Game theory and Python: A tutorial on open-source tools for repeated games

EGAI 2025

Nikoleta E. Glynatsi

Axelrod-Python

https://github.com/Axelrod-Python/Axelrod

repeated-play

https://github.com/Nikoleta-v3/repeated_play

Share Slides on WeChat!

1. Introduction to repeated games
2. Computer tournaments
3. Axelrod-Python
4. Memory-n strategies
5. repeated-play

   repeated-games

Introduction to repeated games

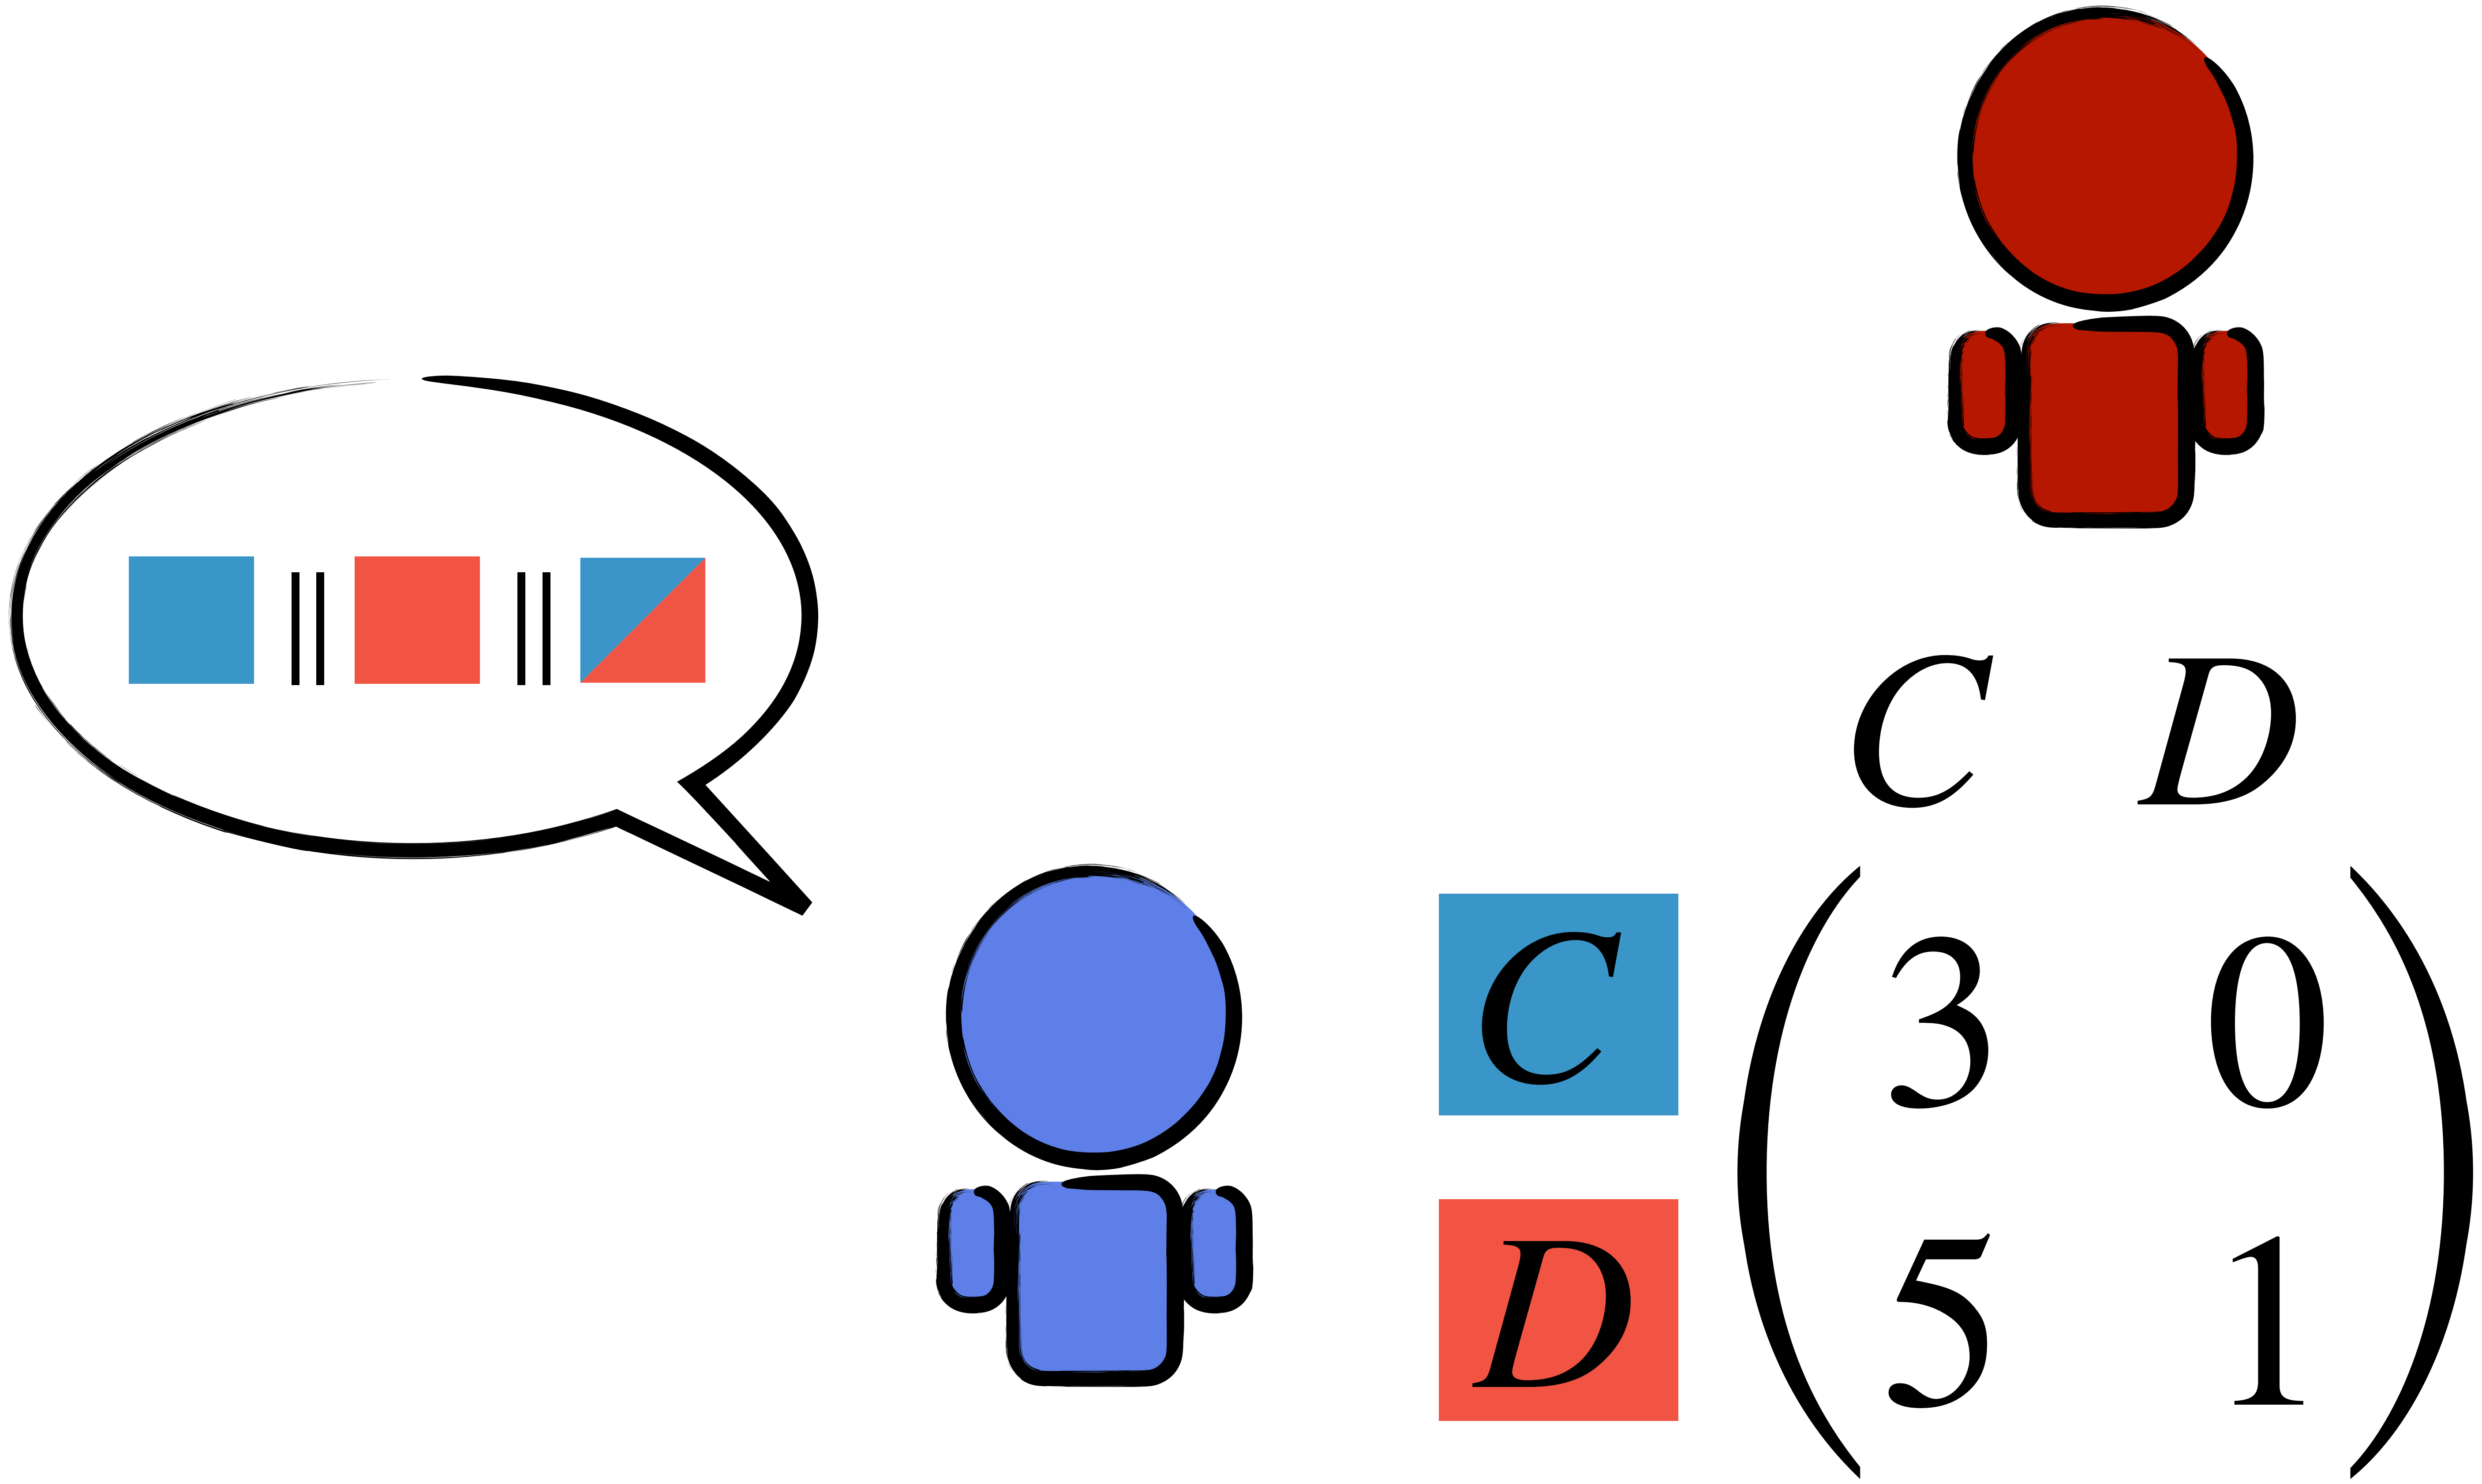

Introduction to repeated games

$$\boxed{1}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \boxed{2}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \boxed{3}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \cdots \quad \boxed{n-1}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \boxed{n}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \boxed{n+1}\ \begin{array}{c|cc} & C & D \\ \hline C & r & s \\ D & t & p \end{array} \quad \infty$$

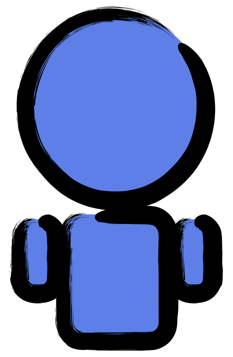	C	D	C	C	C	?
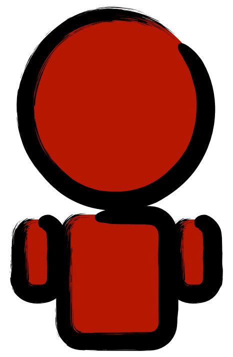	D	C	C	D	C	

What strategy should Blue play to be successful?

