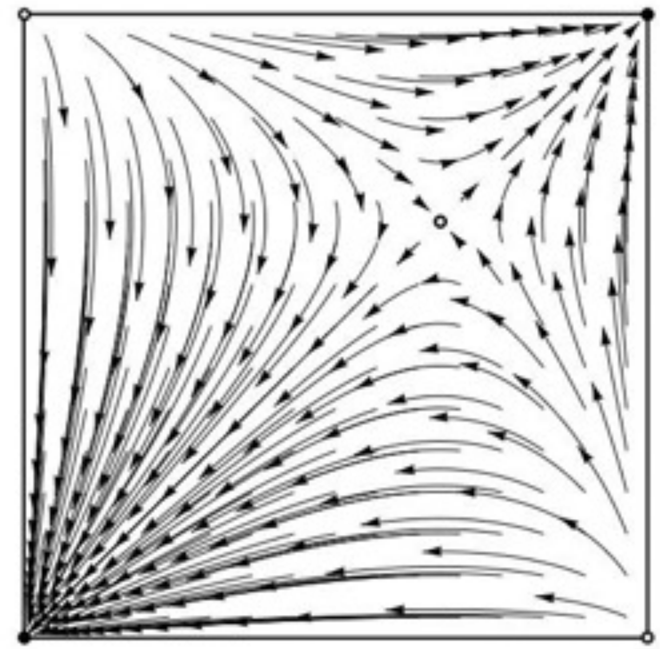

In potential games, for replicator dynamics (as well as many other dyn.) there exists a function L that strictly increases/decreases implying convergence to NE.

Does this imply that CRS=NE? Not automatically.

Why? (Counterexample by Conley).
One has to make sure that this small deviations cannot add up in unpredictable ways. (Sard's thm, smoothness, e.t.c.)

Sanity check 1: Coordination games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

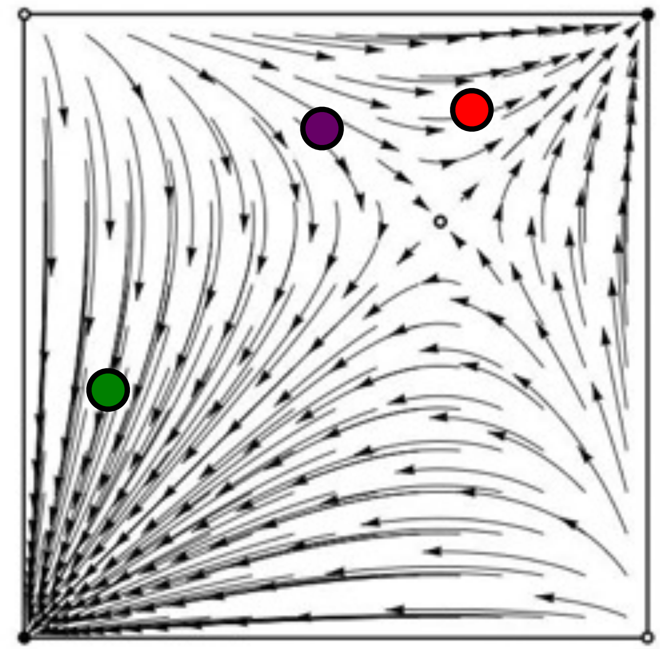

In potential games, for replicator dynamics (as well as many other dyn.) there exists a function L that strictly increases/decreases implying convergence to NE.

Does this imply that CRS=NE? Not automatically.

Why? (Counterexample by Conley).
One has to make sure that this small deviations cannot add up in unpredictable ways. (Sard's thm, smoothness, e.t.c.)

Sanity check 1: Coordination games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

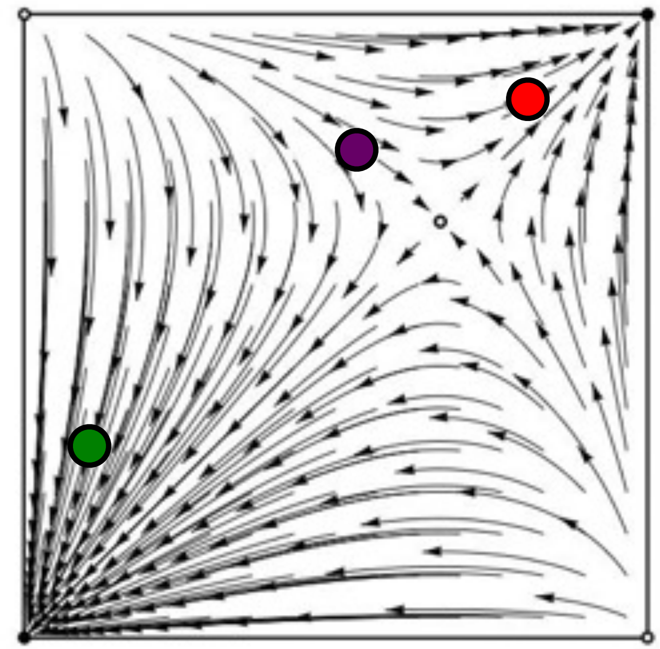

In potential games, for replicator dynamics (as well as many other dyn.) there exists a function L that strictly increases/decreases implying convergence to NE.

Does this imply that CRS=NE? Not automatically.

Why? (Counterexample by Conley).
One has to make sure that this small deviations cannot add up in unpredictable ways. (Sard's thm, smoothness, e.t.c.)

Sanity check 1: Coordination games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

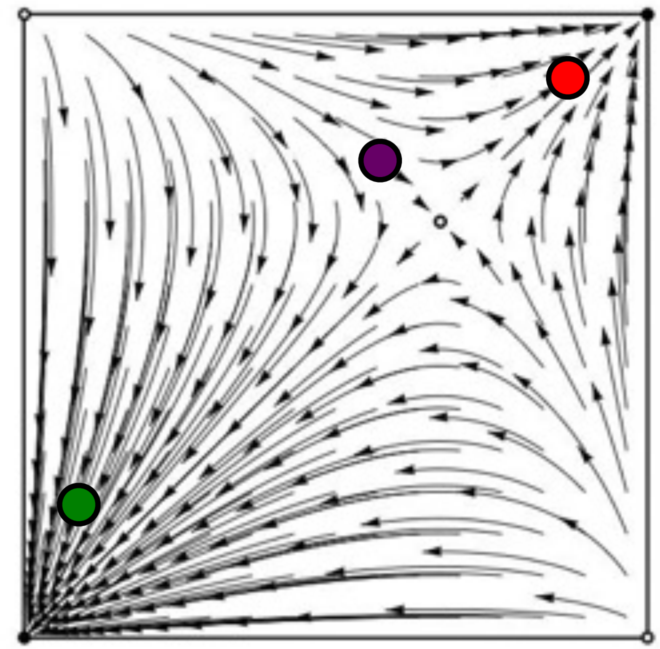

In potential games, for replicator dynamics (as well as many other dyn.) there exists a function L that strictly increases/decreases implying convergence to NE.

