Evolutionary Stochastic Games

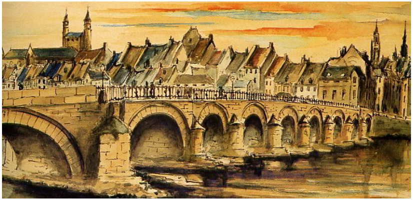

Frank Thuijsman

joint work with J. Flesch, T. Parthasarathy, P. Uyttendaele

Dyn Games Appl (2013) 3, 207–219

Outline

1928, John von Neumann

2-Person Zerosum Games

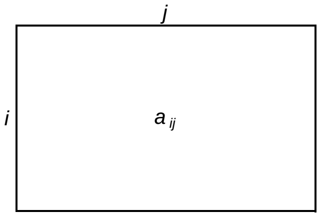

Existence of Value and Optimal Strategies

1951, John Nash

n-Person Non-Zerosum Games

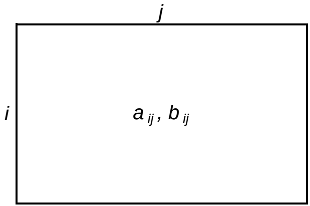

Existence of Equilibria

1953, Lloyd Shapley

2-Person Zerosum Stochastic Games

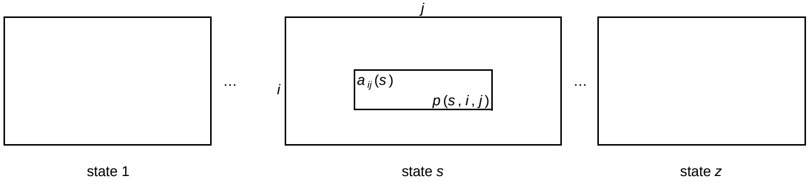

Existence of Value and Optimal Stationary Strategies
for Stopping Stochastic Games

1973, John Maynard Smith and George Price

Evolutionary Games

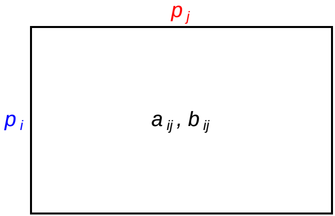

with $b_{ij} = a_{ji}$

- Population consisting of Individuals of Different Types playing against Itself
- Population Distribution $p = (p_1, p_2, \ldots, p_n)$
- Individuals of Type k have Fitness $e_k A p^\top$ in Population p
- Concept of Evolutionarily Stable Strategies (ESS)

Question

What if the Fitness of Population Members corresponds to the Average Rewards in a Stochastic Game rather than to the Expected Payoffs in a One-Shot Game?

The Stochastic Game Model

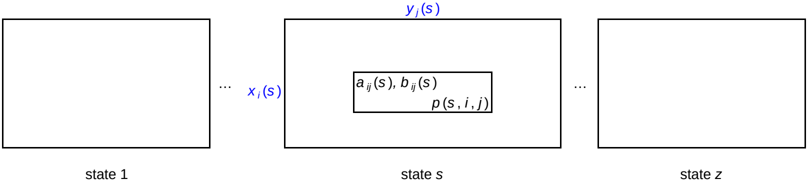

- Finitely Many States, Finitely Many Actions for each Player
- Payoffs and Transitions at each Stage $1, 2, 3, 4, \ldots$
- Each State can serve as Initial State
- Complete Information and Perfect Recall
- Discounting or Averaging the Stage Payoffs

Some Highlights of Stochastic Game Theory

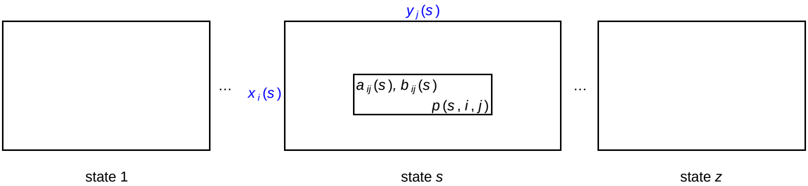

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value
- 1957, H. Everett / D. Gillette:
  2-Person Zerosum Undiscounted Stochastic Games
- 1964, A.M. Fink / M. Takahashi:
  n-Person β-Discounted Stochastic Games - Equilibria
- 1981, J.-F. Mertens and A. Neyman:
  2-Person Zerosum Undiscounted Stochastic Games - Value
- 2000, N. Vieille:
  2-Person Undiscounted Stochastic Games - ε-Equilibria