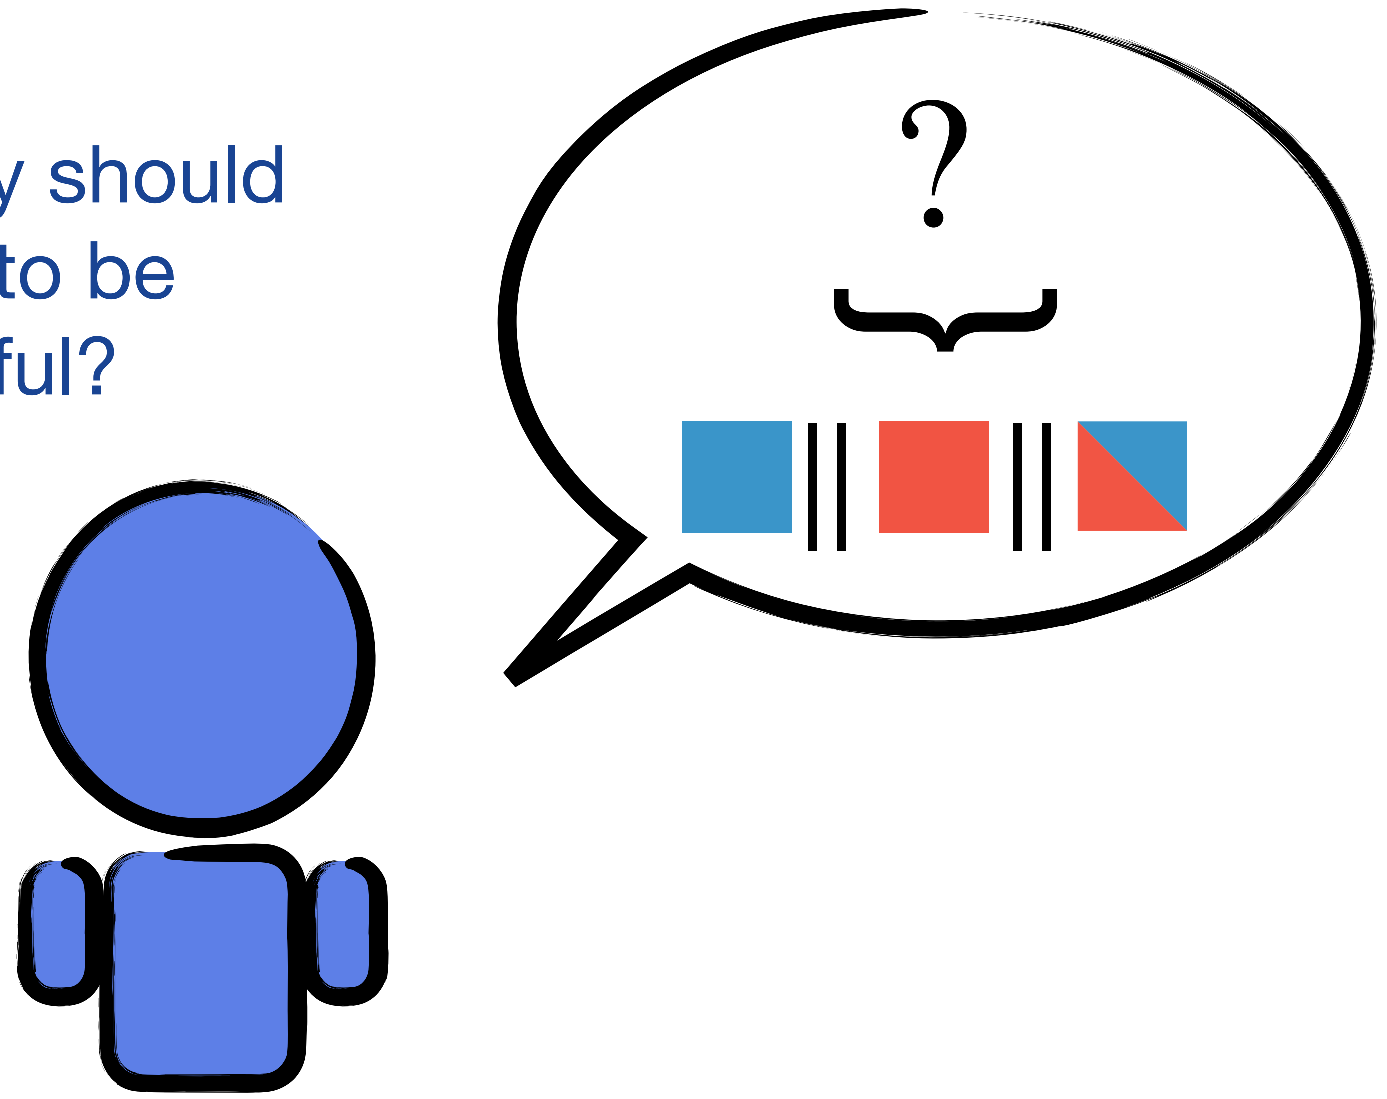

Computer tournaments

Robert Axelrod

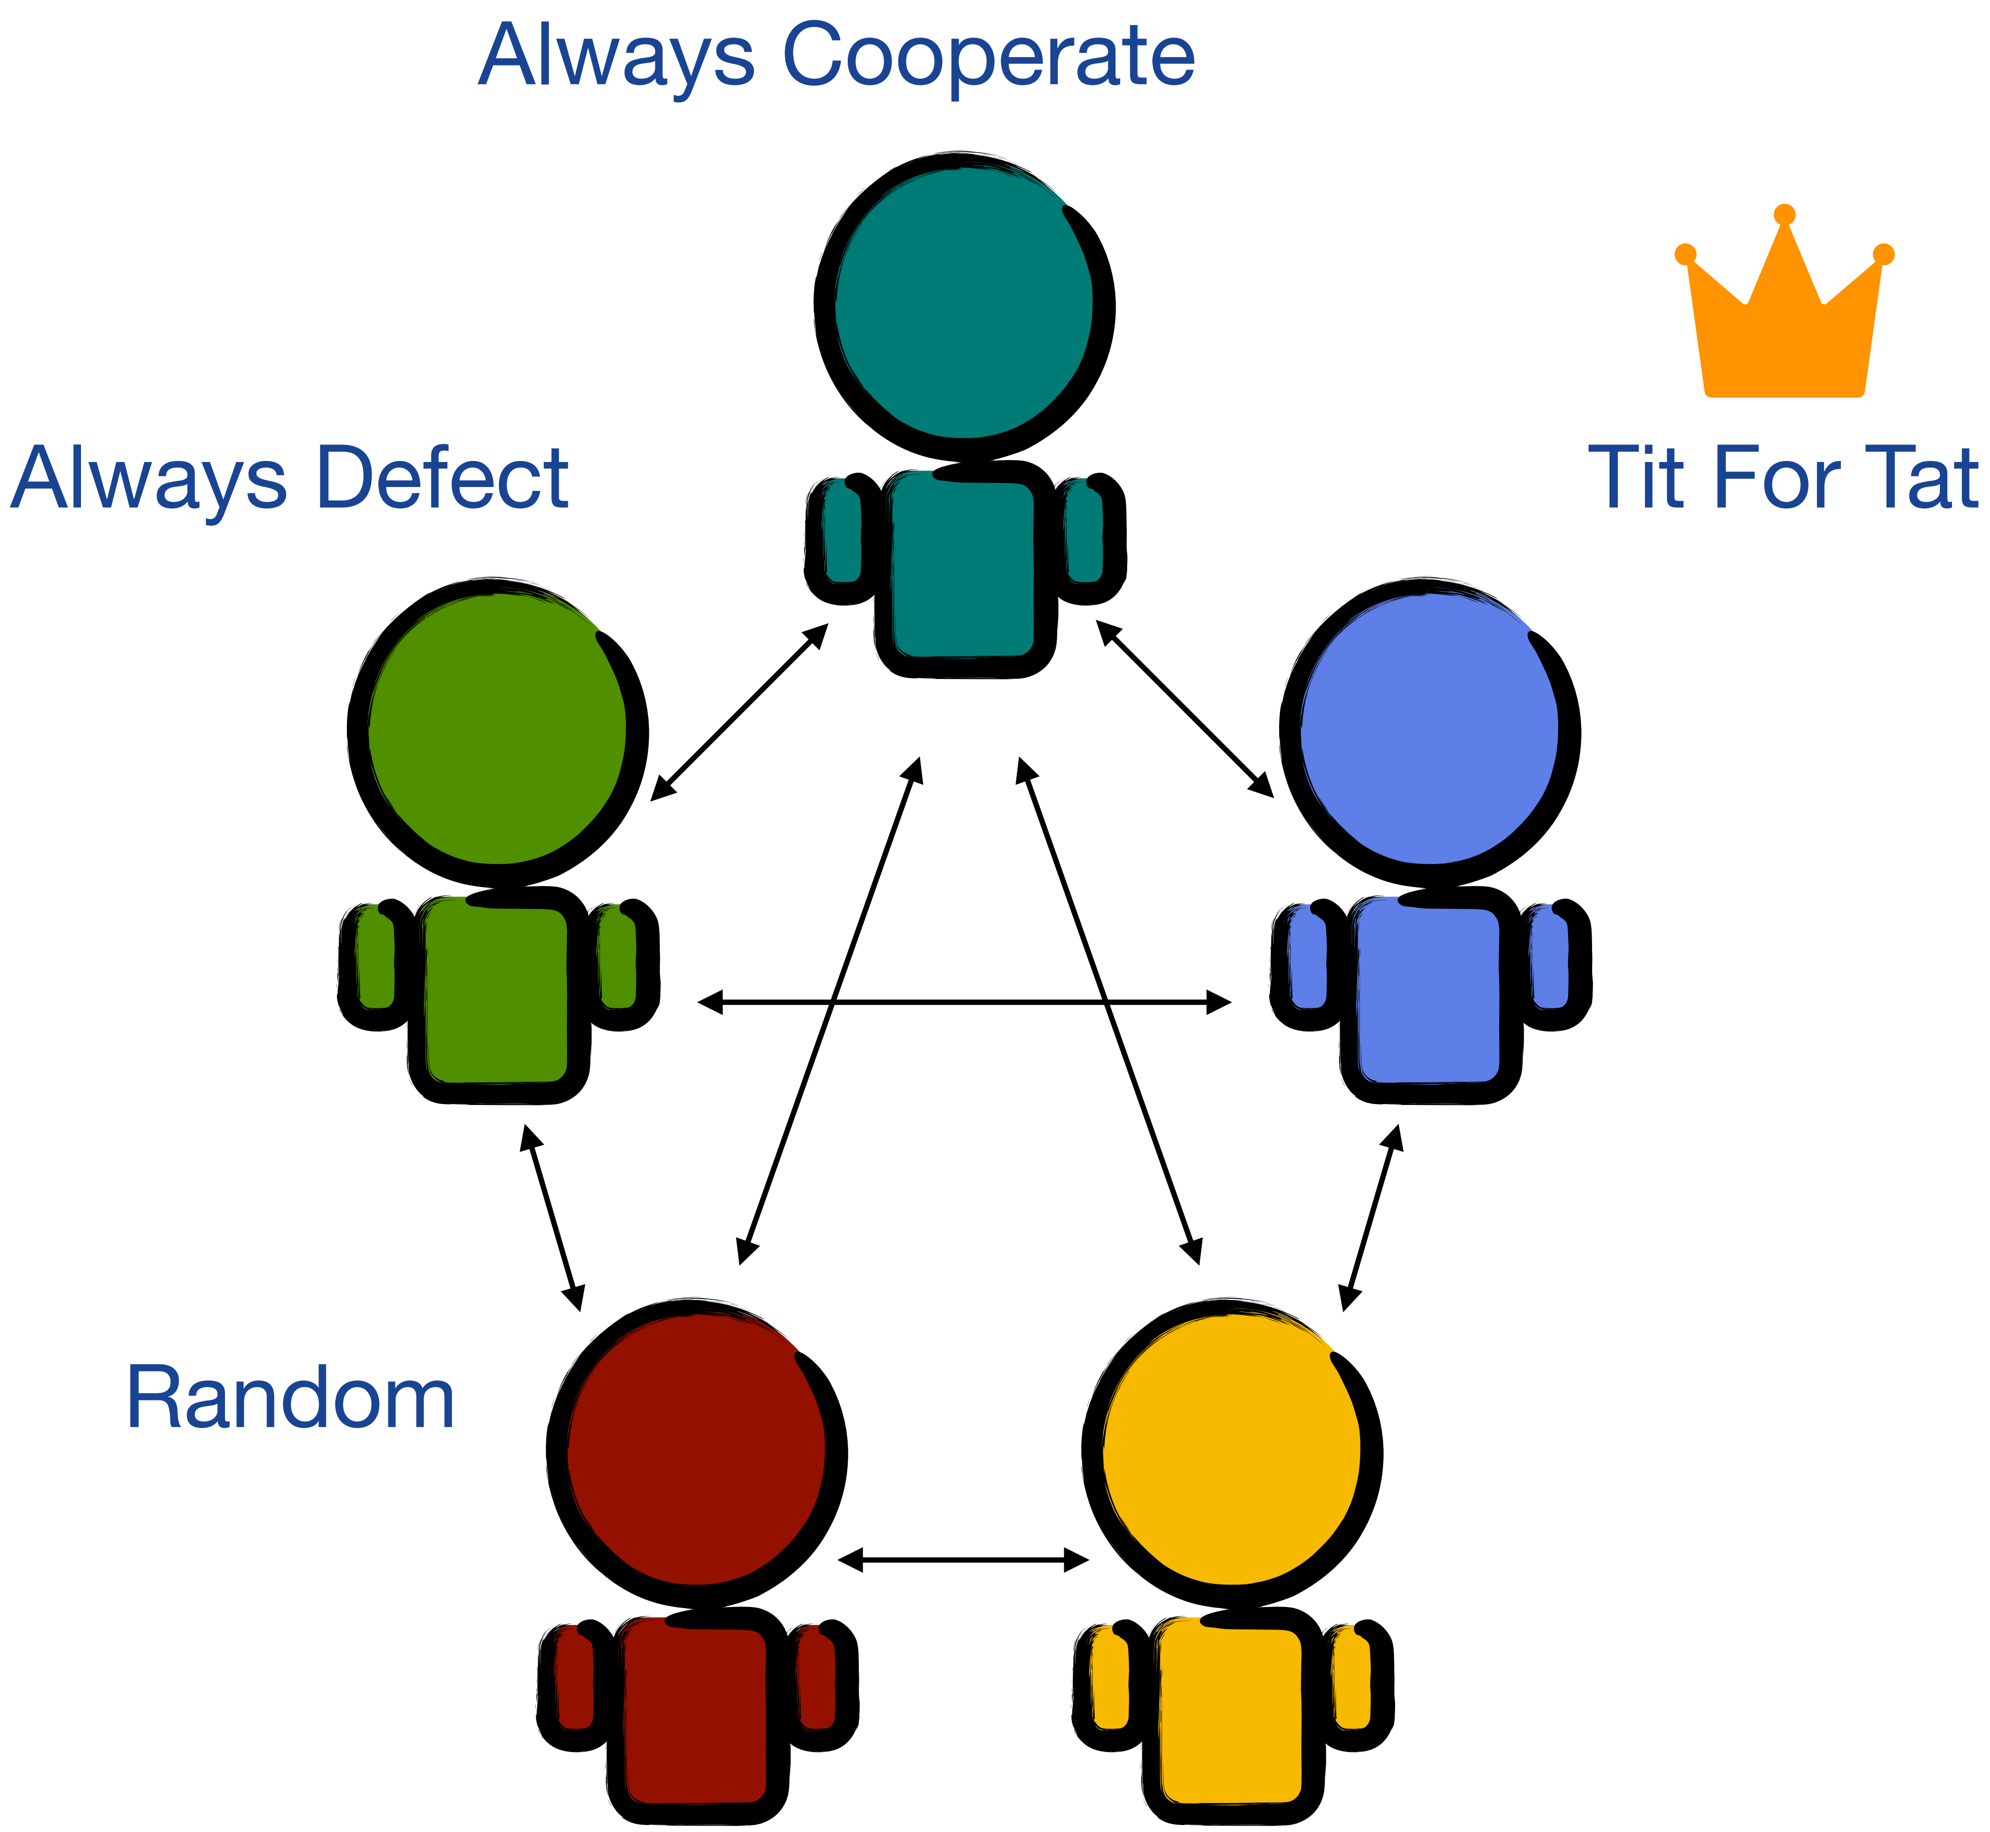

RULES

- Pairwise matches
- Each match plays for a given number of turns
- Tournament is repeated five times
- Strategy with the higher average payoff wins

Computer tournaments

Effective Choice in the Prisoner's Dilemma

ROBERT AXELROD
Institute of Public Policy Studies
University of Michigan

This is a "primer" on how to play the iterated Prisoner's Dilemma game effectively. Existing research approaches offer the participant limited help in understanding how to cope effectively with such interactions. To gain a deeper understanding of how to be effective in such a partially competitive and partially cooperative environment, a computer tournament was conducted for the iterated Prisoner's Dilemma. Decision rules were submitted by entrants who were recruited primarily from experts in game theory from a variety of disciplines: psychology, political science, economics, sociology, and mathematics. The results of the tournament demonstrate that there are subtle reasons for an individualistic pragmatist to cooperate as long as the other side does, to be somewhat forgiving, and to be optimistic about the other side's responsiveness.