Does this imply that CRS=NE? Not automatically.

Why? (Counterexample by Conley).
One has to make sure that this small deviations cannot add up in unpredictable ways. (Sard's thm, smoothness, e.t.c.)

Sanity check 1: Coordination games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

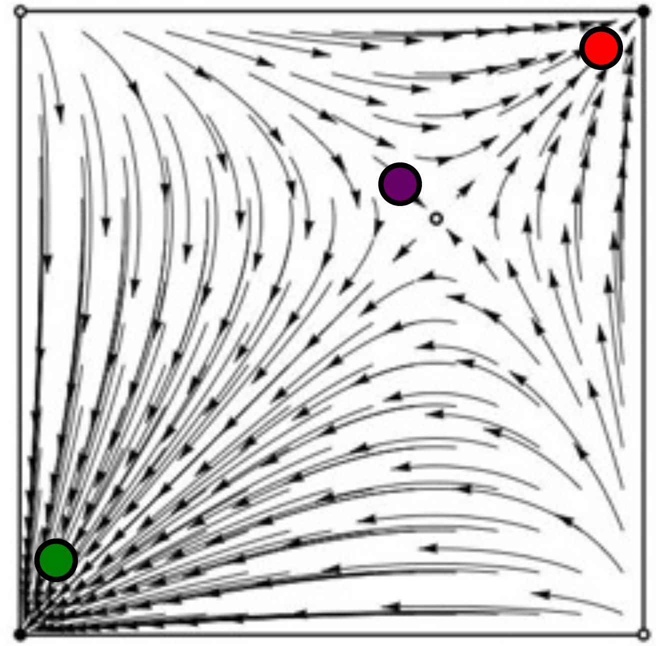

In potential games, for replicator dynamics (as well as many other dyn.) there exists a function L that strictly increases/decreases implying convergence to NE.

Does this imply that CRS=NE? Not automatically.

Why? (Counterexample by Conley).
One has to make sure that this small deviations cannot add up in unpredictable ways. (Sard's thm, smoothness, e.t.c.)

Sanity check 2: 2X2 ZS games!

	A	B
A	1,-1	-1,1
B	-1,1	1,-1

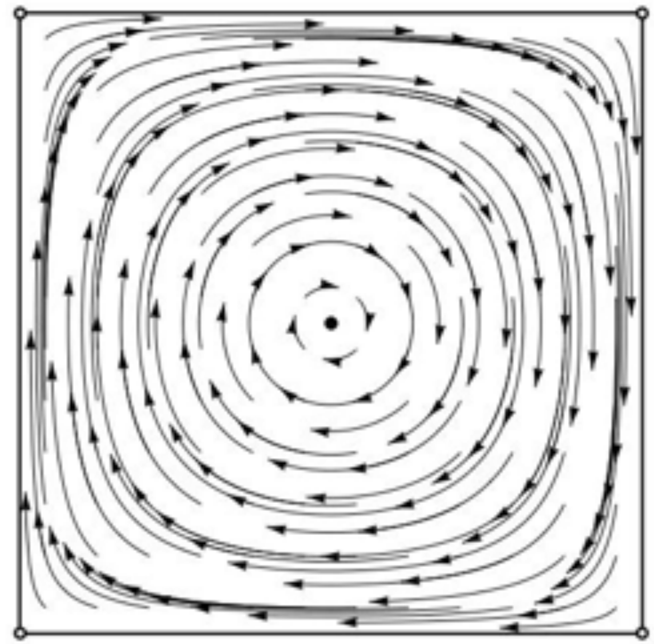

In ZS with fully mixed NE (e.g. Matching Pennies game), the set of CRS is the whole state space.

What does this mean? In the presence of infinitesimal small and infrequent noise absolutely no prediction can be made about the actual state of the system. (Effective chaos: Google vs. Microsoft, or Google vs. Facebook)

Equilibria: **Trivial to compute, unique, do not capture** behavior.

Invariants via Information Theory

Kullback-Leibler divergence is a non-symmetric measure of the difference between two probability distributions *p* and *q*. The K-L divergence of *q* from *p*, denoted $D_{\mathrm{KL}}(p||q)$ is a measure of the information lost when *q* is used to approximate *p*.

$$D_{\mathrm{KL}}(p\|q) = \sum_i \ln\left(\frac{p(i)}{q(i)}\right) p(i)$$

K-L divergence is a **pseudometric**:
Non-negative ($D_{\mathrm{KL}}(p||q)$ =0 iff p=q)
Well-defined (i.e., finite) if p,q have the same support.

In the case of zero-sum games with fully mixed NE we have that:

$$D_{\mathrm{KL}}(p^1_{\mathrm{NE}}\|x^1(t)) + D_{\mathrm{KL}}(p^2_{\mathrm{NE}}\|x^2(t)) = ct$$ [Hofbauer '98]

Similar conditions hold for arbitrary large networks.

Poincaré–Bendixson theorem (1901)

Two dimensional systems with real timeset cannot be chaotic.

