Computational Game Theory in Julia

Eric Pacuit, University of Maryland

Lecture 2

ESSLLI 2023

Plan

- Finding Nash equilibria in GameTheory.jl
- Agent based modeling in Julia: Agents.jl

Game in Normal Form

A **game in normal form** is a tuple $\langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ where:

- N is a finite set of players.
- For each $i \in N$, S_i is a (finite) set of actions, or strategies, for player i.
- For each $i \in N$, $u_i : \Pi_{i \in N} S_i \to \mathbb{R}$

Notation

Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a game in strategic form.

- For $s \in \Pi_{i \in N} S_i$, s_i is the ith component of s and $s_{-i} = (s_1, \ldots, s_{i-1}, s_{i+1}, \ldots s_n)$ is the tuple of all strategies except s_i
- For $i \in N$, let $S = \Pi_{i \in N} S_i$ be the set of **strategy profiles**, also called the outcomes of G.
- For $i \in N$, let $S_{-i} = \Pi_{j \in N, j \neq i} S_j$.

Notation

Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a game in strategic form.

- For $s \in \Pi_{i \in N} S_i$, s_i is the ith component of s and $s_{-i} = (s_1, \dots, s_{i-1}, s_{i+1}, \dots s_n)$ is the tuple of all strategies except s_i
- For $i \in N$, let $S = \Pi_{i \in N} S_i$ be the set of **strategy profiles**, also called the outcomes of G.
- For $i \in N$, let $S_{-i} = \Pi_{j \in N, j \neq i} S_j$.
- For a set X, let $\Delta(X)$ be the set of probability measures on X.
- $m \in \Delta(S_i)$ is called a **mixed strategy** for player i.
- A mixed strategy profile is an element of $\Pi_{i \in N} \Delta(S_i)$.

Expected Utility, Best Response

Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a game in strategic form.

For $a \in S_i$ and $p \in \Delta(S_{-i})$ the **expected utility of** a **with respect to** p is

$$EU_i(a, p) = \sum_{t \in S_{-i}} p(t) u_i(a, t)$$

Expected Utility, Best Response

Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a game in strategic form.

For $a \in S_i$ and $p \in \Delta(S_{-i})$ the **expected utility of** a **with respect to** p is

$$EU_i(a, p) = \sum_{t \in S_{-i}} p(t) u_i(a, t)$$

For $X \subseteq \Delta(S_{-i})$, the **best response set for player** i, $BR_i : X \to \wp(S_i)$, is defined as follows: for $p \in X$,

$$BR_i(p) = \{a \mid a \in S_i \text{ and } \forall a' \in S_i : EU_i(a, p) \geqslant EU_i(a', p)\}$$

Identify S_{-i} with the set $\{p \mid p \in \Delta(S_{-i}), p(s) = 1 \text{ for some } s \in S_{-i}\}$,

A strategy profile $s \in \Pi_{i \in N} S_i$ is a (pure strategy) **Nash equilibrium** provided that for all $i \in N$, $s_i \in BR_i(s_{-i})$

Mixed Extension

Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a game in strategic form.

The **mixed extension of** G is the tuple $\langle N, (\Delta(S_i))_{i \in N}, (U_i)_{i \in N} \rangle$, where for $m \in \Pi_{i \in N} \Delta(S_i)$

$$U_i(m) = \sum_{s \in S} u_i(s) \Pi_{i \in N} m_i(s_i)$$

A **mixed strategy Nash equilibrium** in G is a tuple $m \in \Pi_{i \in N} \Delta(S_i)$ that is a Nash equilibrium in the mixed extension of G.

Symmetric Games

Both players rewarded

One player tempted; the other is a sucker

		Player 2	
		c	d
Player 1	c	R, R	S, T
	d	T, S	P, P

One player tempted; the other is a sucker

Both players punished

Symmetric Games

Both players rewarded

One player tempted; the other is a sucker

		Player 2	
		c	d
Player 1	c	R, R	S, T
	d	T, S	P, P

One player tempted; the other is a sucker

Both players punished

Symmetric games are classified in terms of the relationship between R (reward), T (temptation), S (sucker) and P (punishment):

