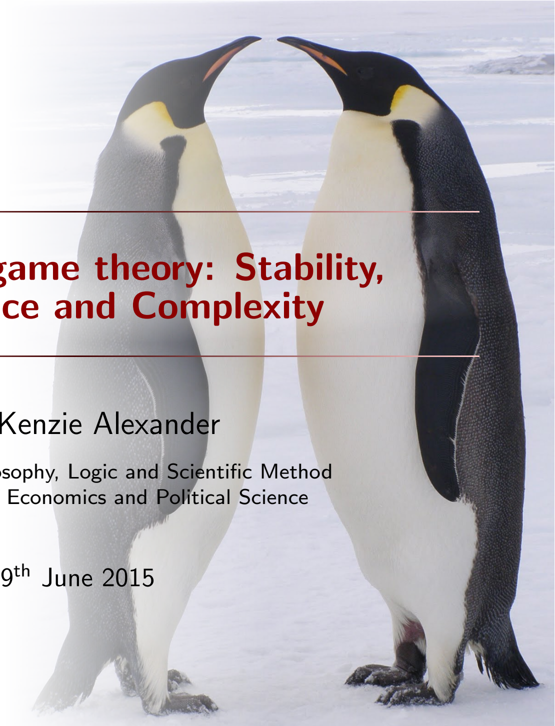

Evolutionary game theory: Stability, Convergence and Complexity

J. McKenzie Alexander

Department of Philosophy, Logic and Scientific Method
London School of Economics and Political Science

19th June 2015

Outline

1. Evolutionarily stable strategies
   - Basic concepts
   - Alternative characterisations of ESS
2. Dynamics and convergence
3. The complexity of ESS

Outline

1. Evolutionarily stable strategies
   - Basic concepts
   - Alternative characterisations of ESS
2. Dynamics and convergence
3. The complexity of ESS

Outline

1. Evolutionarily stable strategies
   - Basic concepts
   - Alternative characterisations of ESS
2. Dynamics and convergence
3. The complexity of ESS

Basic concepts

When talking about evolutionary games, we will almost always work with two-player symmetric games.

		Column	
		S_1	S_2
Row	S_1	(a, a)	(b, c)
	S_2	(c, b)	(d, d)

Symmetric

		Maggie	
		Boxing	Ballet
Billy	Boxing	$(1, 3)$	$(0, 0)$
	Ballet	$(0, 0)$	$(3, 1)$

Not symmetric

Let $S = \{ s_1, \ldots, s_n \}$ denote the set of pure strategies available to all players.

Let $\pi(s_i | s_j)$ denote the payoff received by a player who uses strategy s_i against someone using strategy s_j.

Basic concepts

If $\sigma = p_1 s_1 + \cdots + p_n s_n$ and $\mu = q_1 s_1 + \cdots + q_n s_n$ are *mixed strategies*, then

$$\begin{aligned}\pi(\sigma|\mu) &= \sum_{i=1}^{n} \sum_{j=1}^{n} p_i q_j \pi(s_i|s_j) \\ &= p_1 \pi(s_1|\mu) + p_2 \pi(s_2|\mu) + \cdots + s_n \pi(s_n|\mu) \\ &= q_1 \pi(\sigma|s_1) + q_2 \pi(\sigma|s_2) + \cdots + q_n \pi(\sigma|s_n).\end{aligned}$$

Evolutionarily stable strategies I

The central solution concept for traditional game theory is the following:

Nash equilibrium

$\sigma^* = (\sigma_1^*, \ldots, \sigma_n^*)$ is a *Nash equilibrium* iff, for all players i,

$$\pi_i\left(\sigma_i^*, \sigma_{-i}^*\right) \geq \pi_i\left(\sigma_i, \sigma_{-i}^*\right), \text{ for all } \sigma_i \in S_i.$$

A Nash equilibrium is a set of mutual best responses with respect to unilateral deviations. That is, no individual can gain by switching.

Evolutionarily stable strategies II

Intuitively, an *evolutionarily stable strategy* is one with the property that, if everyone in the population follows it, no other strategy (known as a *mutant*) can invade.