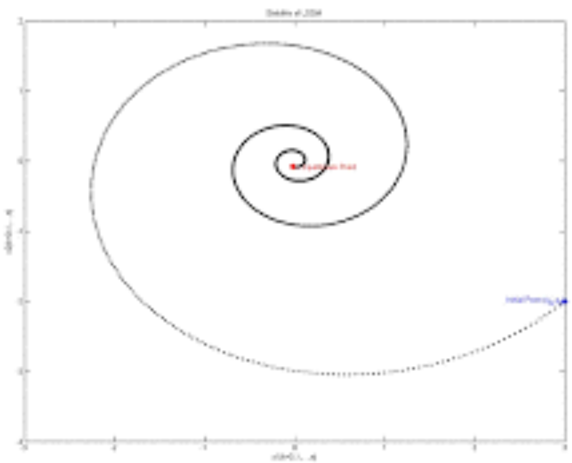

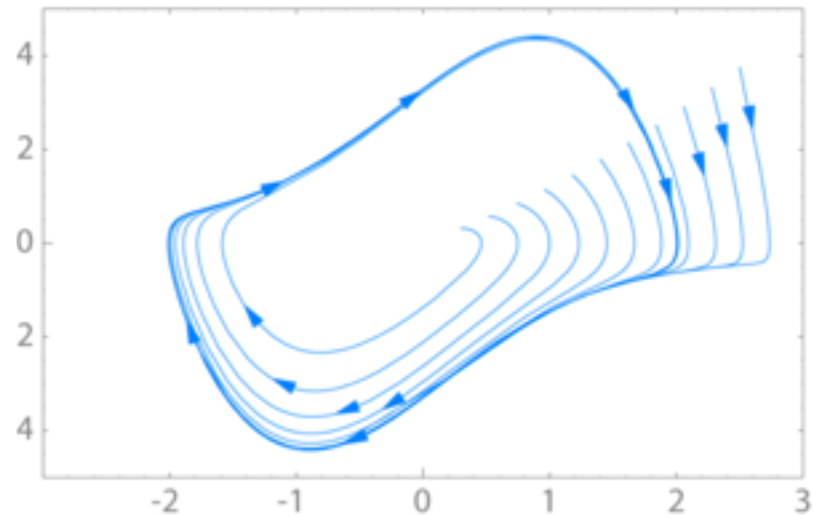

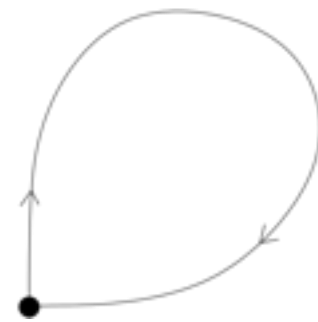

Given a differentiable real dynamical system defined on an open subset of the plane, then every non-empty compact ω-limit set of an orbit, which contains only finitely many fixed points, is either

- a fixed point,
- a periodic orbit, or
- a connected set composed of a finite number of fixed points together with homoclinic and heteroclinic orbits connecting these.

Analysis: 2X2 ZS games

	A	B
A	1,-1	-1,1
B	-1,1	1,-1

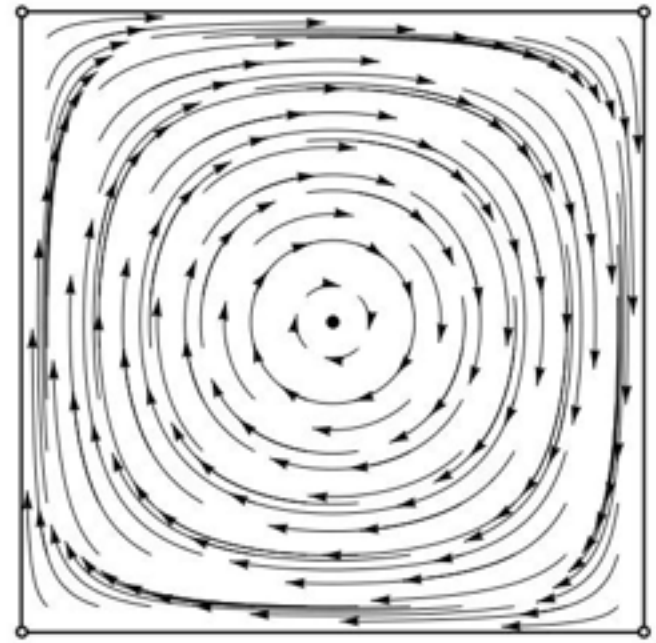

In ZS with fully mixed NE (e.g. Matching Pennies game), the set of CRS is the whole state space.

The existence of the specific invariants

$$D_{\mathrm{KL}}(p^1_{\mathrm{NE}} \| x^1(t)) + D_{\mathrm{KL}}(p^2_{\mathrm{NE}} \| x^2(t)) = ct$$

can be shown to exclude all other possibilities allowed by the Poincaré–Bendixson theorem expect for periodicity. Specifically, every initial condition lies on a cycle where its distance from its center in terms of KL-divergence remains constant.

Nash Equilibria: The Nash equilibrium of solution concepts

A ***lot of work*** has been invested in reconciling these problems to the extent possible:

Non-uniqueness: **Axiomatic refinements of Nash equilibria** (e.g. risk dominant, payoff dominant eq, ESS, SSS, proper, trembling hand, perfect, e.t.c.)

Computational hardness of computing a sample (even approximate) equilibrium in large games: Exploring the tractability of computing ϵ**-approximate equilibria**.

Learning dynamics typically do not converge to equilibria even in simple small games: Focusing on games where dynamics behave well (e.g. **potential, zero-sum games)**

Chain Recurrent Sets > Nash eq.
But which CRS?

A ***lot of work*** has been invested in reconciling these problems to the extent possible:

Non-uniqueness: **Still a problem for CRS.**

Computational hardness of computing a sample (even approximate) CRS in k-person games: Trivially in **P** for replicator dynamics.

Learning dynamics typically **do converge** to CRS even in all games: By definition.

Revisiting Coordination Games

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

Multiple Nash/Chain Recurrent Points: Which one to choose?
``*No other task may be more significant within game theory.*''
Ariel Rubinstein

Worst Case Analysis (a.k.a. **Price of Anarchy**)

$$\textbf{PoA} = \frac{\textbf{Social Welfare (OPT)}}{\textbf{Social Welfare (worst equilibrium)}} \geq 1$$

