Multi-Agent Systems and Evolution

Day 1: Introduction to Game Theory

Elias Fernández Domingos

We want you to participate!

We want you to participate!

Notebooks with examples

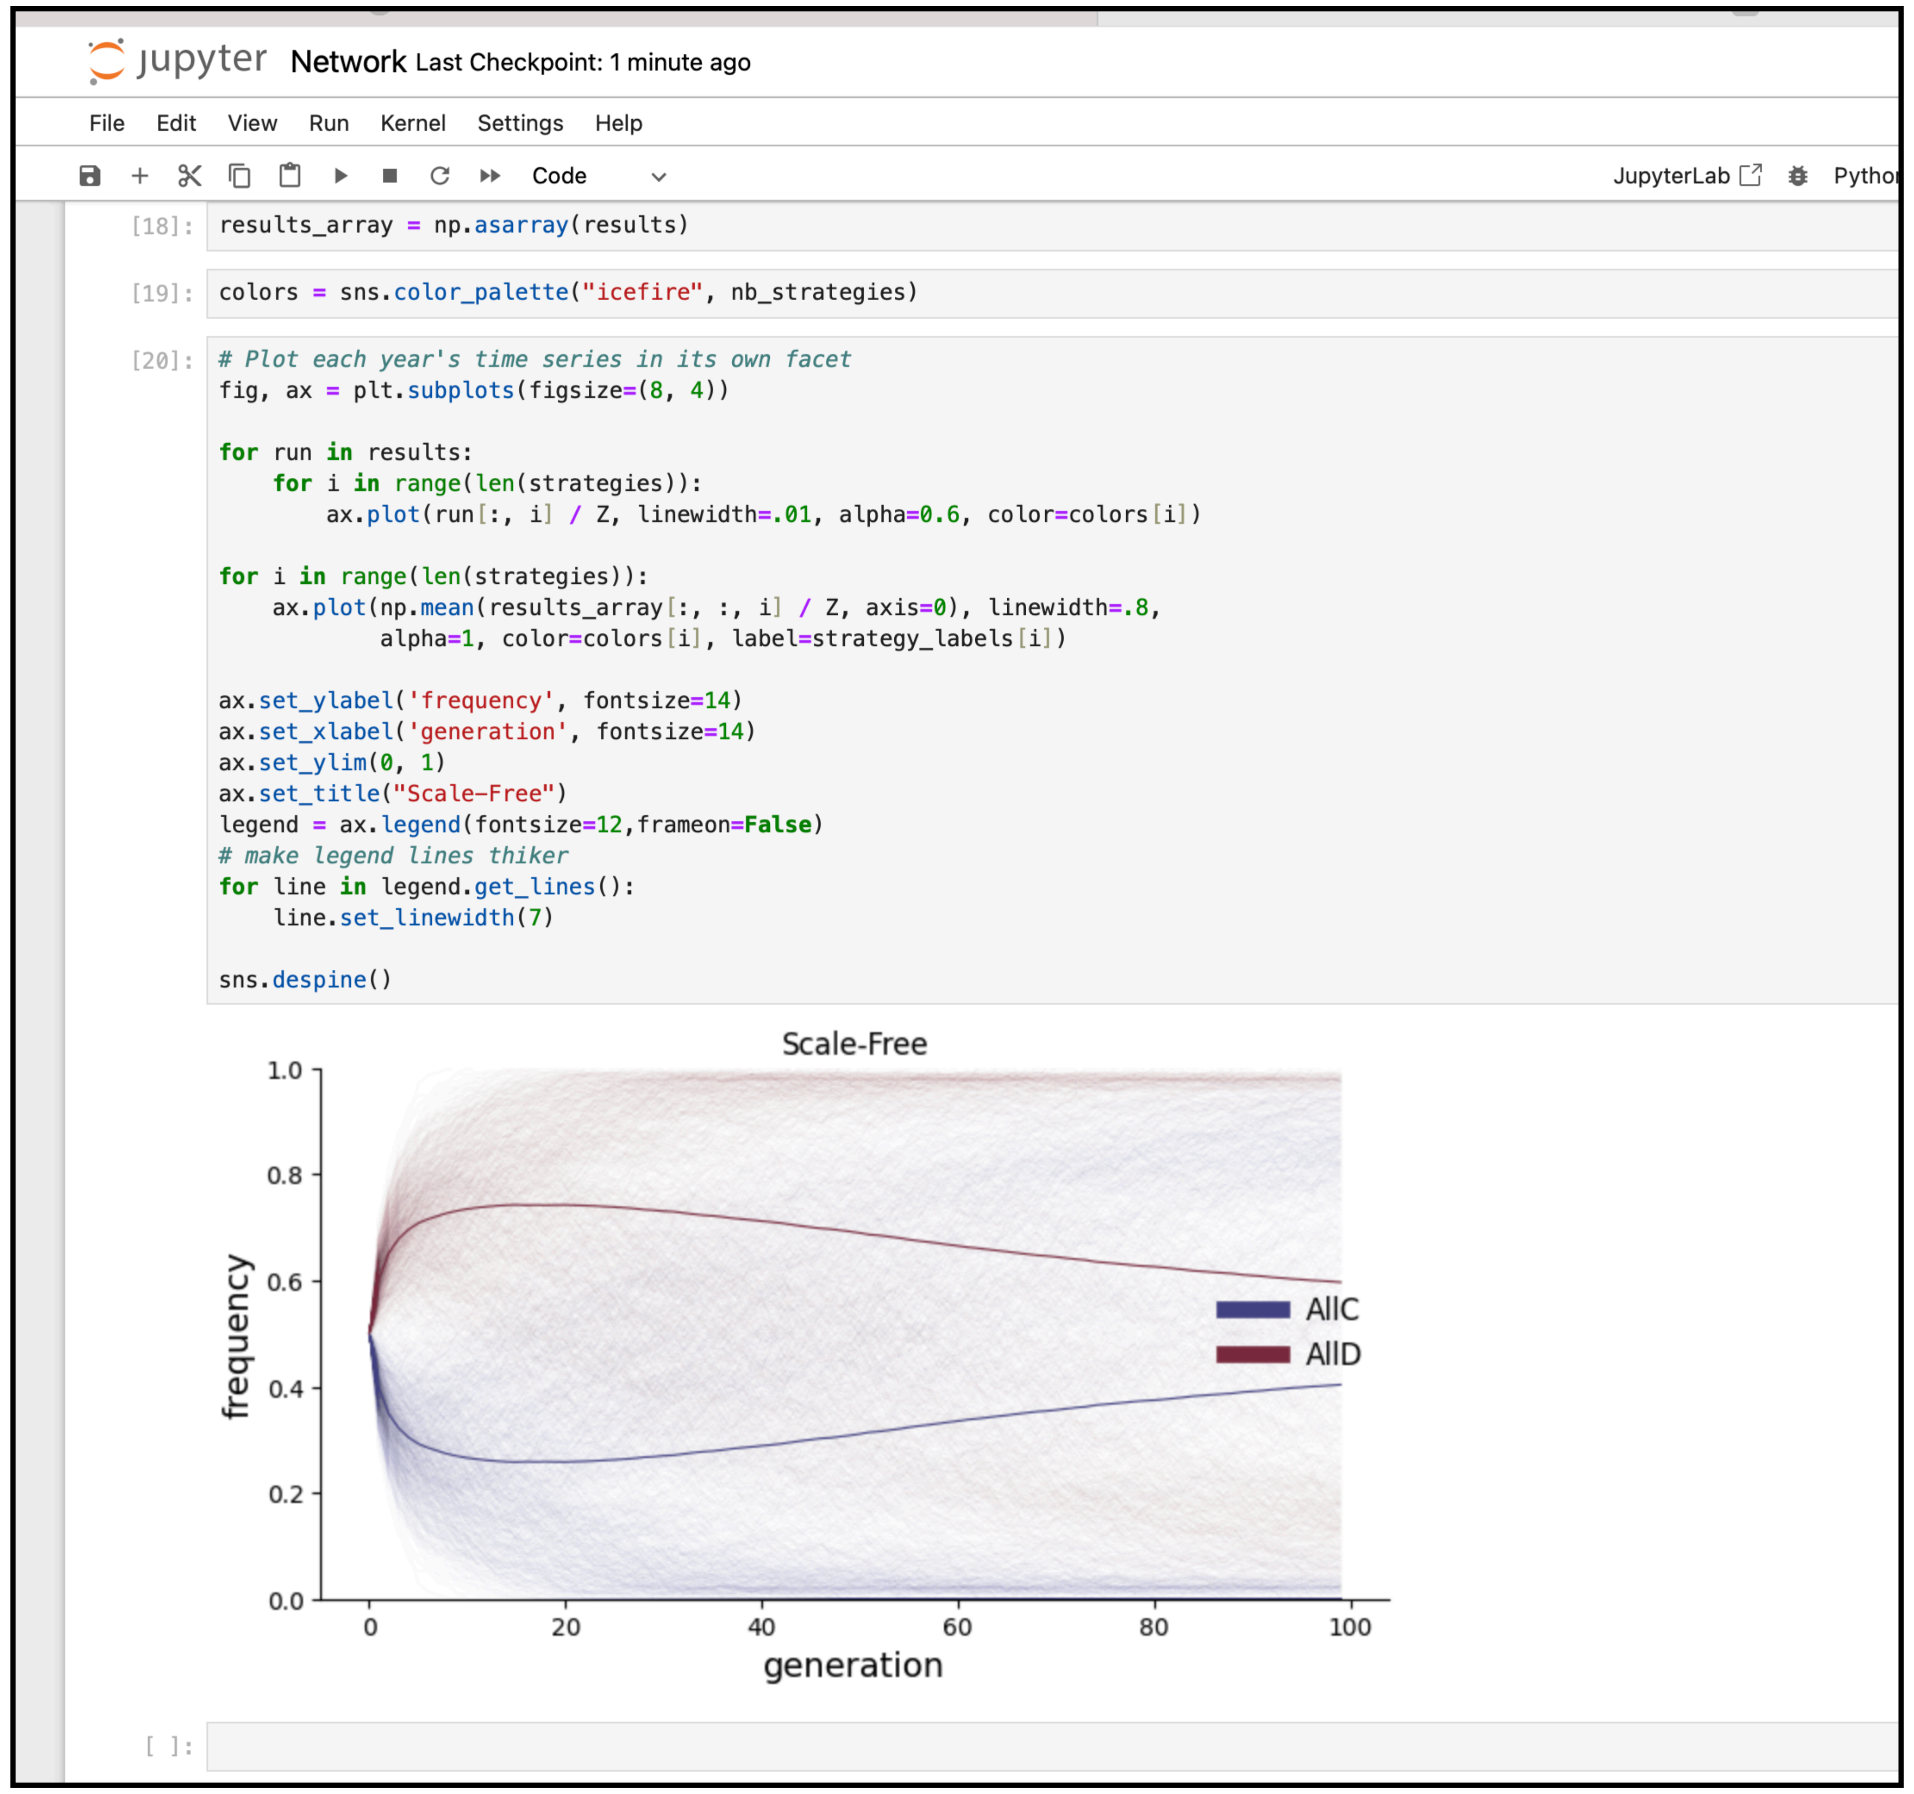

Group project

Outline of the course

- **Day 1: Introduction to Game Theory**
- Day 2: Evolutionary Game Theory
- Day 3: Games on Networks
- Day 4: Practical challenges and connecting theory to Behavioural Experiments
- Day 5: Final remarks and Project presentations