To see why the Nash equilibrium concept does not suffice for analysing evolutionary stability, consider the following game:

	A	B
A	$(2,2)$	$(1,1)$
B	$(1,1)$	$(1,1)$

A Nash equilibrium does not exclude the possibility of drift!

Evolutionarily stable strategies III

There are two logically equivalent ways of writing the definition:

A strategy s^* is *evolutionarily stable* iff for all strategies $s \neq s^*$

First form (Maynard Smith and Price, 1973)

either $\pi(s^*|s^*) > \pi(s|s^*)$

or $\pi(s^*|s^*) = \pi(s|s^*)$ and $\pi(s^*|s) > \pi(s|s)$.

Second form

1. $\pi(s^*|s^*) \geq \pi(s|s^*)$
2. If $\pi(s^*|s^*) = \pi(s|s^*)$ then $\pi(s^*|s) > \pi(s|s)$.

An alternative definition of ESS

Another common definition is as follows:

Alternate definition

s^* is an evolutionarily stable strategy iff for every strategy $s \neq s^*$, there exists an $\bar{\epsilon}_s \in (0,1)$ such that for all $0 < \epsilon < \bar{\epsilon}_s$,

$$\pi(s^*|\epsilon s + (1-\epsilon)s^*) > \pi(s|\epsilon s + (1-\epsilon)s^*)$$

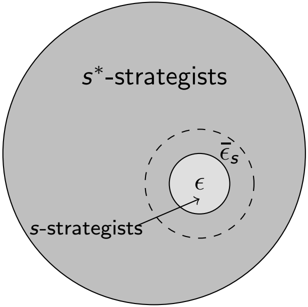

The population

Where $\epsilon s + (1-\epsilon)s^*$ denotes the strategy which, ϵ of the time, plays s and $1-\epsilon$ of the time, plays s^*.

Equivalence of the two definitions

Alternate definition

s^* is an evolutionarily stable strategy if for every strategy $s \neq s^*$, there exists an $\bar{\epsilon}_s \in (0,1)$ such that for all $0 < \epsilon < \bar{\epsilon}_s$,

$$\pi(s^*|\epsilon s + (1-\epsilon)s^*) > \pi(s|\epsilon s + (1-\epsilon)s^*)$$

Rewrite the inequality as:

$$(1-\epsilon)\Big(\pi(s^*|s^*) - \pi(s|s^*)\Big) + \epsilon\Big(\pi(s^*|s) - \pi(s|s)\Big) > 0.$$

As $\epsilon \to 0$,

$$\begin{aligned} \text{either } & \pi(s^*|s^*) - \pi(s|s^*) > 0 \\ \text{or } & \pi(s^*|s^*) = \pi(s|s^*) \text{ and } \pi(s^*|s) - \pi(s|s) > 0. \end{aligned}$$

An ESS is strictly stronger than a NE, I

Consider the following game:

	A	B
A	$(1, 1)$	$(100, 0)$
B	$(0, 100)$	$(100, 100)$

The two pure strategy Nash equilibria are (A, A) and (B, B).

Definition (second form)

1. $\pi(B|B) \geq \pi(A|B)$ ✓
2. If $\pi(B|B) = \pi(A|B)$ then $\pi(B|A) > \pi(A|A)$ ✗

An ESS is strictly stronger than a NE, II

A strategy A is said to *weakly dominate* a strategy B if it is A sometimes better to play A than B, and never worse.

Theorem

No weakly dominated strategy is an ESS.

We shall return to this point when we consider dynamics and convergence.

Is an ESS unique?

Consider the following game:

	A	B	C
A	$(3,3)$	$(0,0)$	$(0,0)$
B	$(0,0)$	$(1,1)$	$(1,2)$
C	$(0,0)$	$(2,1)$	$(1,1)$

$$\left.\begin{array}{l}\pi(A|A) > \pi(B|A)\\ \pi(A|A) > \pi(C|A)\end{array}\right\} \implies A \text{ is an ESS.}$$

$$\left.\begin{array}{l}\pi(C|C) > \pi(A|C)\\ \pi(C|C) = \pi(B|C) \text{ and } \pi(C|B) > \pi(B|B)\end{array}\right\} \implies C \text{ is an ESS.}$$

Does an ESS always exist?