Prisoner's Dilemma

Ann \ Bob	c	d
c	3,3	0,4
d	4,0	1,1

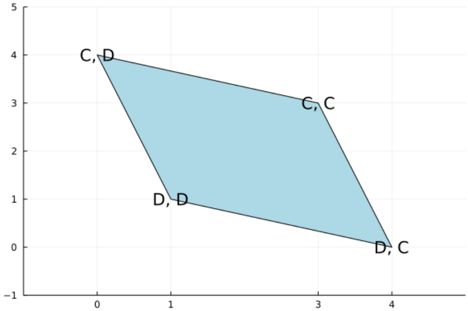

If $T > R > P > S$, then the game is a Prisoner's Dilemma.

d strictly dominates c

(c, c) Pareto dominates (d, d)

(d, d) is the unique Nash equilibrium

Stag Hunt

		Bob c	Bob d
Ann	c	3,3	0,2
Ann	d	2,0	1,1

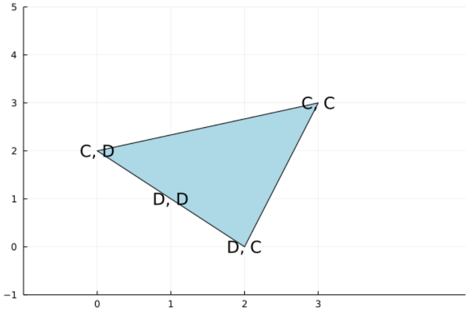

If $R > T$ and $P > S$, then the game is called Stag Hunt.

d is a less "risky" option than c

(c, c) Pareto dominates (d, d)

(c, c) and (d, d) are both Nash equilibria

Chicken

		Bob c	Bob d
Ann	c	2,2	1,3
	d	3,1	0,0

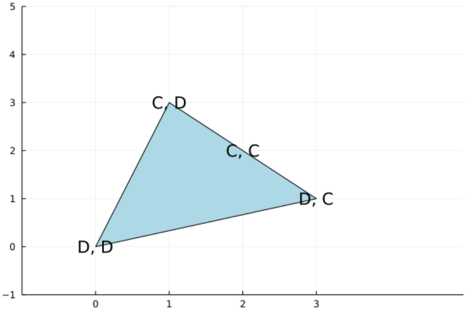

If $T > R$ and $S > P$, then the game is called Chicken (or Hawk-Dove).

c is a less "risky" option than d

(c, c) Pareto dominates (d, d)

(c, d) and (d, c) are both Nash equilibria

Games on a Grid

Fix a set of n agents and a game G.

Put each agent at a point on a grid and randomly assign a strategy C or D to each agent.

During each stage of the simulation:

1. Randomly select an agent a
2. Each of a's neighbors plays G against their neighbors and records the total payout.
3. Agent a imitates the strategy of the player with the maximum total payout.

Modifications

- Add a mutation rate: randomly mutate the chosen strategy.
- Choose 8 random players rather than interacting only with your neighbors.
- Different update rules: choose the agent imitate based on the proportion of average payouts of the neighbors.
- Mix agents with different update rules.
- Use networks rather than a grid.