More Effective Choice in the Prisoner's Dilemma

ROBERT AXELROD
Institute of Public Policy Studies
The University of Michigan

This study reports and analyzes the results of the second round of the computer tournament for the iterated Prisoner's Dilemma. The object is to gain a deeper understanding of how to perform well in such a setting. The 62 entrants were able to draw lessons from the results of the first round and were able to design their entries to take these lessons into account. The results of the second round demonstrate a number of subtle pitfalls which specific types of decision rules can encounter. The winning rule was once again TIT FOR TAT, the rule which cooperates on the first move and then does what the other player did on the previous move. The analysis of the results shows the value of not being the first to defect, of being somewhat forgiving, but also the importance of being provocable. An analysis of hypothetical alternative tournaments demonstrates the robustness of the results.

A strategy of win–stay, lose–shift that outperforms tit-for-tat in the Prisoner's Dilemma game

Martin Nowak* & Karl Sigmund†

*Department of Zoology, University of Oxford, South Parks Road, Oxford OX1 3PS, UK
†Institut für Mathematik, Universität Wien, Strudlhofgasse 4, A-1090 Vienna, Austria

The Prisoner's Dilemma is the leading metaphor for the evolution of cooperative behaviour in populations of selfish agents, especially since the well-known computer tournaments of Axelrod[1] and their application to biological communities[2,3]. In Axelrod's simulations, the simple strategy tit-for-tat did outstandingly well and subsequently became the major paradigm for reciprocal altruism[4 12]. Here we present extended evolutionary simulations of heterogeneous ensembles of probabilistic strategies including mutation and selection, and report the unexpected success of another protagonist: Pavlov. This strategy is as simple as tit-for-tat and embodies the fundamental behavioural mechanism win–stay, lose–shift, which seems to be a widespread rule[13]. Pavlov's success is based on two important advantages over tit-for-tat: it can correct occasional mistakes and exploit unconditional cooperators. This second feature prevents Pavlov populations from being undermined by unconditional cooperators, which in turn invite defectors. Pavlov seems to be more robust than tit-for-tat, suggesting that cooperative behaviour in natural situations may often be based on win–stay, lose–shift.

The conspicuous success of the tit-for-tat (TFT) strategy (start with a C, and then use your co-players previous move) relies in part on the clinical neatness of a deterministic cyber-world. In natural populations, errors occur[7,12]. TFT suffers from stochastic perturbations in two ways: (1) a TFT population can be 'softened up' by random drift introducing unconditional cooperators, which allow exploiters to grow (TFT is not an evolutionarily stable strategy[15,16]); and (2) occasional mistakes between two TFT players cause long runs of mutual backbiting. (Such mistakes abound in real life: even humans are apt to vent frustrations upon innocent bystanders.)

Within the restricted world of strategies reacting only to the co-players previous move, TFT has a very important, but transitory role: in small clusters, it can invade populations of defectors, but then bows out to a related strategy, 'generous tit for tat' (GTFT), which cooperates after a co-player's C, but also with a certain probability after a D[9].

But as soon as one admits strategies which take into account the moves of both players in the previous round, evolution becomes much less transparent[17]. We first conjectured that GTFT (or variants thereof) would win the day, but are forced to admit, after extensive simulations, that the strategy Pavlov did much better in the long run. A Pavlov player cooperates if and only if both players opted for the same alternative in the previous move. The name[18] stems from the fact that this strategy embodies an almost reflex-like response to the payoff: it repeats its former move if it was rewarded by R or T points, but switches behaviour if it was punished by receiving only P or S points. This strategy, which went by the name of 'simpleton'[19], fares poorly against inveterate defectors: in every second round, it switches to cooperation. It cannot gain a foothold in a defector's world; defectors have to be invaded by other strategies, like TFT[9]. But Pavlov has two important advantages over TFT:

Our Meeting With Gradual: A *Good* Strategy For The Iterated Prisoner's Dilemma

Bruno Beaufils, Jean-Paul Delahaye and Philippe Mathieu
Université des Sciences et Technologies de Lille
Laboratoire d'Informatique Fondamentale de Lille – U.R.A. 369 C.N.R.S.
U.F.R. d'I.E.E.A. Bât. M3 – F-59655 Villeneuve d'Ascq Cedex
e-mail: {beaufils, mathieu, delahaye}@lifl.fr

Abstract

In this paper, after a short return to the description of the classical version of the Iterated Prisoner's Dilemma and its application to the study of cooperation, we present a new strategy we have found named *gradual*, which outperforms the *tit-for-tat* strategy, on which are based a lot of works in the Theory of Cooperation. Since no pure strategy is evolutionarily stable in the IPD, we cannot give a mathematical proof of the absolute superiority of *gradual*, but we present a convergent set of facts that must be viewed as strong experimental evidences of

In section 3 we try to improve the strength of this strategy by using a genetic algorithm, on a genotype we have created and which includes lots of well-known strategies (in fact our genotype can cover more than 8×10^{15} strategies). We present our ideas on a tree representation of the strategies space and finally we propose a new view of evolution of cooperation in which complexity plays a major role.

In the last sections we describe our results and we discuss about the natural behavior of this strategy and its good robustness in ecological competitions.

1.1 IPD and Artificial Life

OPEN

Check for updates

Five rules for friendly rivalry in direct reciprocity

Yohsuke Murase[1] & Seung Ki Baek[2]✉

Direct reciprocity is one of the key mechanisms accounting for cooperation in our social life. According to recent understanding, most of classical strategies for direct reciprocity fall into one of two classes, 'partners' or 'rivals'. A 'partner' is a generous strategy achieving mutual cooperation, and a 'rival' never lets the co-player become better off. They have different working conditions: For example, partners show good performance in a large population, whereas rivals do in head-to-head matches. By means of exhaustive enumeration, we demonstrate the existence of strategies that act as both partners and rivals. Among them, we focus on a human-interpretable strategy, named 'CAPRI' after its five characteristic ingredients, i.e., cooperate, accept, punish, recover, and defect otherwise. Our evolutionary simulation shows excellent performance of CAPRI in a broad range of environmental conditions.

Article
https://doi.org/10.1038/s43588-022-00334-w

Evolution of cooperation through cumulative reciprocity

Received: 6 October 2021
Accepted: 14 September 2022
Published online: 20 October 2022
Check for updates

Juan Li[1,2], Xiaowei Zhao[1,3], Bing Li[1], Charlotte S. L. Rossetti[4], Christian Hilbe[4,5]✉ and Haoxiang Xia[1,2,5]✉

Reciprocity is a simple principle for cooperation that explains many of the patterns of how humans seek and receive help from each other. To capture reciprocity, traditional models often assume that individuals use simple strategies with restricted memory. These memory-1 strategies are mathematically convenient, but they miss important aspects of human reciprocity, where defections can have lasting effects. Here we instead propose a strategy of cumulative reciprocity. Cumulative reciprocators count the imbalance of cooperation across their previous interactions with their opponent. They cooperate as long as this imbalance is sufficiently small. Using analytical and computational methods, we show that this strategy can sustain cooperation in the presence of errors, that it enforces fair outcomes and that it evolves in hostile environments. Using an economic experiment, we confirm that cumulative reciprocity is more predictive of human behaviour than several classical strategies. The basic principle of cumulative reciprocity is versatile and can be extended to a range of social dilemmas.