Consider the following game:

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

The only Nash equilibrium is the one in mixed strategies: $\sigma = \left\langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right\rangle$.

Since $\pi(\text{Rock}|\sigma) = \pi(\sigma|\sigma)$ and $\pi(\sigma|\text{Rock}) = \pi(\text{Rock}|\text{Rock})$, there is no ESS for this game.

How many ESS can there be?

We can answer this via the following theorem.

Definition

Let σ be a mixed strategy. The *support* of σ, denoted $\mathrm{supp}(\sigma)$, is the set of all pure strategies played by σ with positive probability.

Theorem

Suppose that σ is an ESS. If μ is a strategy and $\mathrm{supp}(\mu) \subset \mathrm{supp}(\sigma)$, then μ is not an ESS. If μ is an ESS with $\mathrm{supp}(\mu) = \mathrm{supp}(\sigma)$, then $\mu = \sigma$.

How many ESS can there be? II

Since we are considering games with a finite number of strategies, the number of possible support sets is finite.

Theorem

The number of ESS is finite (and possibly zero).

Theorem

A completely mixed ESS is the only ESS of the game.

Outline

1. Evolutionarily stable strategies
    - Basic concepts
    - Alternative characterisations of ESS
2. Dynamics and convergence
3. The complexity of ESS

Uniform invasion barriers

Definition

A strategy σ is said to have a *uniform invasion barrier* if there exists an $\bar{\epsilon} > 0$ such that, for all $\mu \neq \sigma$ and every $\epsilon < \bar{\epsilon}$,

$$\pi(\sigma|\epsilon\mu + (1-\epsilon)\sigma) > \pi(\mu|\epsilon\mu + (1-\epsilon)\sigma)$$

Theorem

A strategy σ is an ESS if and only if σ has a uniform invasion barrier.

That is, an ESS can drive out any other type in the population, provided that the frequency of the other type is below this threshold.

Local superiority

Definition

A strategy σ is said to be *locally superior* if it has a neighbourhood U such that $\pi(\sigma|\mu) > \pi(\mu|\mu)$ for all $\mu \in U$ where $\mu \neq \sigma$.

Theorem

A strategy σ is an ESS if and only if σ is locally superior.

Refinements of Nash equilibrium

Extensive-form refinements:

- Subgame-perfect equilibrium
- (Weak) Perfect Bayesian equilibrium
- Sequential equilibrium
- Trembling-hand perfect equilibrium
- Proper equilibrium

Strategic-form refinements:

- Iterative elimination of dominated strategies
- Perfect equilibrium
- Proper equilibrium

Refinements of Evolutionary Stability

Consider the following game:

	S_1	S_2	S_3
S_1	1	2	2
S_2	3	0	4
S_3	2	1	3

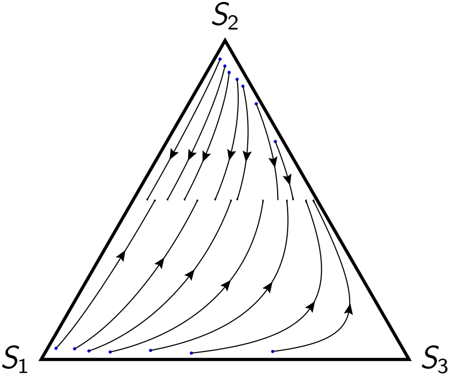

The points on the line connecting $(0, \frac{1}{2}, \frac{1}{2})$ to $(\frac{1}{2}, \frac{1}{2}, 0)$ are all Nash equilibria, but not ESS.

