Population Dynamics in Stochastic Games

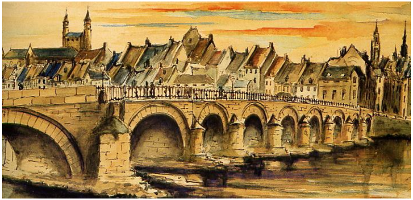

Frank Thuijsman

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

Outline

1928, John von Neumann

2-Person Zerosum Games

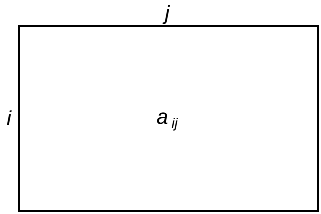

Existence of Value and Optimal Strategies

1951, John Nash

n-Person Non-Zerosum Games

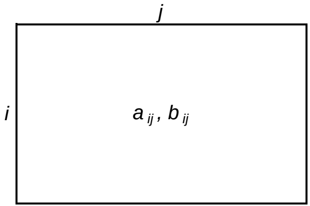

Existence of Equilibria

1953, Lloyd Shapley

2-Person Zerosum Stochastic Games

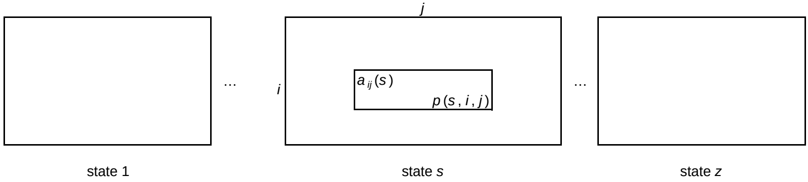

Existence of Value and Optimal Strategies

1953, Lloyd Shapley

2-Person Zerosum Stochastic Games

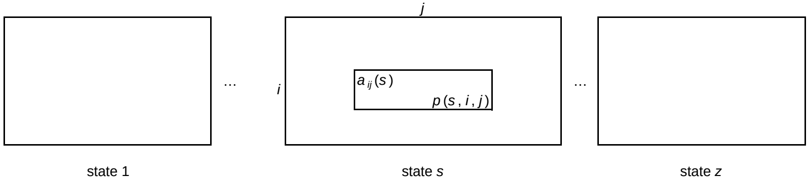

Existence of Value and Optimal Strategies
for Stopping Stochastic Games

1973, John Maynard Smith and George Price

Evolutionary Games

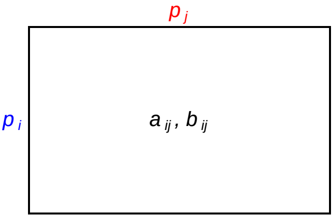

with $b_{ij} = a_{ji}$

- 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 Average Rewards in a Stochastic Game, rather than to Expected Payoffs in a One-Shot Game?