Axelrod-Python

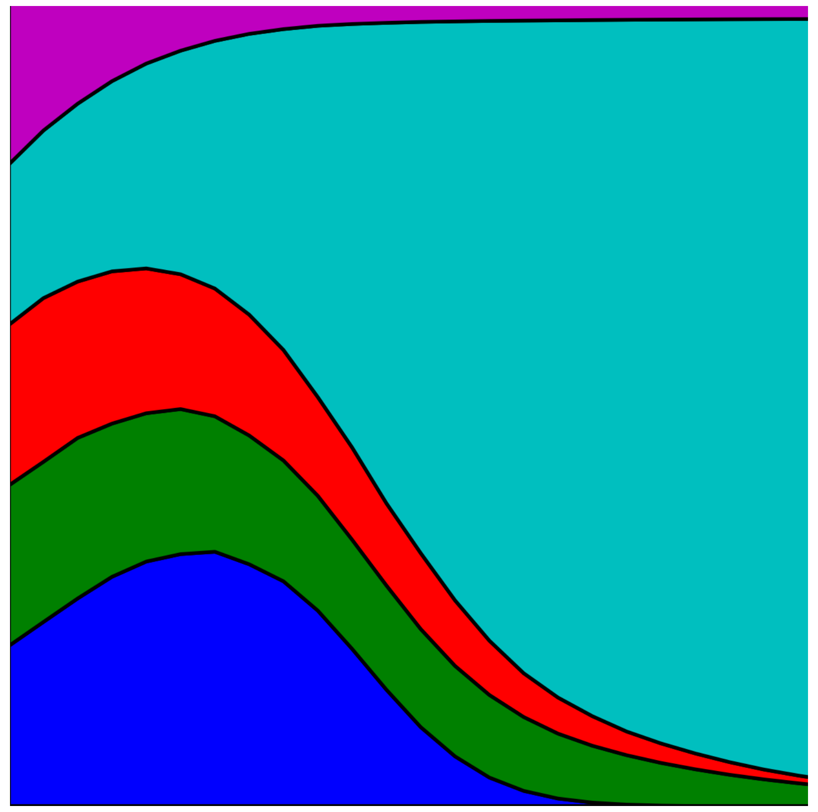

https://axelrod.readthedocs.io/en/stable/

Help desk

Axelrod-Python

Game-theory-Python

https://github.com/Nikoleta-v3/Game-Theory-and-Python

Axelrod-Python

Article https://doi.org/10.1038/s43588-022-00334-w

Evolution of cooperation through cumulative reciprocity

Received: 6 October 2021

Accepted: 14 September 2022

Published online: 20 October 2022

Check for updates

Juan Li[1,2], Xiaowei Zhao[1,3], Bing Li[1], Charlotte S. L. Rossetti[4], Christian Hilbe[4,5] and Haoxiang Xia[1,2,5]

Reciprocity is a simple principle for cooperation that explains many of the patterns of how humans seek and receive help from each other. To capture reciprocity, traditional models often assume that individuals use simple strategies with restricted memory. These memory-1 strategies are mathematically convenient, but they miss important aspects of human reciprocity, where defections can have lasting effects. Here we instead propose a strategy of cumulative reciprocity. Cumulative reciprocators count the imbalance of cooperation across their previous interactions with their opponent. They cooperate as long as this imbalance is sufficiently small. Using analytical and computational methods, we show that this strategy can sustain cooperation in the presence of errors, that it enforces fair outcomes and that it evolves in hostile environments. Using an economic experiment, we confirm that cumulative reciprocity is more predictive of human behaviour than several classical strategies. The basic principle of cumulative reciprocity is versatile and can be extended to a range of social dilemmas.

```
"""
CURE (Cumulative Reciprocity) strategy.

According to this strategy, a player cooperates unless the opponent has defected
significantly more often than they have, exceeding a specified tolerance level.

Parameters:
- tolerance (int): The maximum difference in defections the player is willing to tolerate.
"""
```