Refinements of Evolutionary Stability

Weaker than ESS:

- Neutrally stable strategies
- Robustness against equilibrium entrants

Set-based criteria:

- Evolutionarily stable sets
- Equilibrium evolutionarily stable sets

Outline

1. Evolutionarily stable strategies
   - Basic concepts
   - Alternative characterisations of ESS
2. **Dynamics and convergence**
3. The complexity of ESS

Dynamics for evolutionary games

All the solution concepts identified so far:

- Evolutionarily stable strategies
- Evolutionarily stable sets
- etc.

are *static*.

Shortly after the development of evolutionarily game theory, Taylor and Jonker introduced the *replicator dynamics* to model strategic evolution in populations.

$$\frac{dx_i}{dt} = x_i\left(W_i - \overline{W}\right).$$

Revisiting the refinements problem

There are many possible evolutionary dynamics to choose from, and many ways of generating those dynamics. (See Sandholm (2010) for a comprehensive discussion).

Let $S = \{ 1, \ldots, n \}$ denote the set of pure strategies, π the payoff function and $x = (x_1, \ldots, x_n)$ the state.

Let $\rho_{ij}(\pi, x)$ denote the conditional switch rate.

In general, $\dot{x}_i = \sum_{j \in S} x_j \rho_{ji} - x_i \sum_{j \in S} \rho_{ij}$.

Any specification of ρ_{ij} yields an evolutionary dynamic.

Revisiting the refinements problem

For example, the following all yield the replicator dynamics:

$\rho_{ij}(\pi, x) = x_j(K - \pi_i)$ Dissatisfaction-driven imitation

$\rho_{ij}(\pi, x) = x_j(\pi_j - K)$ Imitation of success

$\rho_{ij}(\pi, x) = x_j\left[\pi_j - \pi_i\right]_+$ Pairwise proportional imitation

The constant K is selected so as to ensure $\rho_{ij} \geq 0$.

There are many other dynamics besides the replicator dynamics.

Revisiting the refinements problem

These dynamics all behave differently, at least with respect to the short- and medium- term population behaviour.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

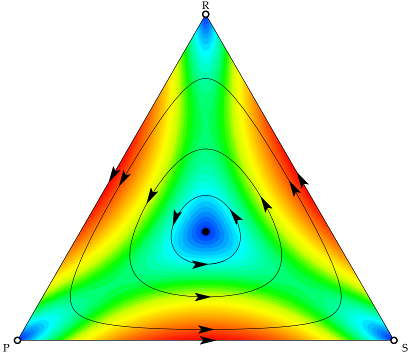

(i) replicator

But do they make an *important* difference?

Revisiting the refinements problem

These dynamics all behave differently, at least with respect to the short- and medium- term population behaviour.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

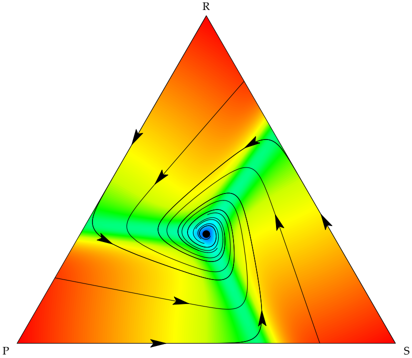

(ii) logit(.08)

But do they make an *important* difference?

Revisiting the refinements problem

These dynamics all behave differently, at least with respect to the short- and medium- term population behaviour.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

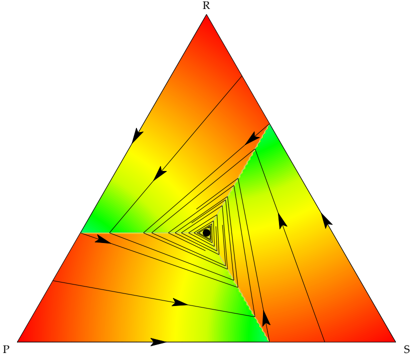

(iii) best response

But do they make an *important* difference?