The Stochastic Game Model

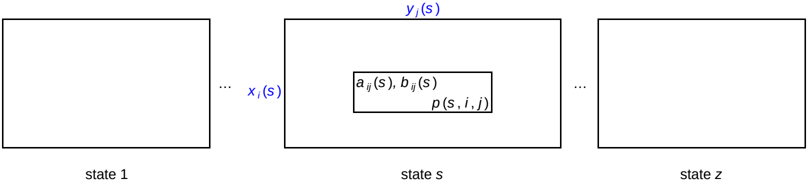

The Stochastic Game Model

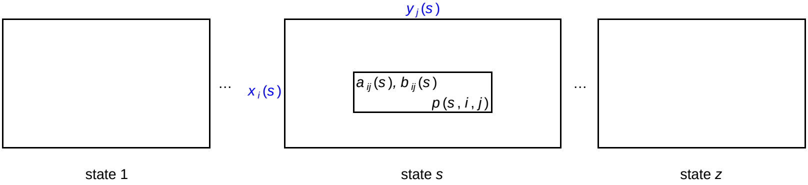

- Finitely Many States, Finitely Many Actions for each Player

The Stochastic Game Model

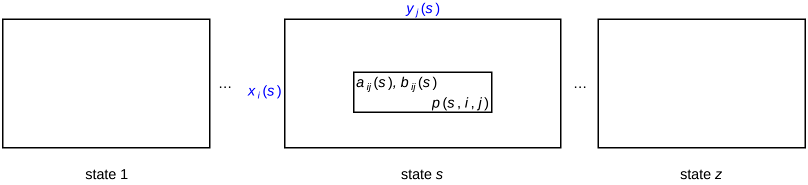

- Finitely Many States, Finitely Many Actions for each Player
- Payoffs and Transitions at each Stage $1, 2, 3, 4, \ldots$

The Stochastic Game Model

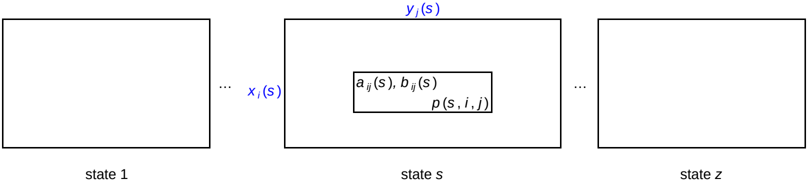

- Finitely Many States, Finitely Many Actions for each Player
- Payoffs and Transitions at each Stage $1, 2, 3, 4, \ldots$
- Each State can serve as Initial State

The Stochastic Game Model

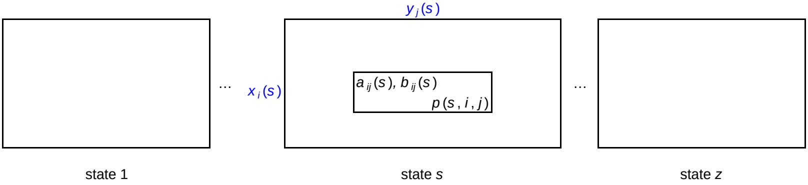

- Finitely Many States, Finitely Many Actions for each Player
- Payoffs and Transitions at each Stage $1, 2, 3, 4, \ldots$
- Each State can serve as Initial State
- Complete Information and Perfect Recall

The Stochastic Game Model

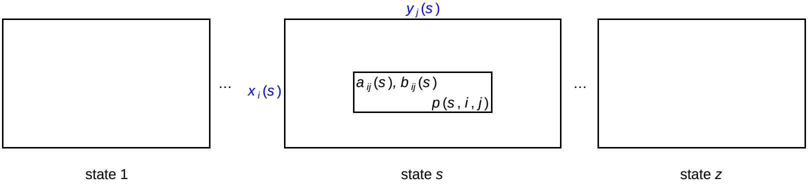

- Finitely Many States, Finitely Many Actions for each Player
- Payoffs and Transitions at each Stage $1, 2, 3, 4, \ldots$
- Each State can serve as Initial State
- Complete Information and Perfect Recall
- Discounting or Averaging the Stage Payoffs