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

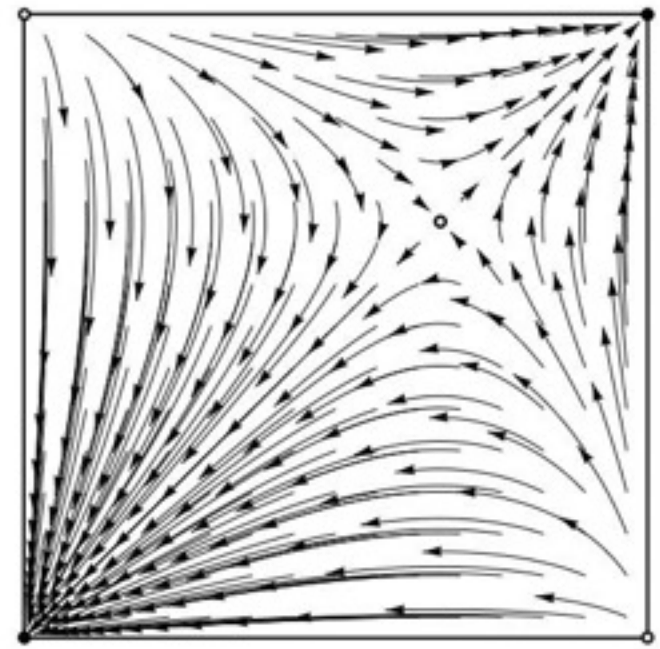

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

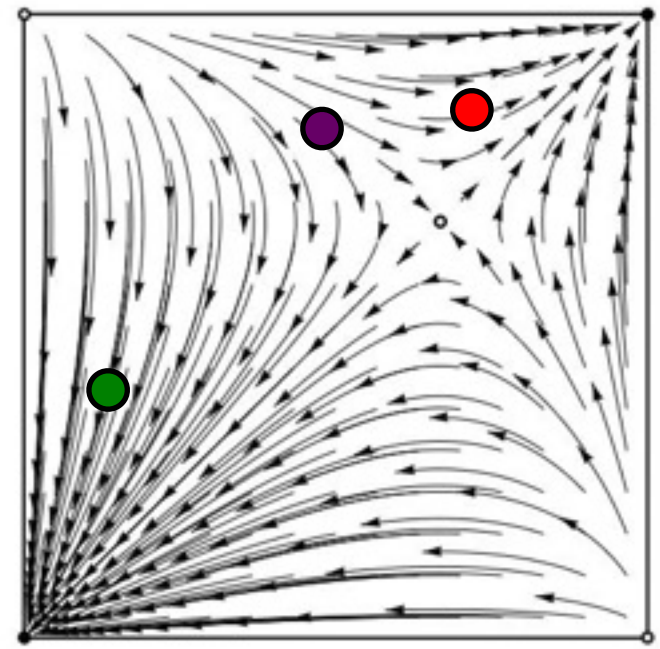

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

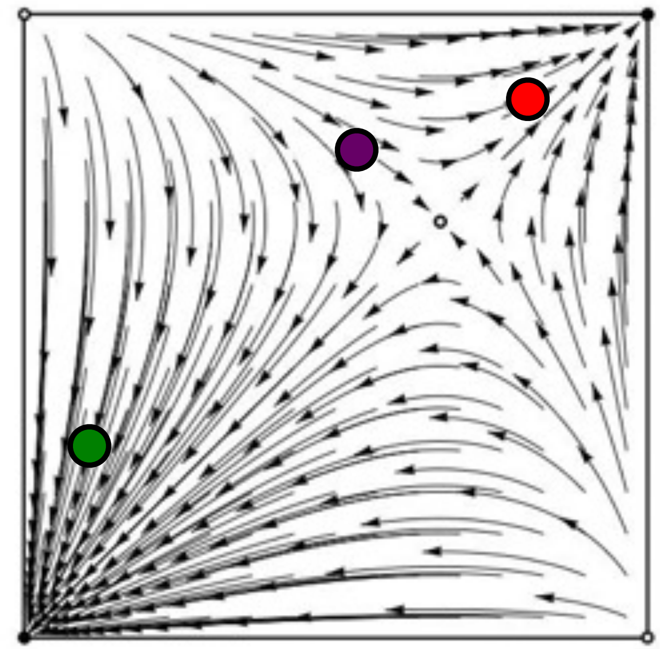

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

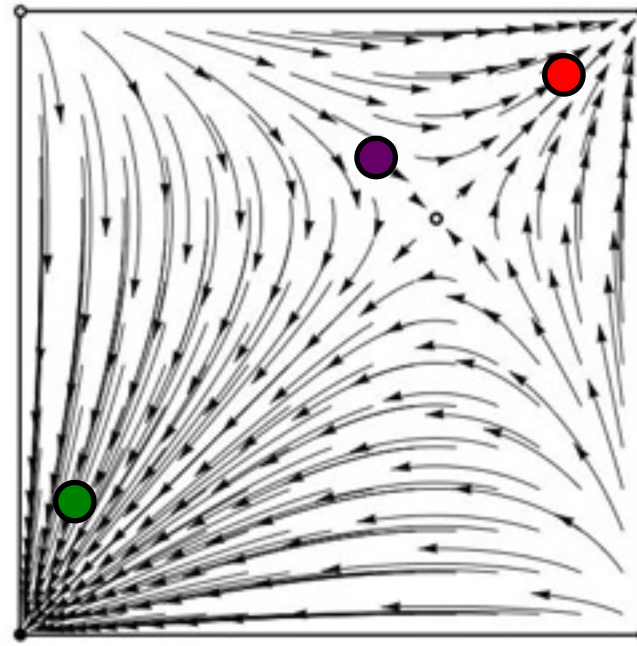

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

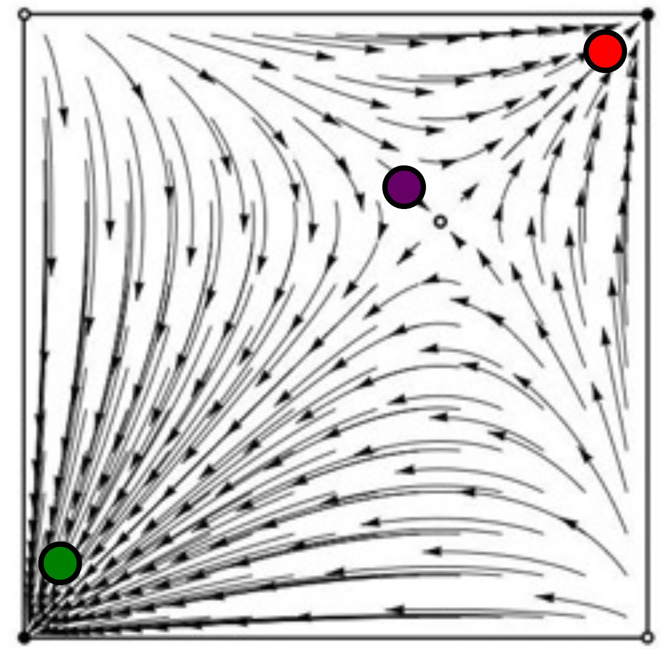

Prediction in Dynamical Systems

How will the system behave?
It depends on the initial conditions.