Revisiting the refinements problem

These dynamics all behave differently, at least with respect to the short- and medium- term population behaviour.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

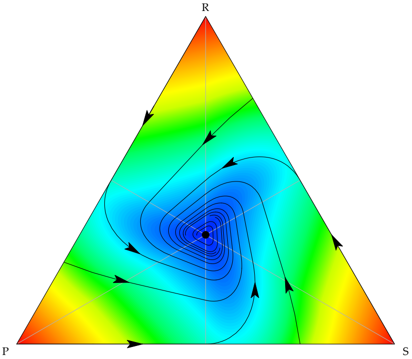

(iv) BNN

But do they make an *important* difference?

Revisiting the refinements problem

These dynamics all behave differently, at least with respect to the short- and medium- term population behaviour.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,1)$	$(1,-1)$
Paper	$(1,-1)$	$(0,0)$	$(-1,1)$
Scissors	$(-1,1)$	$(1,-1)$	$(0,0)$

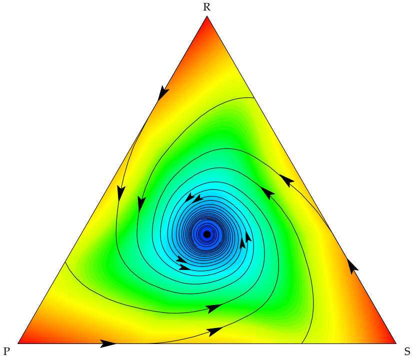

(v) Smith

But do they make an *important* difference?

Preservation of dominated strategies

The replicator dynamics does not eliminate weakly dominated strategies (so it can yield outcomes which are not ESS). It does, however, eliminate strongly dominated strategies.

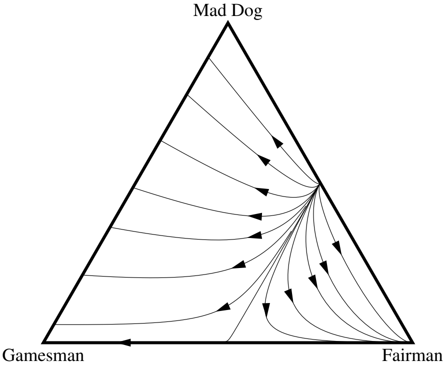

Preservation of dominated strategies

Under the Smith dynamics, even *strictly* dominated strategies can survive. (They are eliminated by the replicator dynamics.)

	R	P	S	T
R	0	-2	1	1
P	1	0	-2	-2
S	-2	1	0	0
T	$-2-\delta$	$1-\delta$	$-\delta$	$-\delta$

Rock-Paper-Scissors with a feeble twin

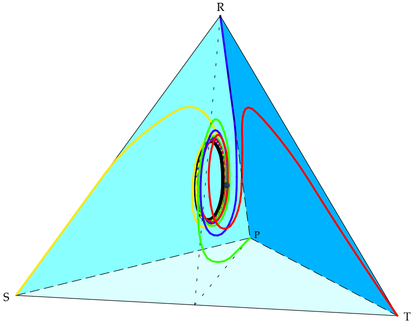

(ii) bad RPS with a feeble twin

The outcomes of evolutionary games may differ radically from classical game theoretic conceptions of rational play.

Dynamic convergence and game equilibria

What is the relationship between the Nash equilibria, or ESS, of a game and the rest points of the dynamics? (Let's concentrate on the replicator dynamics for simplicity.)

Note the following interpretive issue:

- An ESS is a strategy (possibly mixed) used by a *single* player.
- A rest point $\vec{p}$ is a *distribution* of pure strategies in a population which are fixed under the dynamics.

Dynamic convergence and game equilibria

Definition (see Hofbauer and Sigmund, 2002)

Let $\hat{\mathbf{x}} = \langle \hat{x}_1, \ldots, \hat{x}_n \rangle \in S^n$ be a population state for the continuous replicator dynamics with payoff matrix A. Then $\hat{\mathbf{x}}$ is said to be an *evolutionarily stable state* if $\hat{\mathbf{x}} \cdot A\mathbf{x} > \mathbf{x} \cdot A\mathbf{x}$ for all $\mathbf{x} \neq \hat{\mathbf{x}}$ in a neighbourhood of $\hat{\mathbf{x}}$.

An evolutionarily stable state is essentially an evolutionarily stable strategy *interpreted as a population distribution*.