Grid size: 100×100

Mutation rate: 0.0

Update type: Imitator

Number of steps: 100,000

Number of simulations: 20

Prisoner's Dilemma

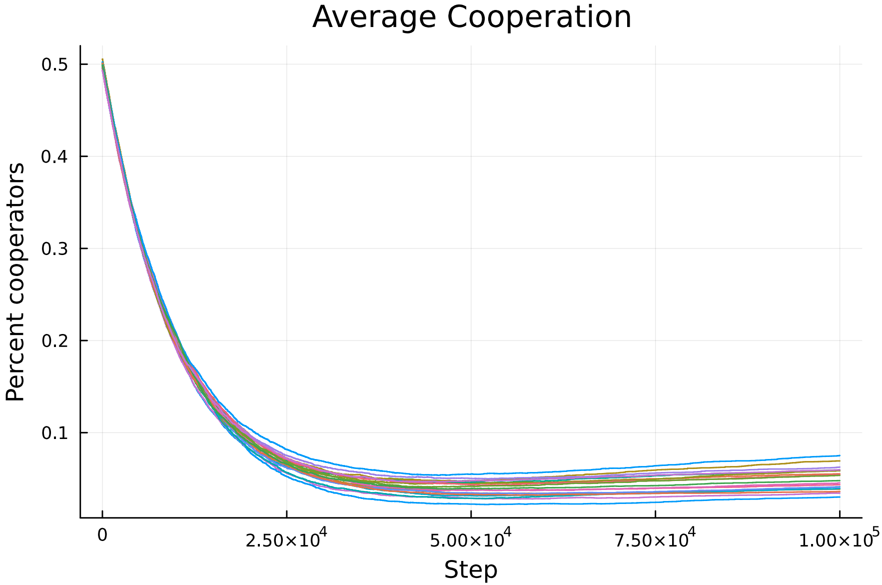

Stag Hunt

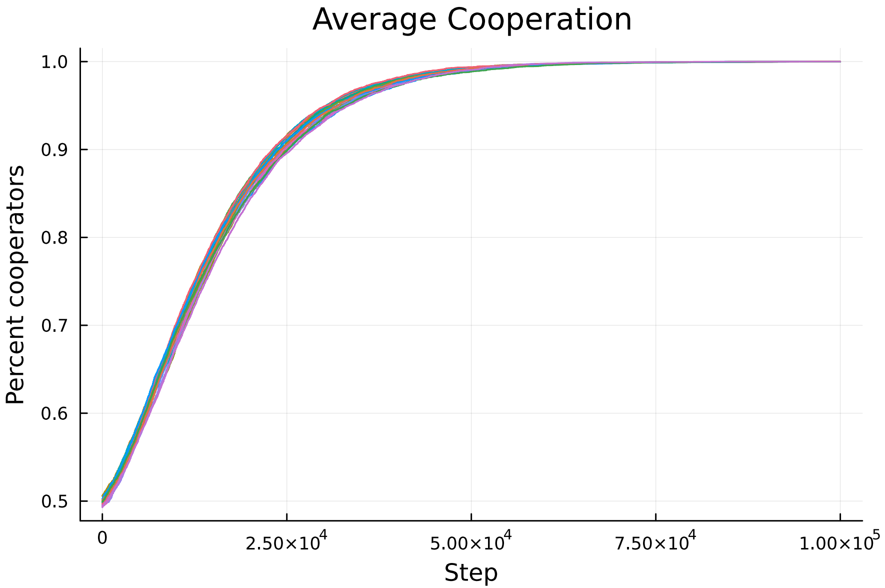
Average Cooperation
Percent cooperators
1.0
0.9
0.8
0.7
0.6
0.5
0
2.50×10^4
5.00×10^4
7.50×10^4
1.00×10^5
Step

Chicken

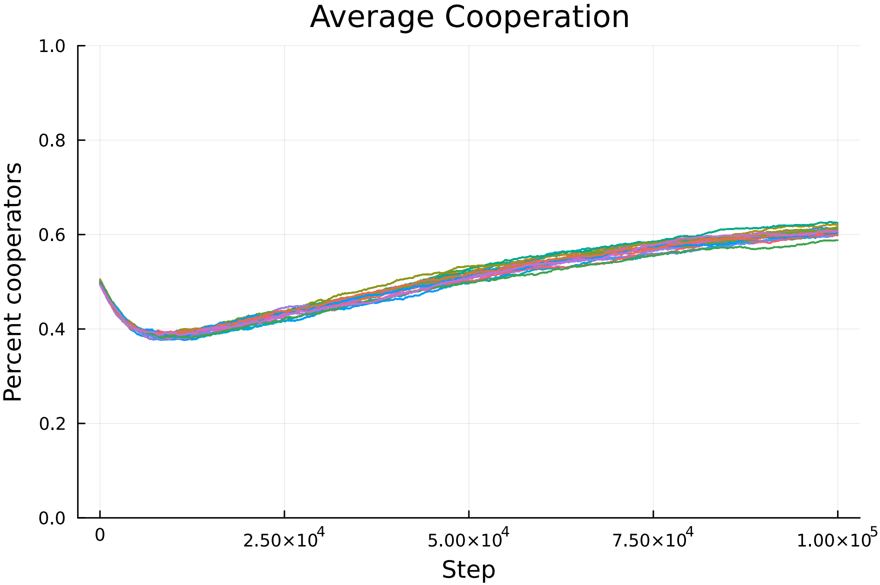

Average Cooperation
Percent cooperators
0.0
0.2
0.4
0.6
0.8
1.0
0
2.50×10^4
5.00×10^4
7.50×10^4
1.00×10^5
Step