Some Highlights of Stochastic Game Theory

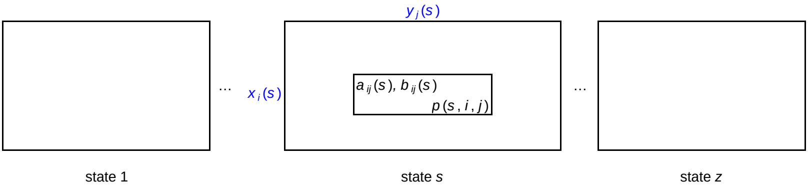

Some Highlights of Stochastic Game Theory

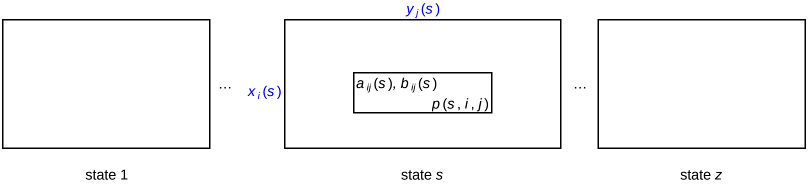

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value

Some Highlights of Stochastic Game Theory

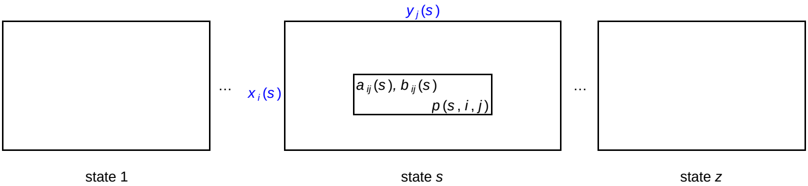

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value
- 1957, H. Everett / D. Gillette:
  2-Person Zerosum Undiscounted Stochastic Games

Some Highlights of Stochastic Game Theory

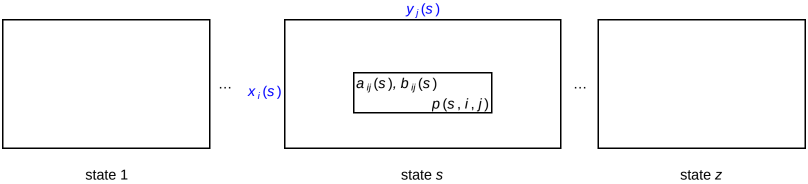

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value
- 1957, H. Everett / D. Gillette:
  2-Person Zerosum Undiscounted Stochastic Games
- 1964, A.M. Fink / M. Takahashi:
  n-Person β-Discounted Stochastic Games - Equilibria

Some Highlights of Stochastic Game Theory

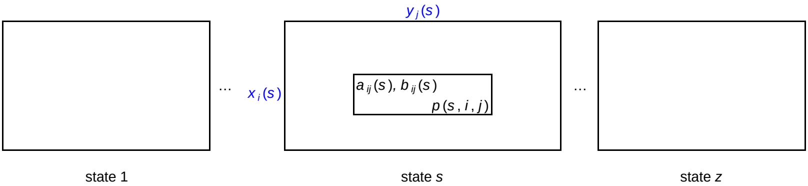

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value
- 1957, H. Everett / D. Gillette:
  2-Person Zerosum Undiscounted Stochastic Games
- 1964, A.M. Fink / M. Takahashi:
  n-Person β-Discounted Stochastic Games - Equilibria
- 1981, J.-F. Mertens and A. Neyman:
  2-Person Zerosum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

Some Highlights of Stochastic Game Theory

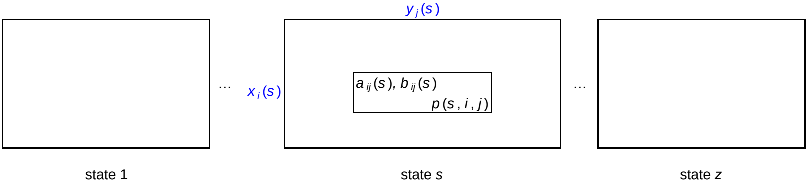

- 1953, L.S. Shapley:
  2-Person Zerosum Stopping Stochastic Games - Value
- 1957, H. Everett / D. Gillette:
  2-Person Zerosum Undiscounted Stochastic Games
- 1964, A.M. Fink / M. Takahashi:
  n-Person β-Discounted Stochastic Games - Equilibria
- 1981, J.-F. Mertens and A. Neyman:
  2-Person Zerosum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

1973, John Maynard Smith and George Price

Evolutionary Games

$$\begin{array}{cc} & p_j \\ p_i & \boxed{a_{ij}, b_{ij}} \end{array} \quad \text{with } b_{ij} = a_{ji}$$