The requirement that $\hat{\mathbf{x}} \cdot A\mathbf{x} > \mathbf{x} \cdot A\mathbf{x}$ is essentially the same as noting that an ESS is locally superior.

Dynamic convergence and game equilibria

One notion of stability is that there is no local "push" away from the fixed point. We can make this precise as follows:

Definition (Lyapunov stability)

Let $\hat{\mathbf{x}} \in S^n$ be a fixed point of the replicator dynamics. Then $\hat{\mathbf{x}}$ is *stable* if, for every $\epsilon > 0$ there exists a $\delta > 0$ such that $\|\mathbf{x}(0) - \hat{\mathbf{x}}\| < \delta$ implies $\|\mathbf{x}(t) - \hat{\mathbf{x}}\| < \epsilon$, for all $t \geq 0$.

Dynamic convergence and game equilibria

Another notion of stability is that there is a local "pull" toward the fixed point. This can be made precise as follows:

Definition

Let $\hat{\mathbf{x}}$ be a fixed point of the replicator dynamics. Then $\hat{\mathbf{x}}$ is *asympotically stable* if it is stable and, in addition, there exists a $\delta > 0$ such that if $\|\mathbf{x}(0) - \hat{\mathbf{x}}\| < \delta$, then $\lim_{t\to\infty} \mathbf{x}(t) = \hat{\mathbf{x}}$.

If a point $\hat{\mathbf{x}}$ is stable but not asymptotically stable, then $\hat{\mathbf{x}}$ is said to be *neutrally stable*.

Dynamic convergence and game equilibria

Theorem

If $\hat{\mathbf{x}} \in S^n$ is an evolutionarily stable state for the game with payoff matrix A, then $\hat{\mathbf{x}}$ is an asymptotically stable rest point of the continuous replicator dynamics.

	Rock	Paper	Scissors
Rock	$(0,0)$	$(-1,2)$	$(2,-1)$
Paper	$(2,-1)$	$(0,0)$	$(-1,2)$
Scissors	$(-1,2)$	$(2,-1)$	$(0,0)$

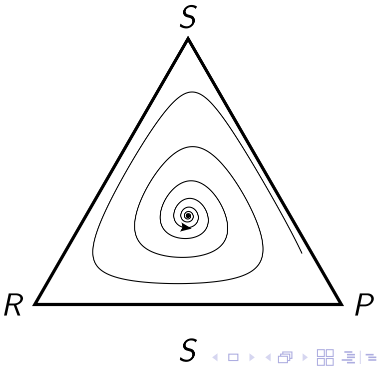

Dynamic convergence and game equilibria

Definition

Suppose that $\mathbf{x}(t) = \langle x_1 t, \ldots, x_n(t) \rangle$ is a solution to the replicator dynamics for some game and some initial condition $\mathbf{x}(0)$. The ω-limit of $\mathbf{x}(t)$, denoted $\omega(\mathbf{x}(t))$ is the set of its accumulation points. That is,

$$\omega(\mathbf{x}(t)) = \{\, \boldsymbol{v} \in S^n \,|\, \mathbf{x}(t_k) \to \boldsymbol{v} \text{ for some sequence } t_k \to +\infty \,\}$$

(The ω-limit of a path is the set of points whose neighbourhoods get visited infinitely often.)

Dynamic convergence and game equilibria

Theorem

1. If σ is a Nash equilibrium of a two-player symmetric game, then σ (interpreted as a population distribution) is a rest point of the continuous replicator dynamics.
2. If $\mathbf{x}$ is the ω-limit of an orbit of the replicator dynamics in the interior of S^n, then $\mathbf{x}$ (interpreted as a mixed strategy) is a Nash equilibrium.
3. If $\mathbf{x}$ is Lyapunov stable, then $\mathbf{x}$ (interpreted as a strategy) is a Nash equilibrium.

Can ② be strengthened further? I.e., if $\mathbf{x}$ is asymptotically stable, then $\mathbf{x}$ is an ESS?