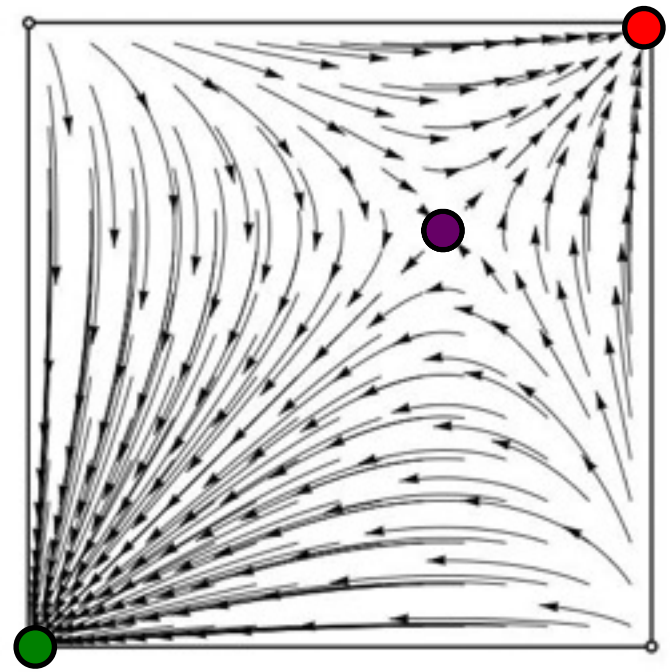

Prediction in Dynamical Systems

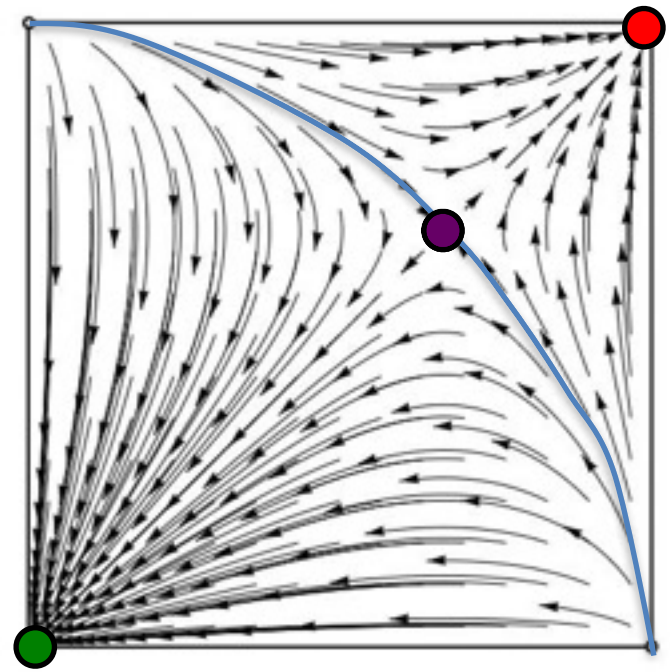

How will the system behave?
It depends on the initial conditions.
This specific dynamical system has **two good** properties:

A) Every trajectory converges to a fixed point.

B) We can predict long term behavior given any initial condition if we test a simple oracle (e.g. is the initial condition above/below the blue curve).

Average case analysis: Given an initial distribution over initial conditions (typically uniform), output the resulting distribution over equilibria (proportional to its region of attraction).

How well does the system behave on average?

	A	B
A	1, 1	0, 0
B	0, 0	2, 2

Convergence to equilibria follows from standard Lyapunov args.

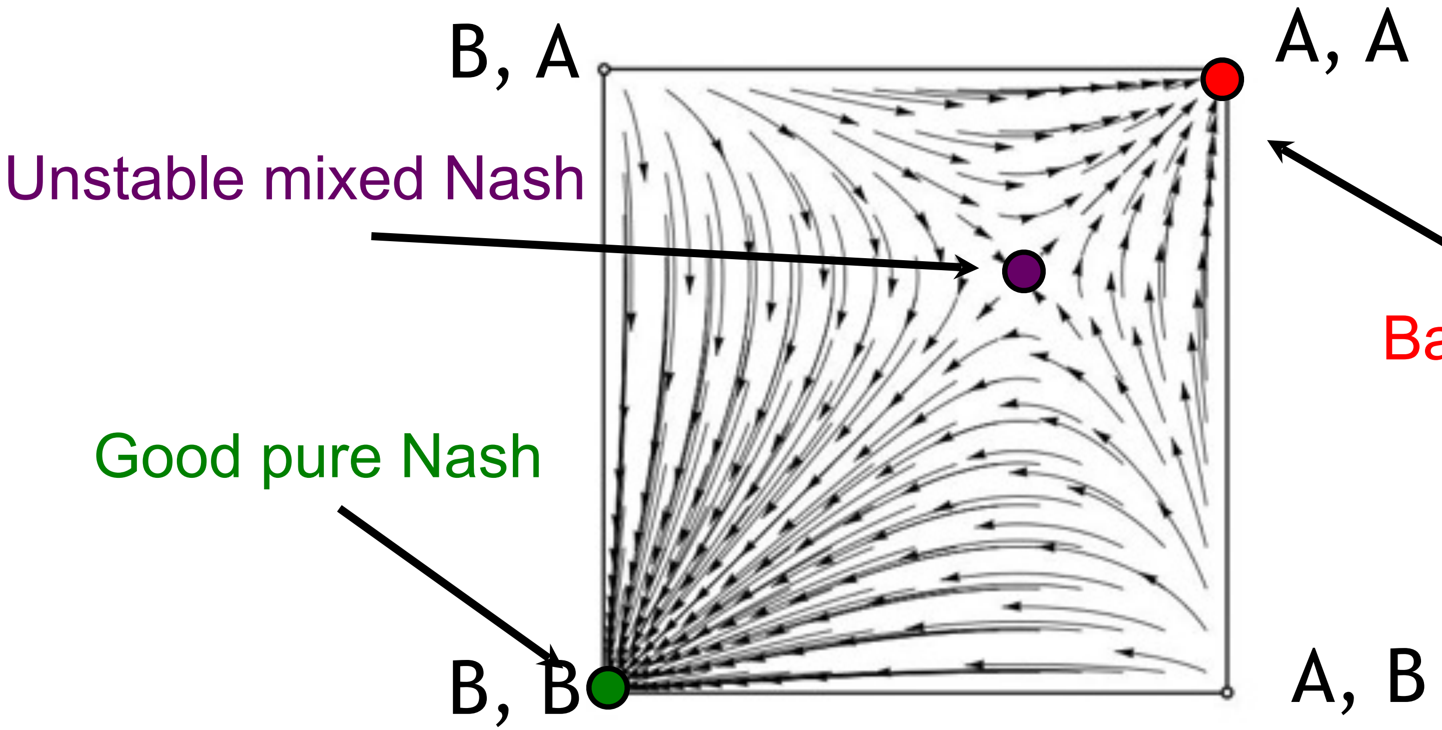

How well does the system behave on average?

	A	B
A	1, 1	0, 0
B	0, 0	3, 3

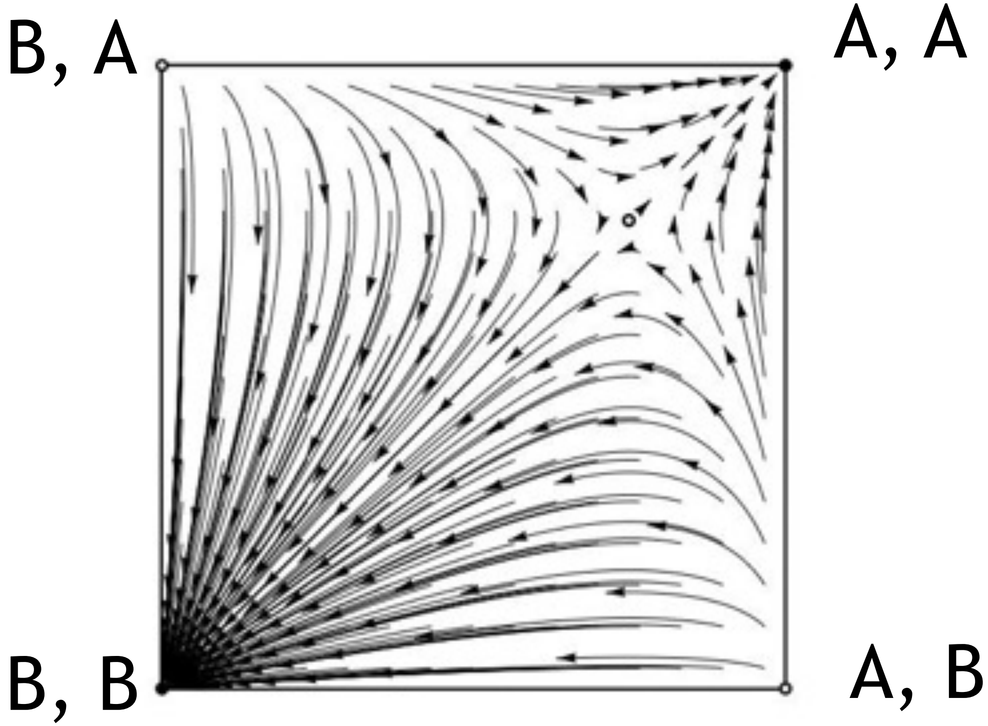

How well does the system behave on average?