- 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

$$\begin{array}{cc} & p_j \\ p_i & \boxed{\quad a_{ij}, b_{ij} \quad} \end{array} \qquad \text{with } b_{ij} = a_{ji}$$

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

$$
\begin{array}{cc}
 & p_j \\
p_i & \boxed{\quad a_{ij}, b_{ij} \quad}
\end{array}
\qquad \text{with } b_{ij} = a_{ji}
$$

Population Development by the Replicator Equation:

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

Remarks on ESS and Asymptotic Stability

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process
- ESS Not Always Exists

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process
- ESS Not Always Exists
- Replicator Dynamic Not Always Converges

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process
- ESS Not Always Exists
- Replicator Dynamic Not Always Converges
- Any ESS is Asymptotically Stable

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process
- ESS Not Always Exists
- Replicator Dynamic Not Always Converges
- Any ESS is Asymptotically Stable
- Limit Points of Dynamic Not Always ESS

Remarks on ESS and Asymptotic Stability

- A Static Concept and a Dynamic Process
- ESS Not Always Exists
- Replicator Dynamic Not Always Converges
- Any ESS is Asymptotically Stable
- Limit Points of Dynamic Not Always ESS

$$\begin{pmatrix} 0 & 6 & -4 \\ -3 & 0 & 5 \\ -1 & 3 & 0 \end{pmatrix}$$ (Hofbauer and Sigmund, 1998)

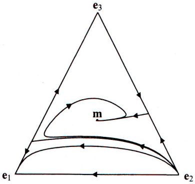

Assumptions for Evolutionary Stochastic Games

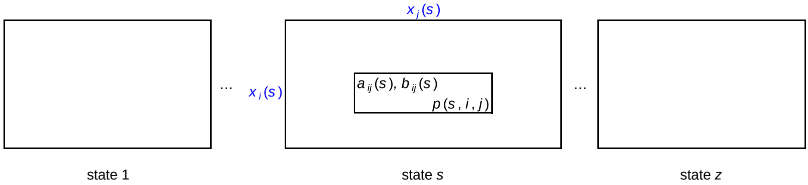

Assumptions for Evolutionary Stochastic Games

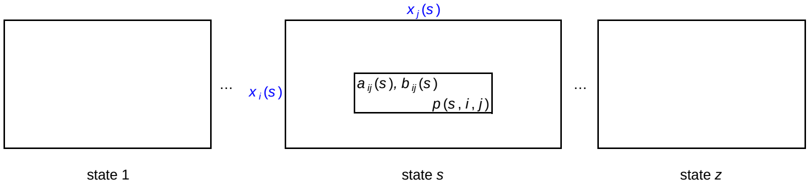

- Population playing against Itself

Assumptions for Evolutionary Stochastic Games

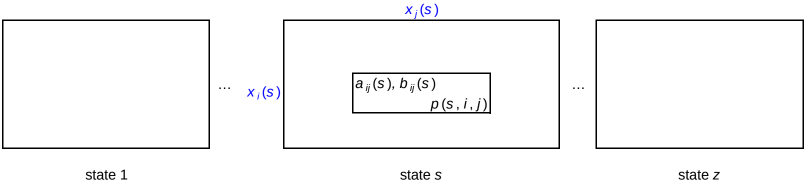

- Population playing against Itself
- Symmetric Payoffs: $b_{ij} = a_{ji}$

Assumptions for Evolutionary Stochastic Games

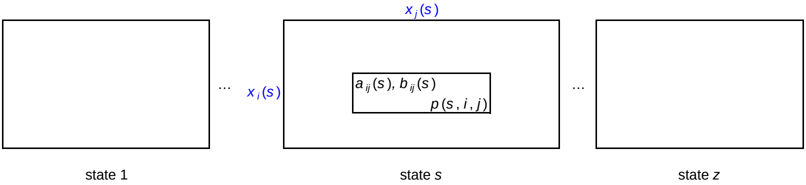

- Population playing against Itself
- Symmetric Payoffs: $b_{ij} = a_{ji}$
- Symmetric Transitions: $p(s, i, j) = p(s, j, i)$