Dynamic convergence and game equilibria

	S_1	S_2	S_3
S_1	0	5	-4
S_2	-7	0	8
S_3	-1	2	0

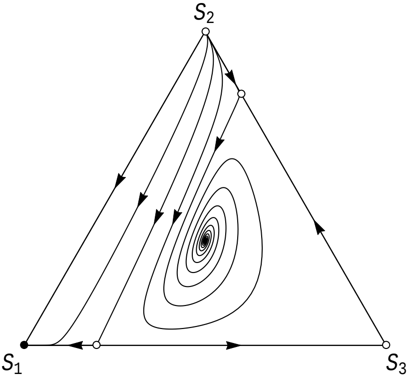

The state $\mathbf{x} = \left\langle \frac{1}{3}, \frac{1}{3}, \frac{1}{3} \right\rangle$ is asymptotically stable, but so is $\hat{\mathbf{x}} = \langle 1, 0, 0 \rangle$.

Since $\hat{\mathbf{x}}$ is an ESS, $\mathbf{x}$ cannot be, since a completely mixed ESS is unique.

Dynamic convergence and game equilibria

However, a partial converse obtains for certain games.

Definition

A game is said to be *doubly symmetric* if, for every pair of pure strategies s_i and s_j, it is the case that $\pi(s_i|s_j) = \pi(s_j|s_i)$.

Theorem

For any doubly symmetric game, the following two statements are equivalent:

1. σ is an ESS.
2. The population state σ is asymptotically stable under the continuous linear replicator dynamics.

Outline

1. Evolutionarily stable strategies
   - Basic concepts
   - Alternative characterisations of ESS
2. Dynamics and convergence
3. The complexity of ESS

Complexity classes

We would like to know the *complexity class* of computing a Nash equilibrium. Complexity classes organise problems based on their difficulty.

Some problems are easy to solve:

Sorting

Input: A list of numbers $L = \{ n_1, \ldots, n_k \}$.

Output: A rearrangement of L from smallest to largest.

Complexity classes

We would like to know the *complexity class* of computing a Nash equilibrium. Complexity classes organise problems based on their difficulty.

Some problems, surprisingly, are easy to solve:

Primality

Input: An n-bit integer p.

Question: Is p prime?

Complexity classes

We would like to know the *complexity class* of computing a Nash equilibrium. Complexity classes organise problems based on their difficulty.

Some problems are surprisingly hard to solve:

Committee formation

Input: A set of employees E, and a list of all pairs who do not get along.

Task: Form a committee of size K, no two of whom are incompatible, or show this is impossible.

We don't know of an efficient way to do this.

Complexity classes

Easy problems are those where an efficient algorithm is known.

“Efficient,” here, means one which takes polynomial time in the size of the input.

This is the complexity class P.

Complexity classes

For some problems, it is easy to *verify* a purported solution, even if it is hard to *find* a solution.

Hamiltonian path

Input: A graph $G = (V, E)$.

Question: Is there a path which visits each vertex exactly once?

Complexity classes

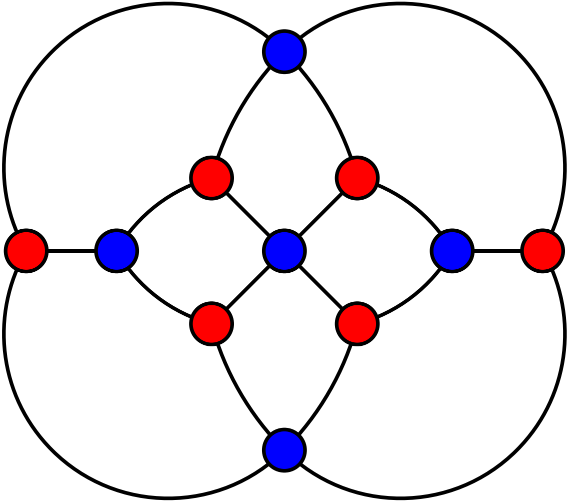

Complexity classes

The class NP

NP is the class of decision problems for which, whenever the answer for a particular instance is “yes,” there is a proof of this fact which can be checked in polynomial time.