	A	B
A	1, 1	0, 0
B	0, 0	4, 4

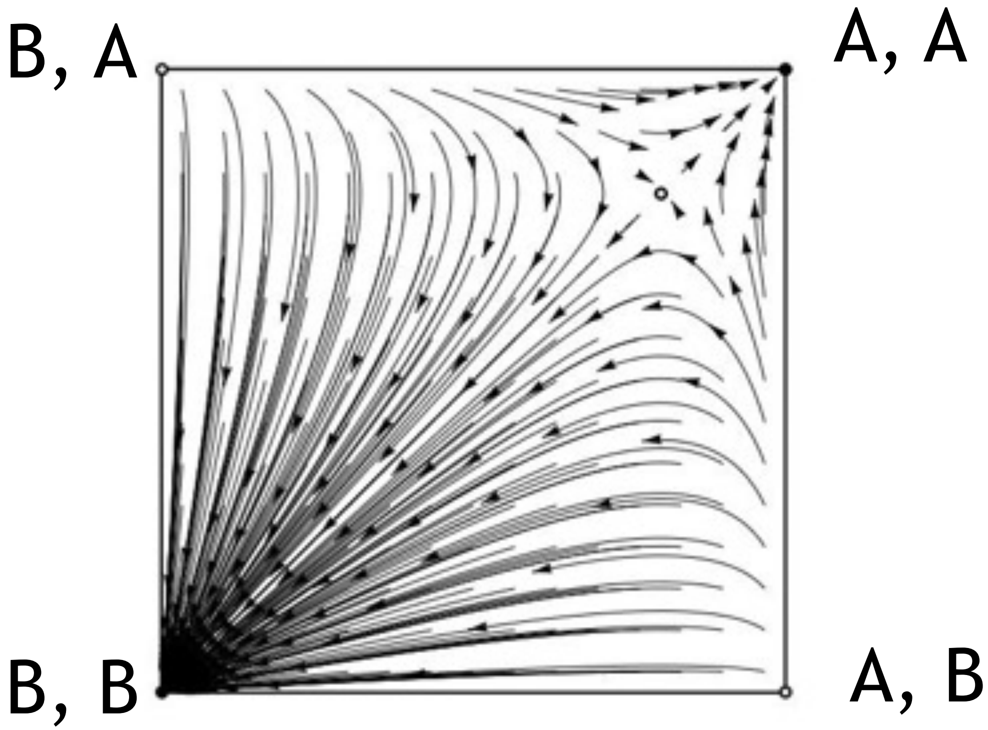

How well does the system behave on average?

	A	B
A	1, 1	0, 0
B	0, 0	10, 10

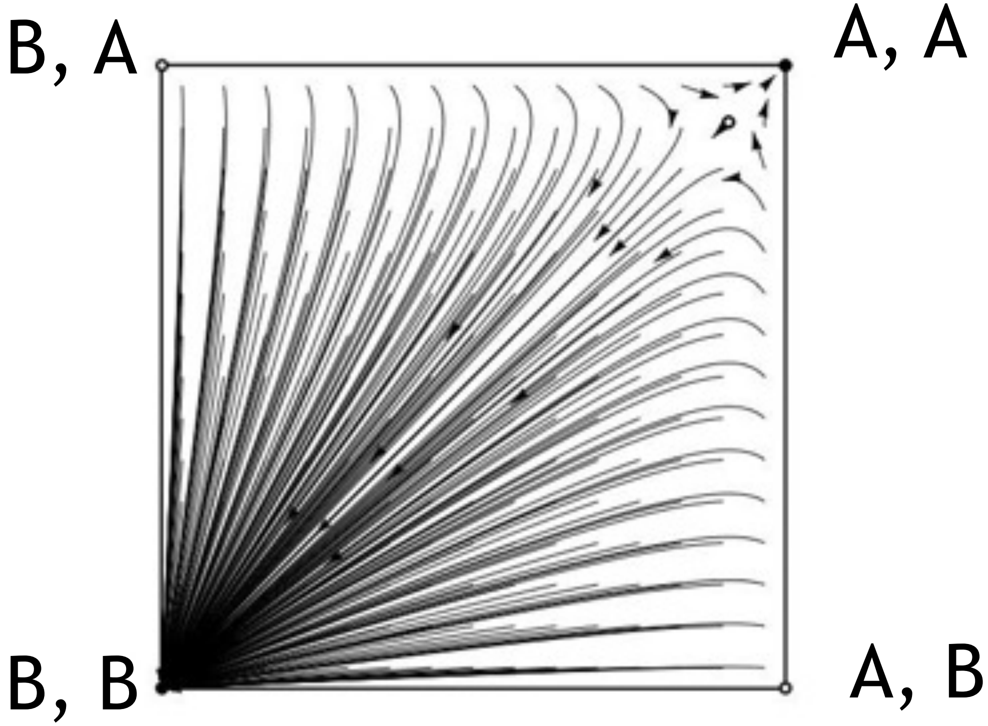

Average Case Performance

	A	B
A	1, 1	0, 0
B	0, 0	w, w

Theorem: [Panageas-P., EC'16]

1. In the case of 2x2 coordination games with w=2, the size of the region of attraction of the optimal equilibrium (B,B) is equal to $(9+2\pi\sqrt{3})/27 \approx 0.7364$.
2. For w≥1, the size of attraction of the optimal equilibrium is at least w/(w+1) and at most $1-2/(w+1)^2$.
3. For w≥1, the average system performance, i.e. where the quality of each equilibrium is weighted by the size of its basin of attraction, is at least 1.653*w. (close to the optimal = 2w, much better than the worst pure Nash = 2)

Invariants via Information Theory

Kullback-Leibler divergence is a non-symmetric measure of the difference between two probability distributions *p* and *q*. The K-L divergence of *q* from *p*, denoted $D_{KL}(p||q)$ is a measure of the information lost when *q* is used to approximate *p*.

$$D_{\mathrm{KL}}(p\|q) = \sum_i \ln\left(\frac{p(i)}{q(i)}\right) p(i)$$

K-L divergence is a **pseudometric**:
Non-negative ($D_{KL}(p||q)$ =0 iff p=q)
Well-defined (i.e., finite) if p,q have the same support.