Assumptions for Evolutionary Stochastic Games

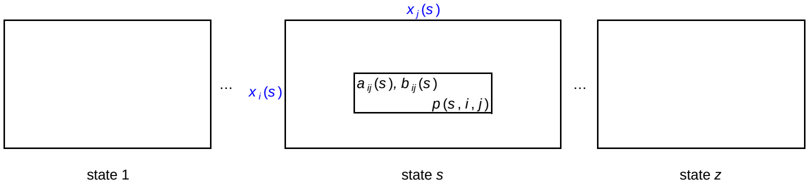

- Population playing against Itself
- Symmetric Payoffs: $b_{ij} = a_{ji}$
- Symmetric Transitions: $p(s, i, j) = p(s, j, i)$
- Irreducible Stochastic Game

Assumptions for Evolutionary Stochastic Games

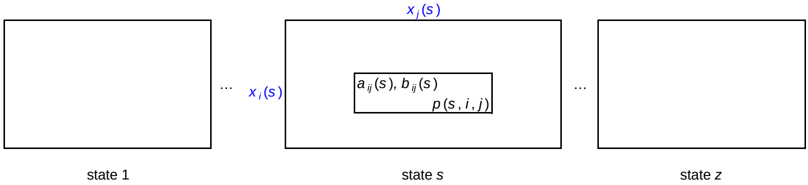

- Population playing against Itself
- Symmetric Payoffs: $b_{ij} = a_{ji}$
- Symmetric Transitions: $p(s, i, j) = p(s, j, i)$
- Irreducible Stochastic Game
  All States communicate for all Stationary Strategies

Assumptions for Evolutionary Stochastic Games

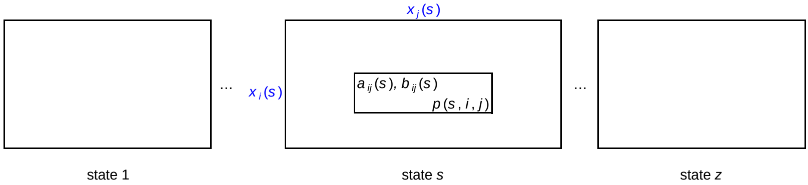

- Population playing against Itself
- Symmetric Payoffs: $b_{ij} = a_{ji}$
- Symmetric Transitions: $p(s, i, j) = p(s, j, i)$
- Irreducible Stochastic Game
  All States communicate for all Stationary Strategies
- Types correspond to Pure Stationary Strategies

Assumptions Continued

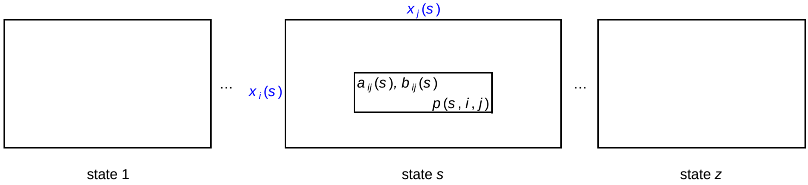

Assumptions Continued

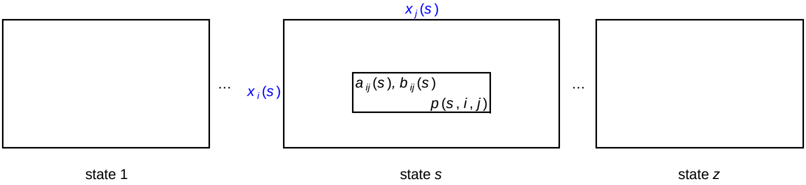

- Fitness of Individual of Type k in Population $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$ is Average Reward $\gamma(e_k, x)$ where x Stationary Strategy determined by $\bar{x}$