Day 1: Introduction to Game Theory

1. Game Theory, Social Dynamics and Artificial Intelligence
2. Introduction to Game Theory
3. Description of Projects

Part 1: Social Dynamics, Game Theory and Artificial Intelligence

Social dilemmas and collective risk

Climate action

Group hunting

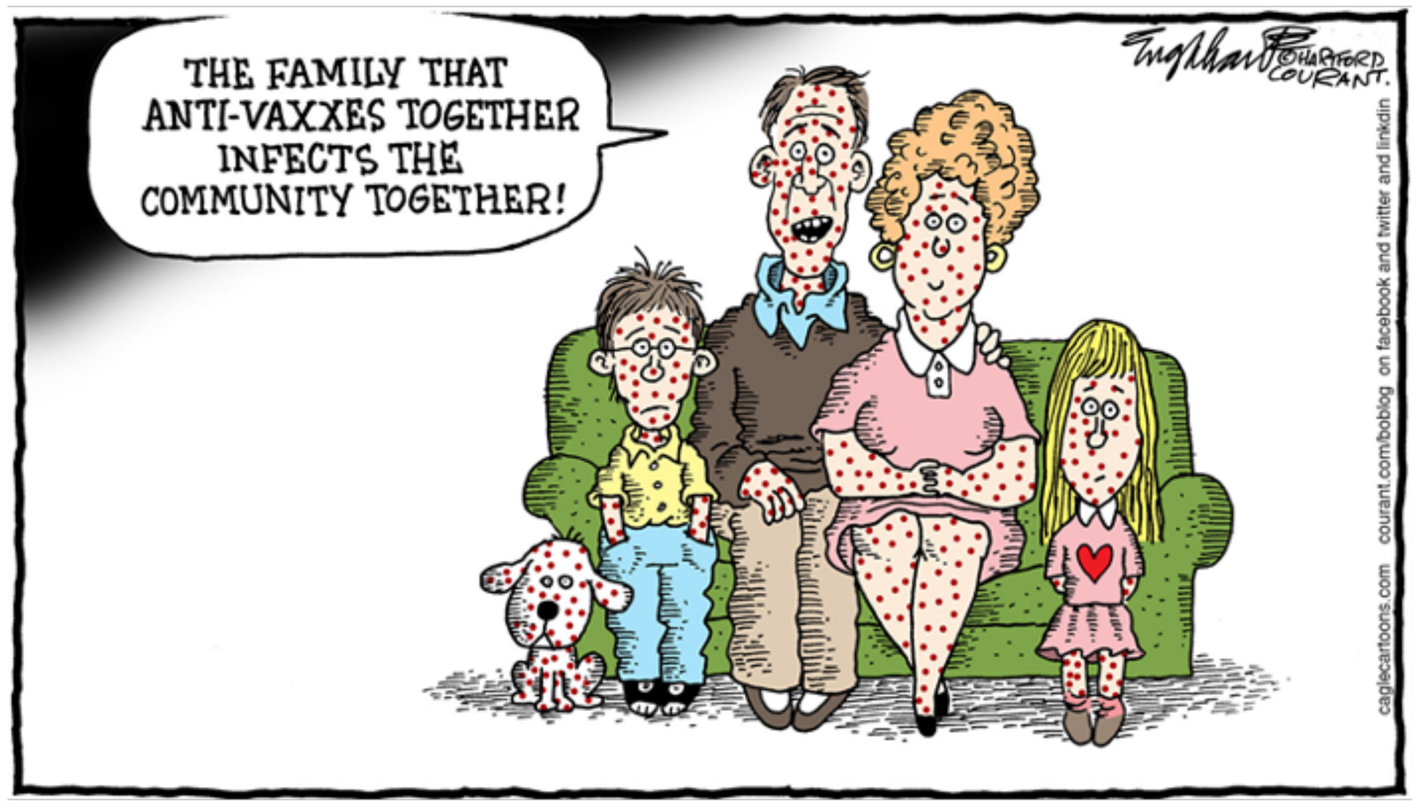

Vaccination resistance

Abuse of antibiotics

Complex strategic interactions

Actors