Abraham Neyman - A Personal Account

1973, John Maynard Smith and George Price

Evolutionary Games

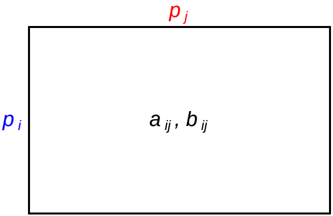

with $b_{ij} = a_{ji}$

- Population consisting of Individuals of Different Types playing against Itself
- Population Distribution $p = (p_1, p_2, \ldots, p_n)$
- Individual of Type k has Fitness $e_k A p^\top$ in Population p
- Concept of Evolutionarily Stable Strategies (ESS)

The ESS Concept

Evolutionary Games

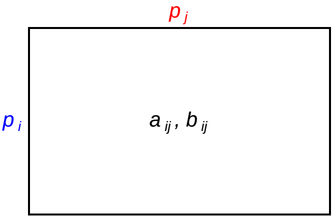

with $b_{ij} = a_{ji}$

ESS: Population Distribution $p = (p_1, p_2, \ldots, p_n)$ with

- $pAp^\top \geq qAp^\top \quad \forall q$
- If $q \neq p$ and $qAp^\top = pAp^\top$, then $pAq^\top > qAq^\top$

The Replicator Dynamic by Taylor and Jonker, 1978

Evolutionary Games

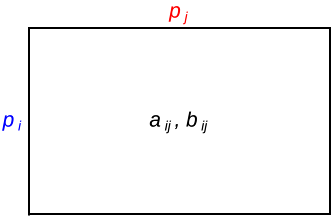

with $b_{ij} = a_{ji}$

Population Development by the Replicator Equation:

- $\dot{p}_k = p_k \left(e_k A p^\top - p A p^\top \right)$

Assumptions for Evolutionary Stochastic Games

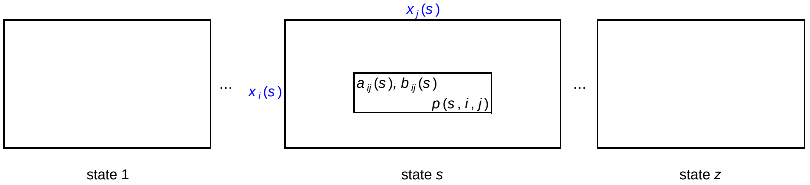

- Population playing against Itself
- *Types now correspond to Pure Stationary Strategies*
- Symmetric Payoffs: $b_{ij} = a_{ji}$
- Symmetric Transitions: $p(s,i,j) = p(s,j,i)$
- Irreducible Stochastic Game
  All States communicate for all Stationary Strategies

Assumptions Continued

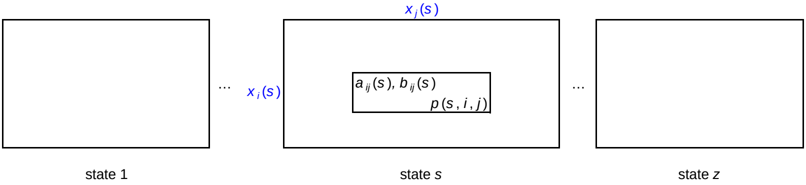

- Fitness of Individual of Type k in Population Distribution $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$, taken over Pure Stationary Strategies, is Average Reward $\gamma(e_k, x)$, where x is Stationary Strategy induced by $\bar{x}$

Assumptions Continued

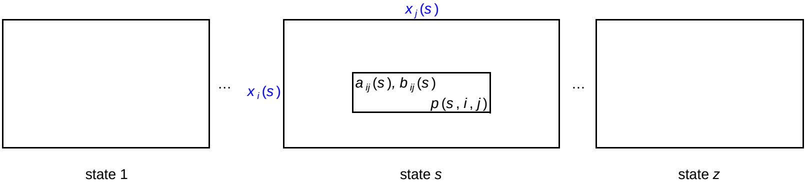

- Fitness of Individual of Type k in Population Distribution $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$, taken over Pure Stationary Strategies, is Average Reward $\gamma(e_k, x)$, where x is Stationary Strategy induced by $\bar{x}$
- Different Populations can give *Same* Stationary Strategy