Assumptions Continued

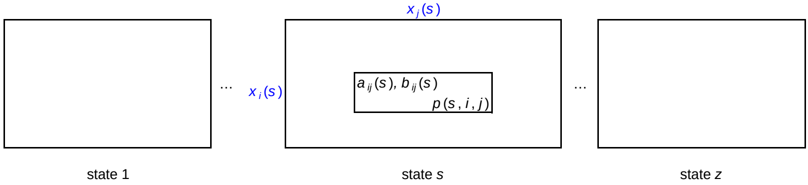

- Fitness of Individual of Type k in Population $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$ is Average Reward $\gamma(e_k, x)$ where x Stationary Strategy determined by $\bar{x}$
- Different Populations can give Same Stationary Strategy

Assumptions Continued

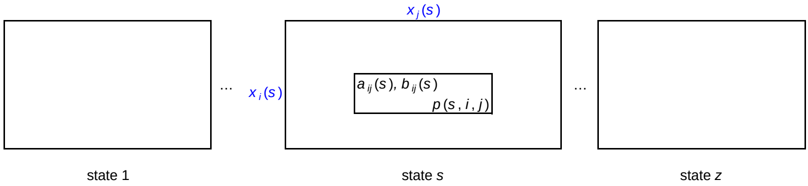

- Fitness of Individual of Type k in Population
  $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$
  is Average Reward $\gamma(e_k, x)$
  where x Stationary Strategy determined by $\bar{x}$
- Different Populations can give Same Stationary Strategy
- *Stationary Strategy* x is ESS if
  - $\gamma(x,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

Assumptions Continued

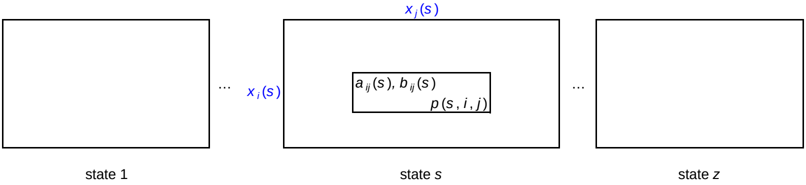

- Fitness of Individual of Type k in Population
  $\bar{x} = (\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_n)$
  is Average Reward $\gamma(e_k, x)$
  where x Stationary Strategy determined by $\bar{x}$
- Different Populations can give Same Stationary Strategy
- *Stationary Strategy* x is ESS if
  - $\gamma(x,x)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

Some Remarks

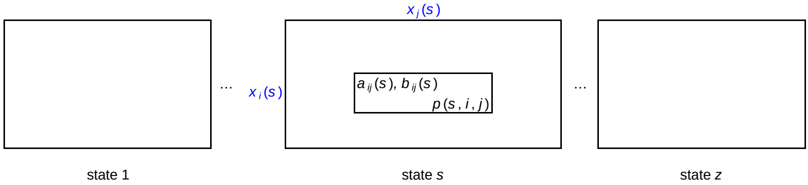

- ESS Not Always Exists
- Replicator Dynamic Not Always Converges
- Limit Points of Dynamic Not Always give ESS

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
- The Infinite Horizon Average Reward approximates the Finite Horizon Average Reward

Different Populations, Same Stationary Strategy

Example

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

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

Different Populations, Same Stationary Strategy

Example

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

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

Question

Suppose $((\frac{1}{2},\frac{1}{2}),(\frac{1}{2},\frac{1}{2}))$ is an ESS,
does it imply $\gamma(T,T)=\gamma(B,B)=\gamma(T,B)=\gamma(B,T)$?

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.

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.

Here Irreducibility is Essential!

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.

Here Irreducibility is Essential!

Example

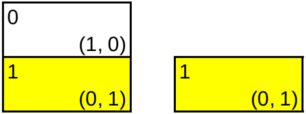

state 1 state 2

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.

Proof follows by showing that:

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.

Proof follows by showing that:
For a Stationary Strategy in an Irreducible MDP
the Average Reward is a Convex Combination
of the Averages 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.