Axelrod-Python

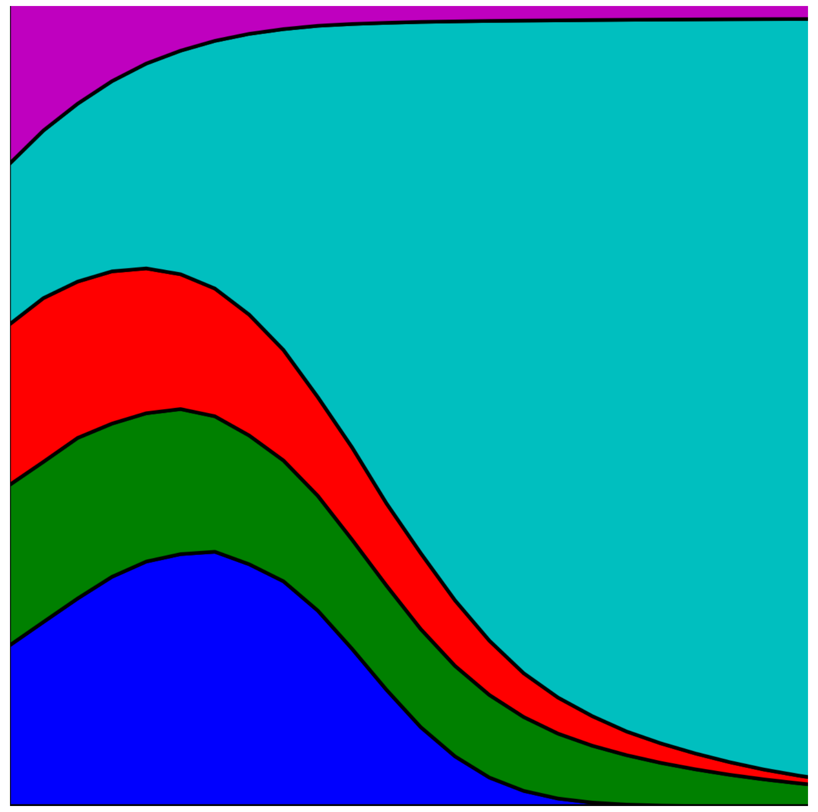

https://axelrod.readthedocs.io/en/stable/

Help desk

Axelrod-Python

https://github.com/Axelrod-Python/Axelrod

Game-theory-Python

https://github.com/Nikoleta-v3/Game-Theory-and-Python

1 Introduction to direct reciprocity/repeated games

2 Computer tournaments

3 Axelrod-Python

1

4 Memory-n strategies

repeated-play

5 repeated-games

Memory-n

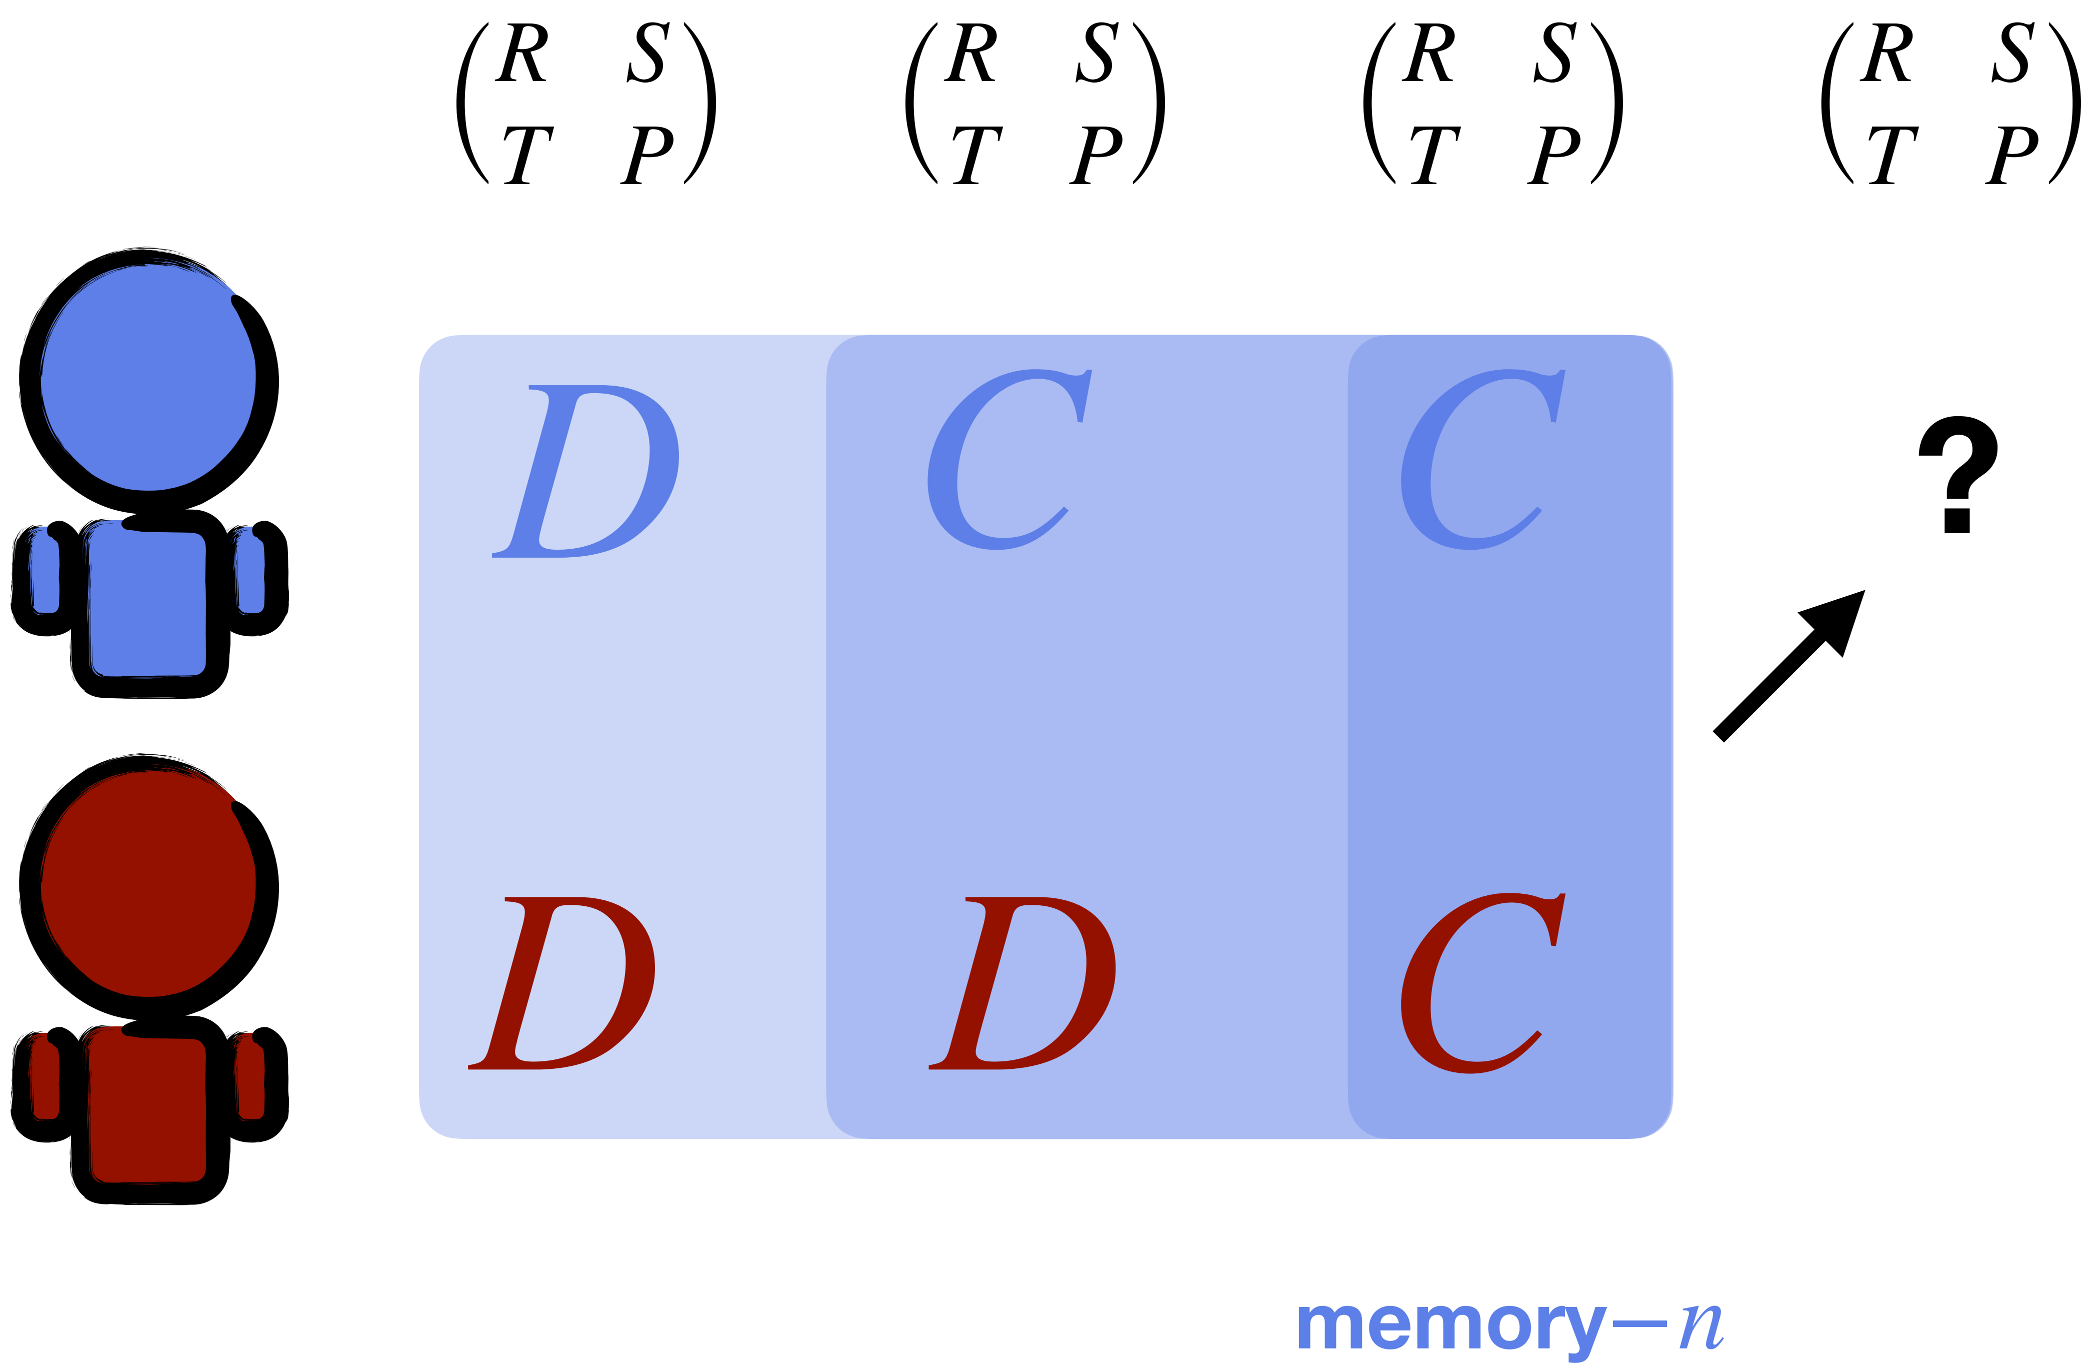
$\begin{pmatrix} R & S \\ T & P \end{pmatrix}$ $\begin{pmatrix} R & S \\ T & P \end{pmatrix}$ $\begin{pmatrix} R & S \\ T & P \end{pmatrix}$ $\begin{pmatrix} R & S \\ T & P \end{pmatrix}$
D C C
D D C
?
memory—n