In the case of coordination games with fully mixed NE we have that: (analogous but more complex relations hold for)

$$D_{\mathrm{KL}}(p^1_{\mathrm{NE}}\|x^1(t)) - D_{\mathrm{KL}}(p^2_{\mathrm{NE}}\|x^2(t)) = ct$$ [Hofbauer '98]

Similar conditions hold for arbitrary large networks.

Computing stable/unstable manifolds

The KL-divergence invariant for the stable/unstable manifold must be equivalent to:

$x^w(1-x)=y^w(1-y)$ where x,y the probabilities that the first, second agent assign to action A.

Indeed, the unstable manifold y=x satisfies this equation.

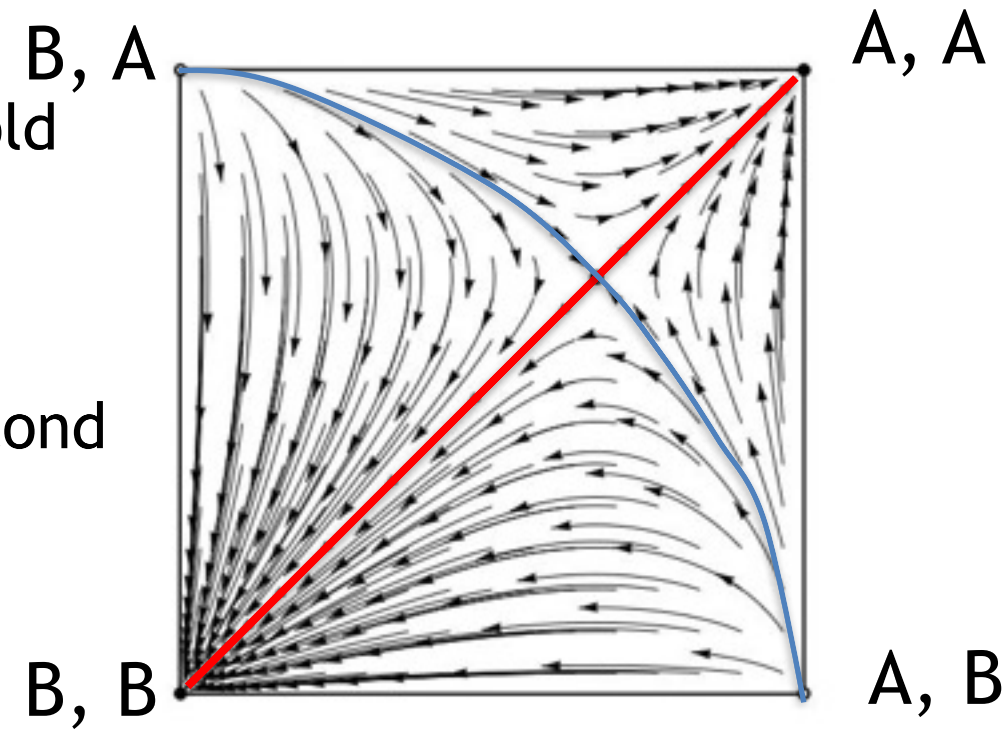

For fixed values of w, e.g. w=2, we can explicitly solve this equation.
The solution for w=2 is $y = (1 - x + \sqrt{1 + 2x - 3x^2})/2$.
We can integrate this function to compute the exact size of the region of attraction for the good eq. (B,B) and hence for (A,A) as well.

Approximating the basin of attraction of pure eq.

The KL-divergence invariant for the stable/unstable manifold must be equivalent to:

$x^w(1-x)=y^w(1-y)$ where x,y the probabilities that the first, second agent assign to action A.

The parametric uncertainty in the exponent does not allow for a general closed solution.

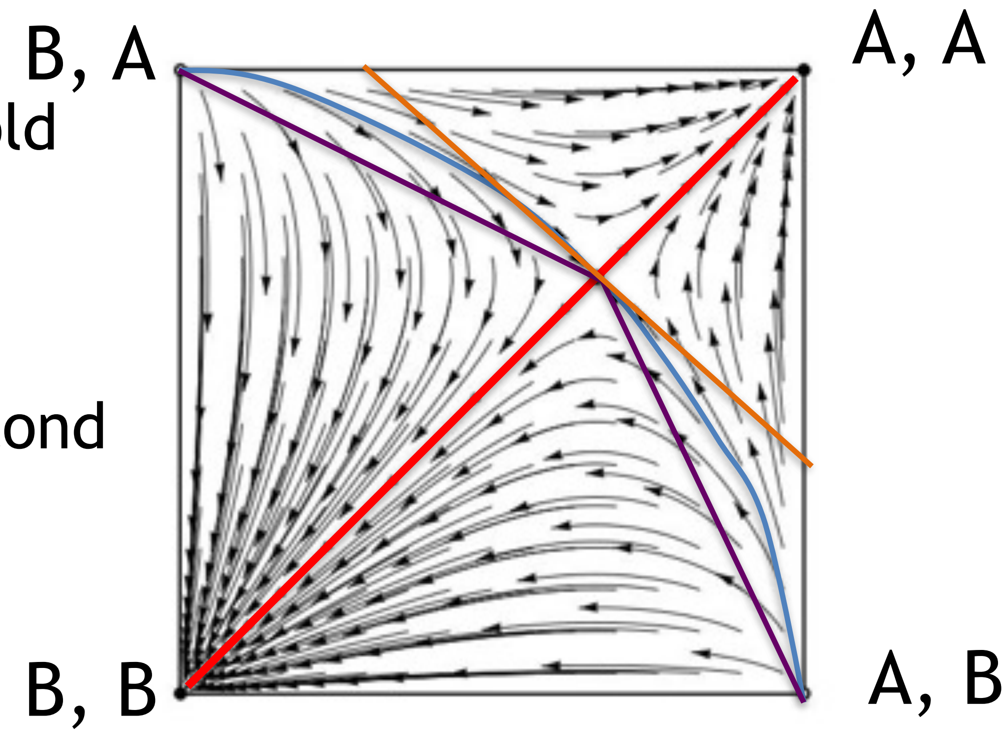

Instead, we produce approximate coverings of the regions of attraction using unions of parametric polytopes. This allows us to compute upper and lower bounds on the size of each basin.