Assumptions Continued

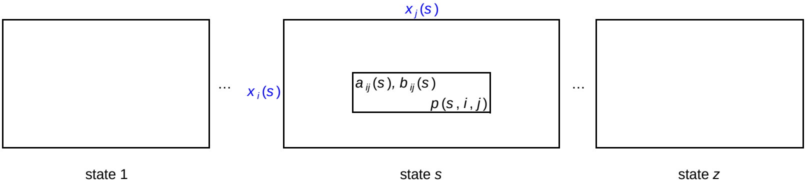

- Fitness of Individual of Type k in Population Distribution $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$, taken over Pure Stationary Strategies, is Average Reward $\gamma(e_k, x)$, where x is Stationary Strategy induced by $\bar{x}$
- Different Populations can give *Same* Stationary Strategy
- *Stationary Strategy* x is ESS if
  - $\gamma(x,x) \geq \gamma(y,x) \ \ \forall y$
  - If $y \neq x$ and $\gamma(y,x) = \gamma(x,x)$, then $\gamma(x,y) > \gamma(y,y)$

Assumptions Continued

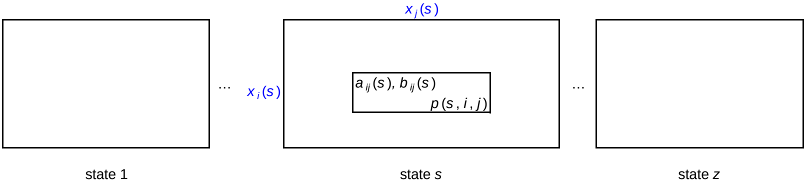

- Fitness of Individual of Type k in Population Distribution $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$, taken over Pure Stationary Strategies, is Average Reward $\gamma(e_k, x)$, where x is Stationary Strategy induced by $\bar{x}$
- Different Populations can give *Same* Stationary Strategy
- *Stationary Strategy* x is ESS if
  - $\gamma(x,x) \geq \gamma(y,x) \quad \forall y$
  - If $y \neq x$ and $\gamma(y,x) = \gamma(x,x)$, then $\gamma(x,y) > \gamma(y,y)$
- Population Development by Replicator Dynamic
  - $\dot{\bar{x}}_k = \bar{x}_k \left(\gamma(e_k, x) - \gamma(x,x)\right)$

Some Remarks

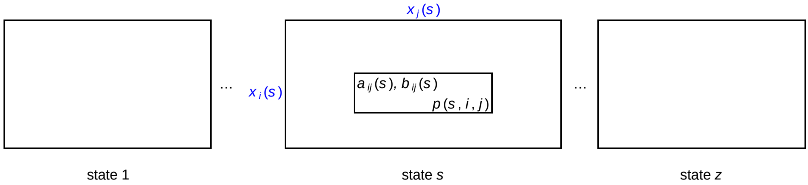

- ESS Not Always Exists
- Replicator Dynamic Not Always Converges
- Limit Points of Dynamic Not Always give ESS

Time Scale Assumption

Time Scale Assumption

- We assume that the Stochastic Game Horizon corresponds to Individual Life Time, which is Negligibly Small compared to the Time Scale of the Evolutionary Population Dynamics

Time Scale Assumption

- We assume that the Stochastic Game Horizon corresponds to Individual Life Time, which is Negligibly Small compared to the Time Scale of the Evolutionary Population Dynamics
- We assume that during Individual Life Time All States are visited Sufficiently Often

Time Scale Assumption

- We assume that the Stochastic Game Horizon corresponds to Individual Life Time, which is Negligibly Small compared to the Time Scale of the Evolutionary Population Dynamics
- We assume that during Individual Life Time All States are visited Sufficiently Often
- The Infinite Horizon Average Reward approximates the Finite Horizon Average Reward

Multi-State Interpretation

Multi-State Interpretation

- The Individual Fitness depends on the Actions taken at *Multiple Situations* encountered in Life *Altogether*

Multi-State Interpretation

- The Individual Fitness depends on the Actions taken at *Multiple Situations* encountered in Life *Altogether*
- The Fraction of Time that *Specific* Situations govern Individual Life, *depends* on the Actions taken by *All* Population Members

Different Populations, Same Stationary Strategy

In a Game with 2 States and in each State 2 Actions, T and B:

$$\frac{1}{2}\cdot(T,T)+\frac{1}{2}\cdot(B,B)=((\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}))=\frac{1}{2}\cdot(T,B)+\frac{1}{2}\cdot(B,T)$$

Question

Suppose $x=((\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}))$ is an ESS, does it imply that any Pure Stationary Strategy in $C(x)$ is also a Best Reply to x?
In other words:
Do $\gamma((T,T),x)$, $\gamma((B,B),x)$, $\gamma((T,B),x)$, $\gamma((B,T),x)$ all equal $\gamma(x,x)$?