repeated-play

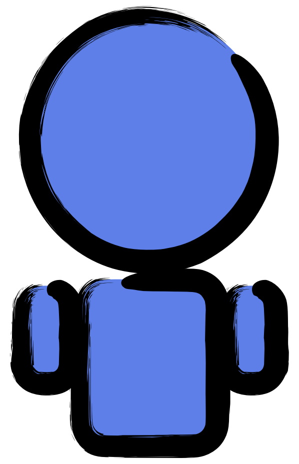

$\mathbf{p} = (p_{CC}, p_{CD}, p_{DC}, p_{DD})$

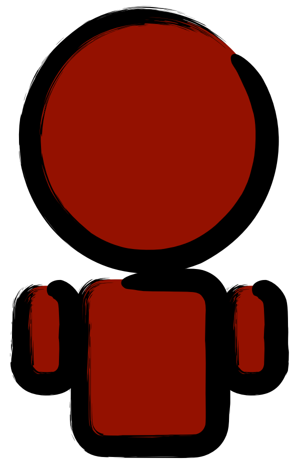

$\mathbf{q} = (q_{CC}, q_{CD}, q_{DC}, q_{DD})$

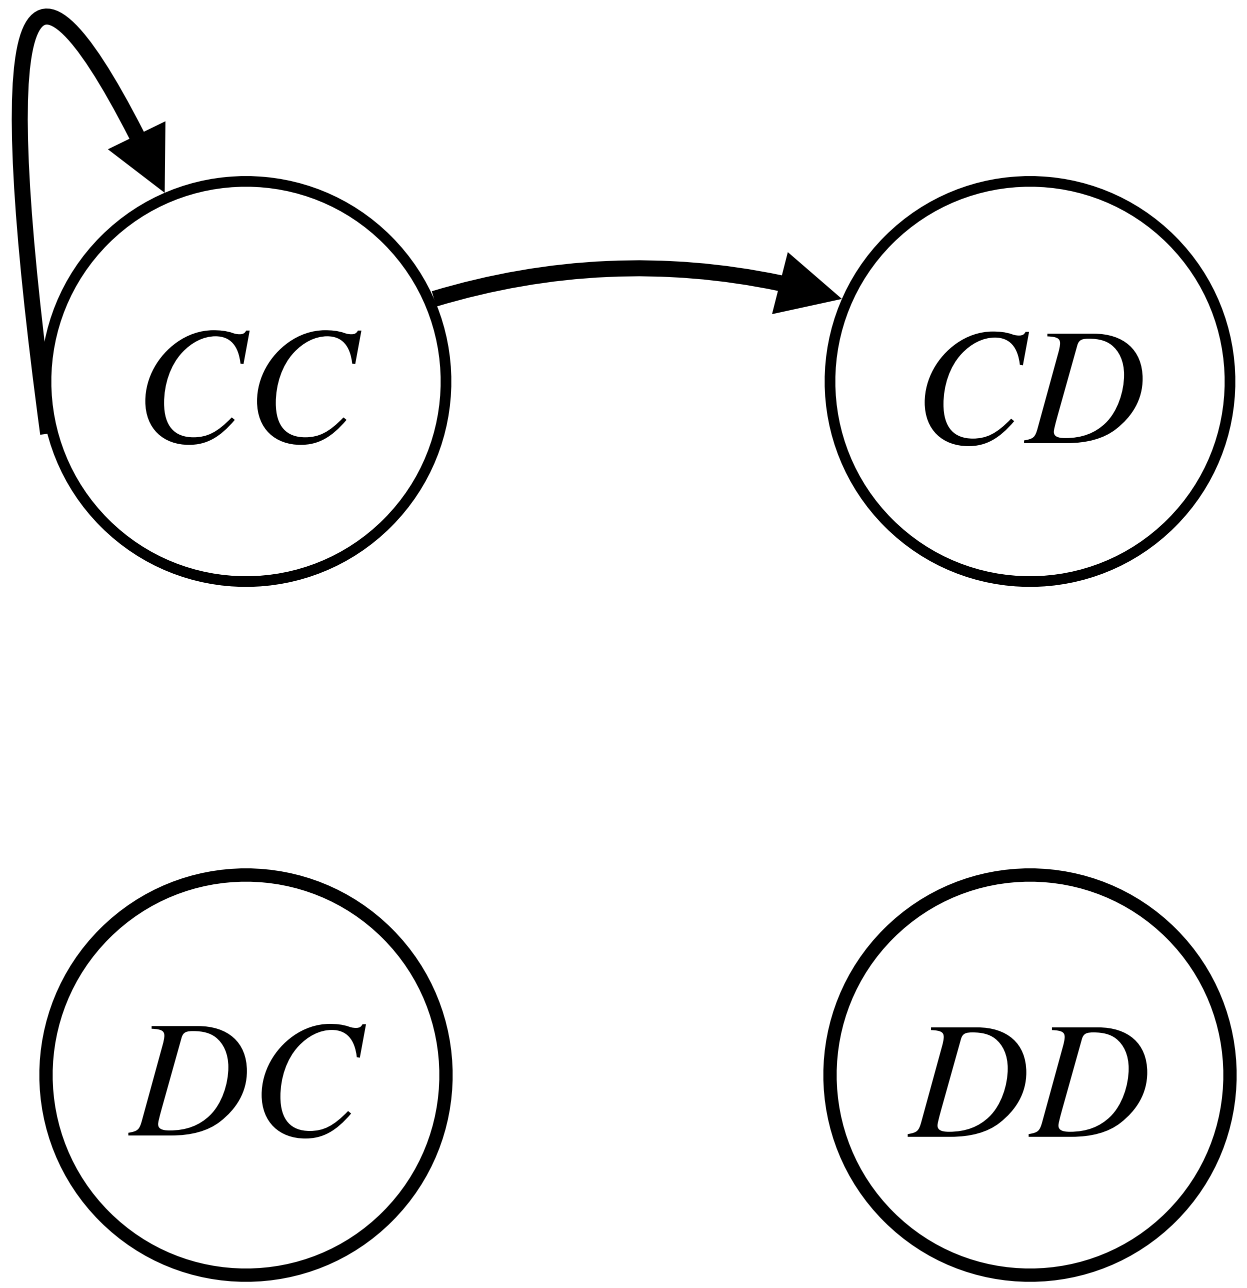

repeated-play

repeated-play

https://github.com/Nikoleta-v3/repeated_play

Help desk

repeated-play

repeated-play

https://github.com/Nikoleta-v3/repeated_play

Summary

Axelrod-Python

https://github.com/Axelrod-Python/Axelrod

repeated-play

https://github.com/Nikoleta-v3/repeated_play

The Journal of Open Source Education

Game Theory and Python: An educational tutorial to game theory and repeated games using Python

Jupyter Notebook | Python | Submitted 05 December 2019 • Published 24 May 2021

JOSE
The Journal of Open Source Education

Game Theory and Python: An educational tutorial to game theory and repeated games using Python

Nikoleta E. Glynatsi[1,2] and Vincent A. Knight[1]

1 Department of Mathematics, Cardiff University, Senghennydd Rd, Cardiff CF24 4AG 2 Max Planck Institute for Evolutionary Biology, August-Thienemann-Straße 2, 24306 Plön

DOI: 10.21105/jose.00078

Software
- Review
- Repository
- Archive

Submitted: 05 December 2019
Published: 24 May 2021

License
Authors of papers retain copyright and release the work under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

Summary

These materials are an open-source educational tutorial aimed at introducing participants to game theory and more specifically to repeated games. The tutorial uses the open-source library called Axelrod-Python and Jupyter Notebooks making this an open-source, reproducible and interactive tutorial.

Statement of Need

We are mathematicians and Python programmers with interests in game theory and peda-

Software repository
Paper review
Download paper
Software archive

Review
Editor: @labarba (all papers)
Reviewers: @sconde (all reviews),

Game-theory-Python

https://github.com/Nikoleta-v3/Game-Theory-and-Python

Nikoleta-v3

THANK YOU!

JOSE **JOSS**

Share Slides on WeChat!

Other open-source tools

NashPy

https://github.com/drvinceknight/Nashpy

EGTtools

https://github.com/Socrats/EGTTools