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:

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 $C(x_s) \subset C(x_s^*)$ $\forall s$,

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,

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^{nd} ESS Condition,

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^{nd} ESS Condition,
which implies x is no ESS. $\square$

Another Reason for Irreducibility

Example

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

state 1

2, 2 (0, 1)

state 2

Another Reason for Irreducibility

Example

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

state 1

2, 2 (0, 1)

state 2

No Symmetric Equilibrium in Stationary Strategies

Theorem

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

Theorem

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

For Discounted as well as for Undiscounted (Average) Rewards

Theorem

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

For Discounted as well as for Undiscounted (Average) Rewards

Proof follows by:

- a Fixed Point Argument for the Discounted Best Reply Map
- 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

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

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

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:

About this Model

Further Research:

- Finding Real Life Applications that Fit

About this Model

Further Research:

- Finding Real Life Applications that Fit
- Exploring the Relation between Replicator Dynamic and Fictitious Play for Symmetric Stochastic Games

1951, George Brown / Julia Robinson

The Fictitious Play Process:
Playing Best Replies against Observed Action Frequencies

1951, George Brown / Julia Robinson

The Fictitious Play Process:
Playing Best Replies against Observed Action Frequencies

- For Matrix Games FP leads to Optimal Strategies

1951, George Brown / Julia Robinson

The Fictitious Play Process:
Playing Best Replies against Observed Action Frequencies

- For Matrix Games FP leads to Optimal Strategies
- No FP Convergence for Bimatrix Games (Shapley, 1964)

A 2 State Example with Fictitious Play

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

A 2 State Example with Fictitious Play

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 Fictitious Play

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

A 3 State Example with Fictitious Play

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)

Other 'Evolutionary' Work in Maastricht

Other ‘Evolutionary’ Work in Maastricht

Examining the effects of periodic fitness in replicator dynamics

(Trajectory)

Other ‘Evolutionary’ Work in Maastricht

Examining the dynamics of local interactions
Dynamic Stability: Predator-Prey Behaviour

$$\begin{pmatrix} 1 & 8 & 1 \\ 1 & 1 & 8 \\ 8 & 1 & 1 \end{pmatrix}$$

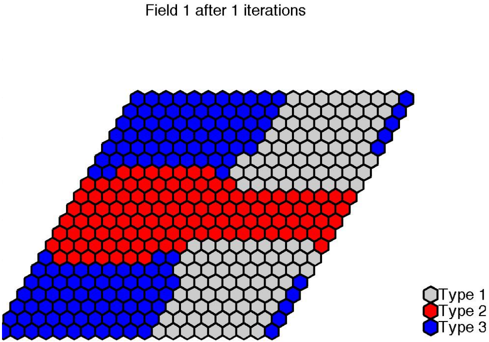

Other ‘Evolutionary’ Work in Maastricht

Studying sex choice ovipositioning behavior of parasitoid wasps

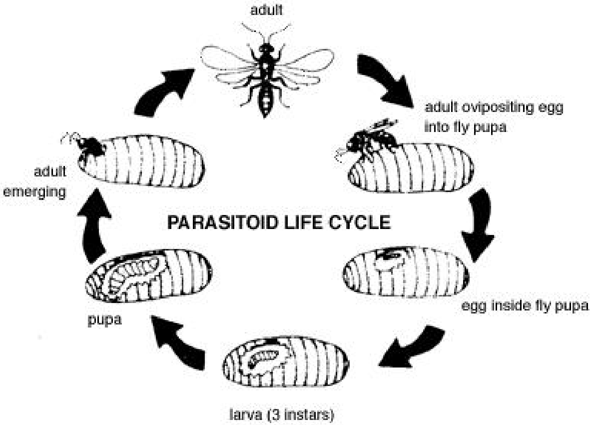

Thanks

Thank you for your attention!
Any comment is welcome!

Paper and presentation will soon be available at
www.personeel.unimaas.nl/F-Thuijsman