Hamiltonian Path and **Committee Formation** are problems in NP.

Complexity classes

Some problems in NP have the remarkable property that they are amongst the *hardest* problems in NP, in the following sense:

The class NP-complete

A problem B is NP-*complete* if, for any problem A in NP, there is a polynomial-time reduction from A to B.

In other words, if we can solve B efficiently, we can solve *any* problem in NP efficiently.

The question “Is P≠NP?” is the greatest unsolved problem in computer science.

How hard is it to find a Nash equilibrium?

Theorem

The following are NP-complete problems, even for symmetric games: Given a two-player game in strategic form, does it have

- At least two Nash equilibria?
- A nash equilibrium in which player 1 has utility at least k?
- A Nash equilibrium with support of size at least N?
- A Nash equilibrium whose support contains strategy s?
- A Nash equilibrium whose support does not contain strategy s?

How hard is it to find a Nash equilibrium?

The general problem of finding a Nash equilibrium is thought not to be NP-complete.

It has been shown to be PPAD-complete ("Polynomial Parity Arguments on Directed graphs").

This makes it equivalent to the following problems: in difficulty:

- Finding a Brouwer fixed point
- Finding an Arrow-Debreu market equilibrium
- Finding a cutting plane as in the Ham Sandwich theorem

How hard is it to find an ESS?

The class NP-hard

A problem B is NP-hard if and only if there exists an NP-complete problem A that has a polynomial-time reduction to B.

That is, NP-hard problems are *at least* as hard as the NP-complete problems (and possibly harder).

Suri (2007) showed that finding an evolutionarily stable strategy is NP-hard.

▸ NP-hard diagram

Evolutionary game dynamics may be uncomputable

In *The Philosophical Computer*, Patrick Grim showed how certain combinations of strategies in the indefinitely iterated spatial Prisoner's Dilemma can be used to construct Minsky machines.

Thus the answer to the question "Will this strategy eventually drive other strategies to extinction?" is, in general, not computable.

▸ Complex evolutionary dynamics

References and miscellaneous readings I

Hofbauer, Josef and Karl Sigmund (2002). *Evolutionary Games and Population Dynamics*. Cambridge University Press.

Maynard Smith, John and George Price (1973). "The Logic of Animal Conflict". *Nature* 246: 15–18.

Papadimitriou, Christos H. (1994). "On the Complexity of the Parity Argument and Other Inefficient Proofs of Existence". *Journal of Computer and System Sciences* 48, 3: 498–532.

Sandholm, William H. (2010). *Population Games and Evolutionary Dynamics*. MIT Press.

Suri, Siddharth (2007). "Computational Evolutionary Game Theory". In Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani, eds., *Algorithmic Game Theory*, Cambridge University Press, chapter 29, pp. 717–735.

How hard is it to find an ESS?

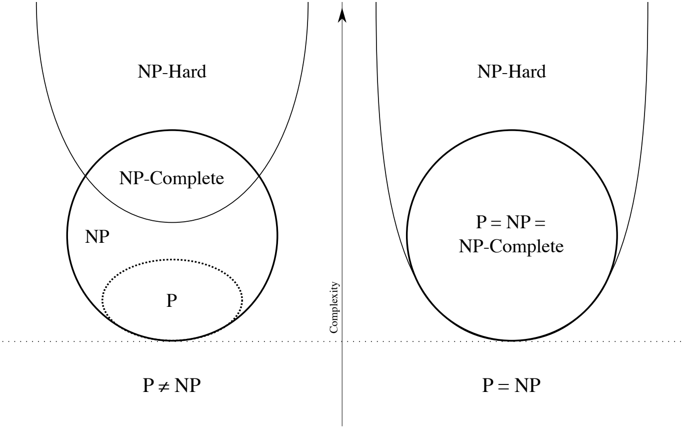

Return

The expressiveness of evolutionary dynamics

The uncomputability of answers to certain evolutionary questions derives from how expressive evolutionary dynamics can be.

▸ Return