Extensions to more agents

(1,1)	(0,0)
(0,0)	(2,2)

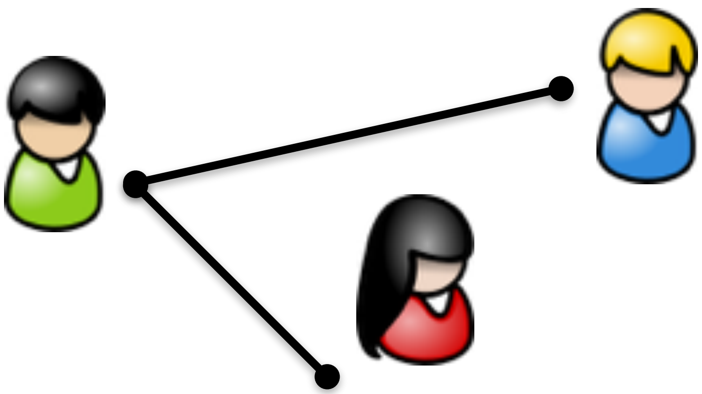

(1,1)	(0,0)
(0,0)	(2,2)

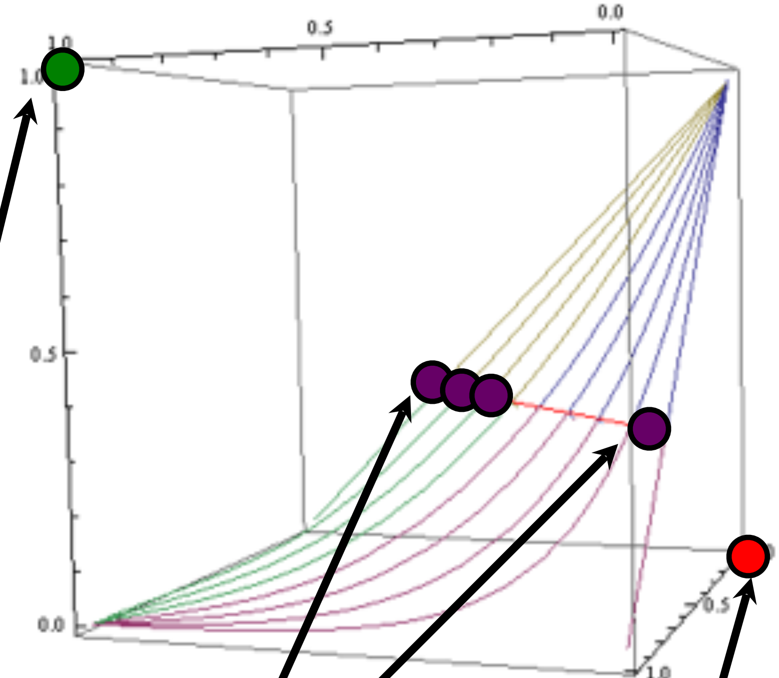

Good pure Nash

Continuum of
Unstable mixed Nash

Bad pure Nash

More agents->more invariants

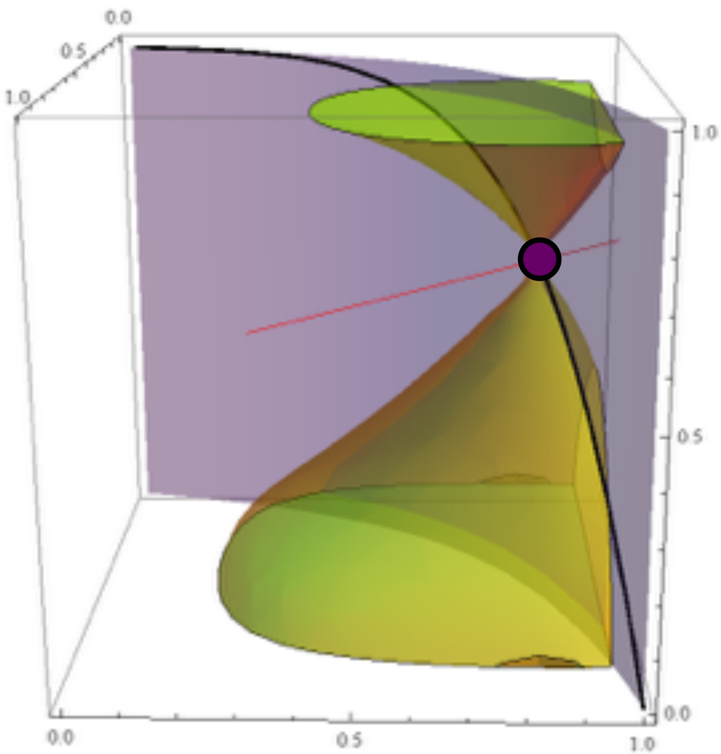

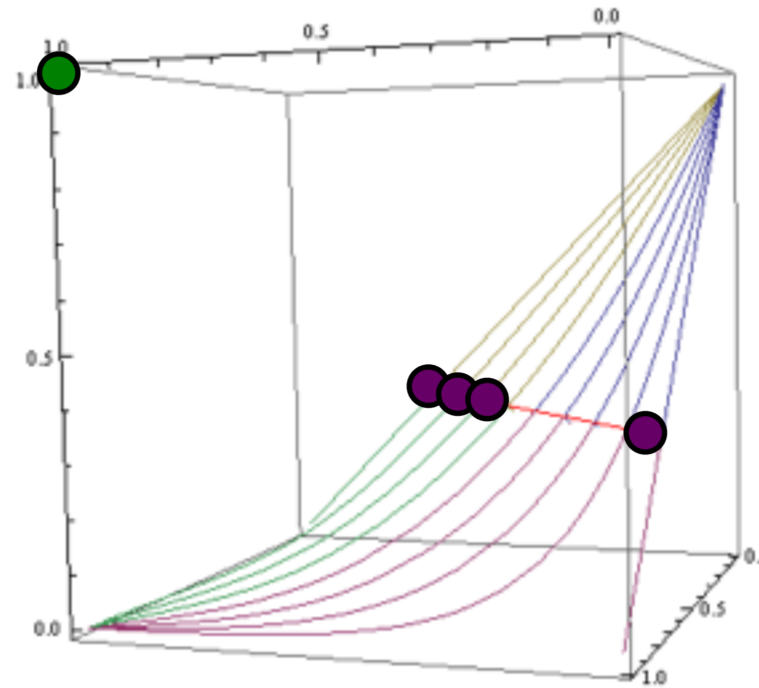

Suppose that we can find two invariants functions I_1,I_2 then given an unstable equilibrium x_0 all points of y on its stable manifold satisfy both:

$I_1(y)=I_1(x_0)$ Yellow set & $I_2(y)= I_2(x_0)$ Purple set

Their intersection identifies the stable manifold.

Open Questions

Chain recurrent sets is the most recent **mathematical theory of studying many-body interactions** and it is **totally unexplored**.

1. The ``combinatorial'' structure of CRS
2. Average Price of Anarchy beyond Potential Games
3. Understanding the mixing properties inside a CRS-component
4. Computational complexity considerations.
5. …

market

1920s theory, Equilibria

Does not scale, unless bad approx.

1980s model, CRS

Thank you