Answer

Yes

Irreducible Markov Decision Problems

Theorem

If x^* is a Stationary Optimal Strategy in an Irreducible MDP and x is a Stationary Strategy with $C(x_s) \subset C(x_s^*)$ for each s, then x is Optimal as well.

Sketch of Proof:

We show that for a Stationary Strategy in an Irreducible MDP the Average Reward is a Convex Combination of the Average Rewards for the Pure Stationary Strategies in its Carrier. □

Evolutionary Stochastic Games

Corollary

If x^* is an ESS in an Evolutionary Stochastic Game and $x \neq x^*$ is a Stationary Strategy with $C(x_s) \subset C(x_s^*)$ for each s, then x is no ESS.

Proof:
If x^* ESS and $C(x_s) \subset C(x_s^*)$ $\forall s$,
then $\gamma(x, x^*) = \gamma(x^*, x^*)$ by previous Theorem,
then $\gamma(x^*, x) > \gamma(x, x)$ by 2$^{\text{nd}}$ ESS Condition for x^*,
which implies x is no ESS by 1$^{\text{st}}$ ESS Condition for x. $\square$

Another Reason for Irreducibility

Example

State 1:

3, 3 (1, 0)	5, 4 (1, 0)
4, 5 (1, 0)	2, 2 (0, 1)

state 1

State 2:

2, 2 (0, 1)

state 2

No Symmetric Equilibrium in Stationary Strategies
Not even Symmetric ε-Equilibrium in Stationary Strategies

Another Reason for Irreducibility

Example

State 1:

	3, 3 (1, 0)	5, 4 (1, 0)
	4, 5 (1, 0)	2, 2 (0, 1)

State 2:

2, 2 (0, 1)

state 1 state 2

No Symmetric Equilibrium in Stationary Strategies
Not even Symmetric ε-Equilibrium in Stationary Strategies
But *there is* a Symmetric ε-Equilibrium

Theorem

For every Symmetric Irreducible Stochastic Game there exists a Symmetric Stationary Equilibrium.

For Discounted as well as for Undiscounted (Average) Rewards

Sketch of Proof:
This follows by applying a Fixed Point Argument
for the Discounted Best Reply Map
and by taking the Limit of Discounted Fixed Points
for the Undiscounted Case. □

A 2 State Example with Replicator Dynamics

1, 1 (0, 1)	4, 3 (.5, .5)
3, 4 (.5, .5)	2, 2 (0, 1)

state 1

3, 3 (1, 0)	5, 4 (.5, .5)
4, 5 (.5, .5)	2, 2 (1, 0)

state 2

(Trajectory)

A 2 State Example with Replicator Dynamics

1, 1 (0, 1)	4, 3 (.5, .5)
3, 4 (.5, .5)	2, 2 (0, 1)

state 1

3, 3 (1, 0)	5, 4 (.5, .5)
4, 5 (.5, .5)	2, 2 (1, 0)

state 2

(Trajectory)

A 3 State Example with Replicator Dynamics

1, 1 (.5, 0, .5)	4, 3 (.5, .5, 0)
3, 4 (.5, .5, 0)	2, 2 (0, .5, .5)

state 1

3, 3 (1, 0, 0)	5, 4 (.5, 0, .5)
4, 5 (.5, 0, .5)	2, 2 (0, 0, 1)

state 2

4, 4 (0, 1, 0)	6, 7 (0, .5, .5)
7, 6 (0, .5, .5)	5, 5 (1, 0, 0)

state 3

(Trajectory)

About this Model

Further Research:

- Finding Real Life Applications that Fit
- Exploring the Existence of Symmetric (ε-)Equilibria for Symmetric Stochastic Games

Local Interactions and Fitnesses in Space/Time

$$\begin{pmatrix} 1 & 2 & 3+1.5\sin(\pi(0.5t+y)/40) \\ 2 & 1.1 & 2 \\ 3-1.5\sin(\pi(0.5t+y)/40) & 2 & 1 \end{pmatrix}$$

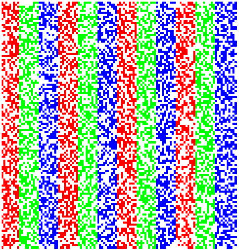

GAMES 2016 & EC'16

5-th World Congress of the Game Theory Society

17-th ACM Conference on Economics and Computation: EC'16

Maastricht, 24-28 July 2016

Thanks!


GAMES
2016
5th World Congress of the Game Theory Society
Maastricht University
24-28 JULY 2016
MAASTRICHT
THE NETHERLANDS
GAMES2016.NL