Complex strategic interactions

Actors

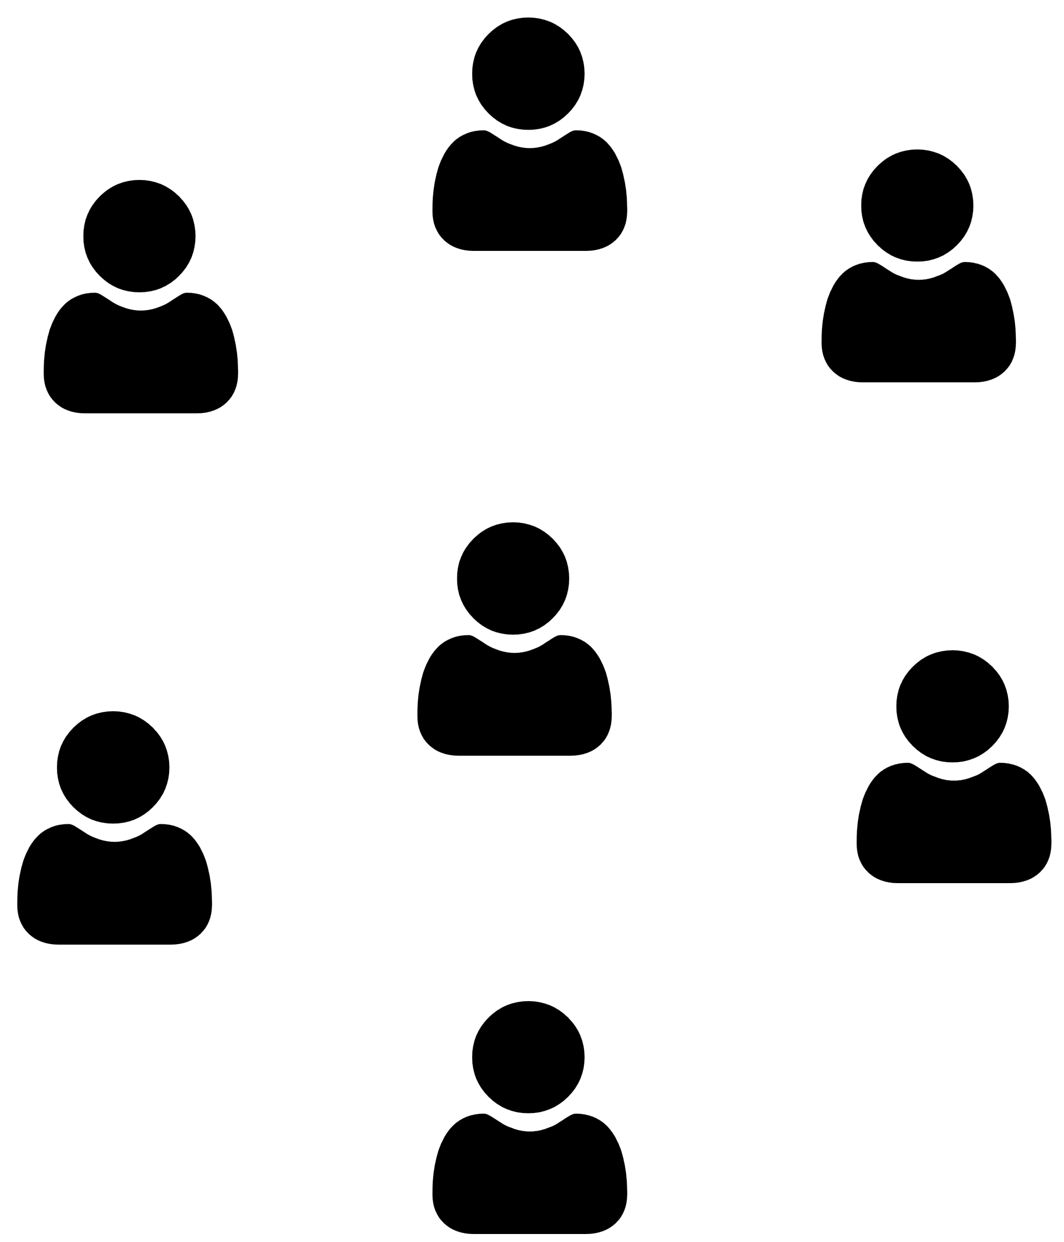

Complex strategic interactions

Actors

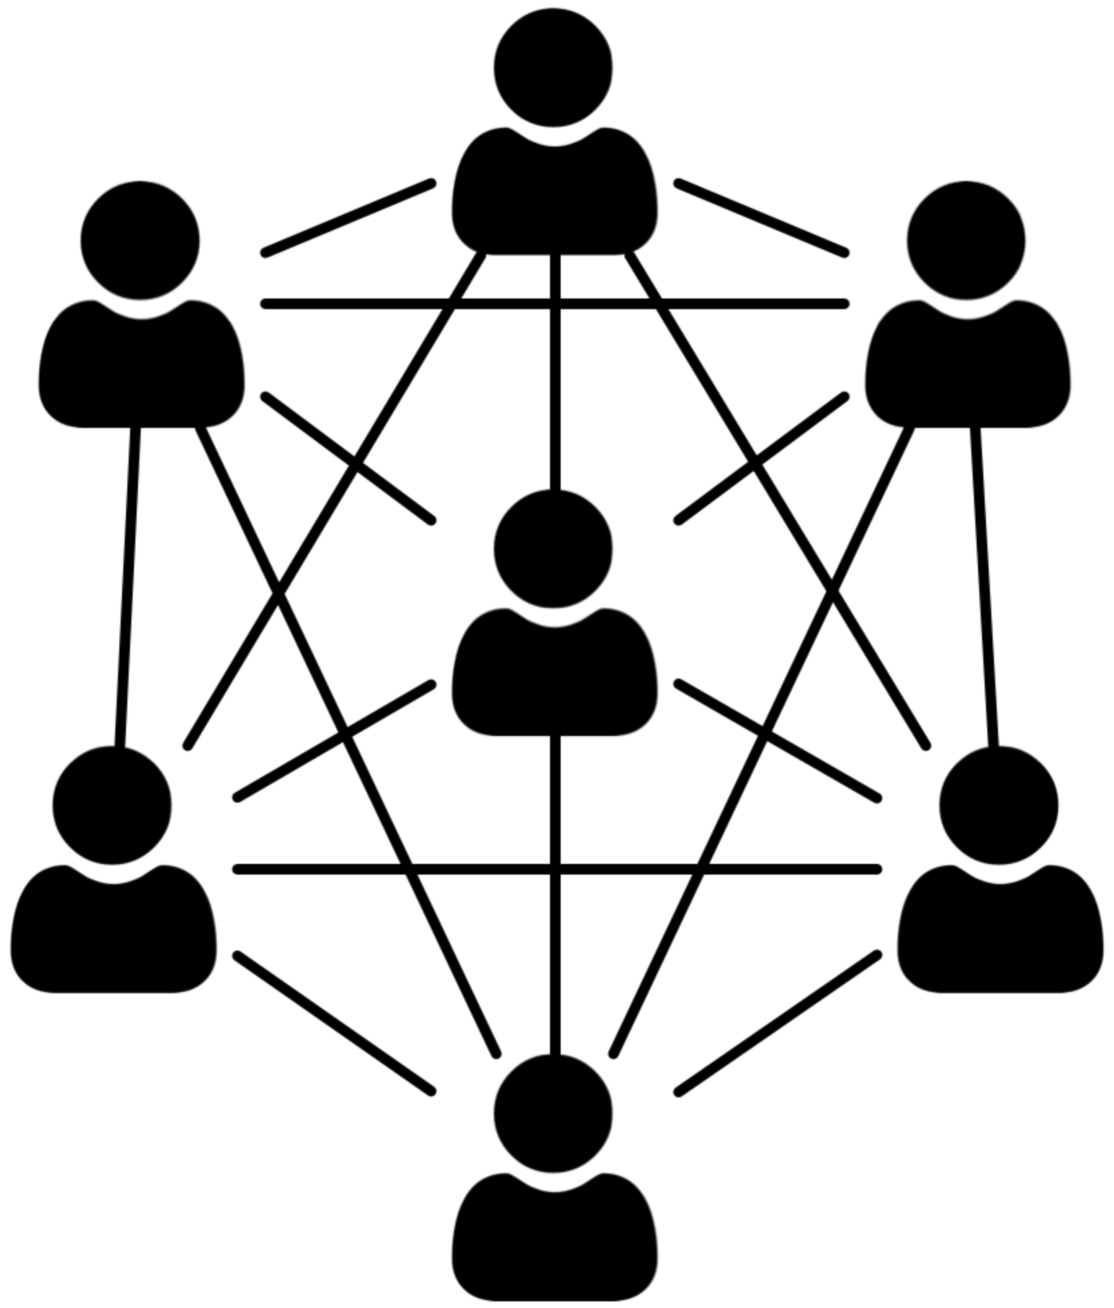

Complex strategic interactions

Actors

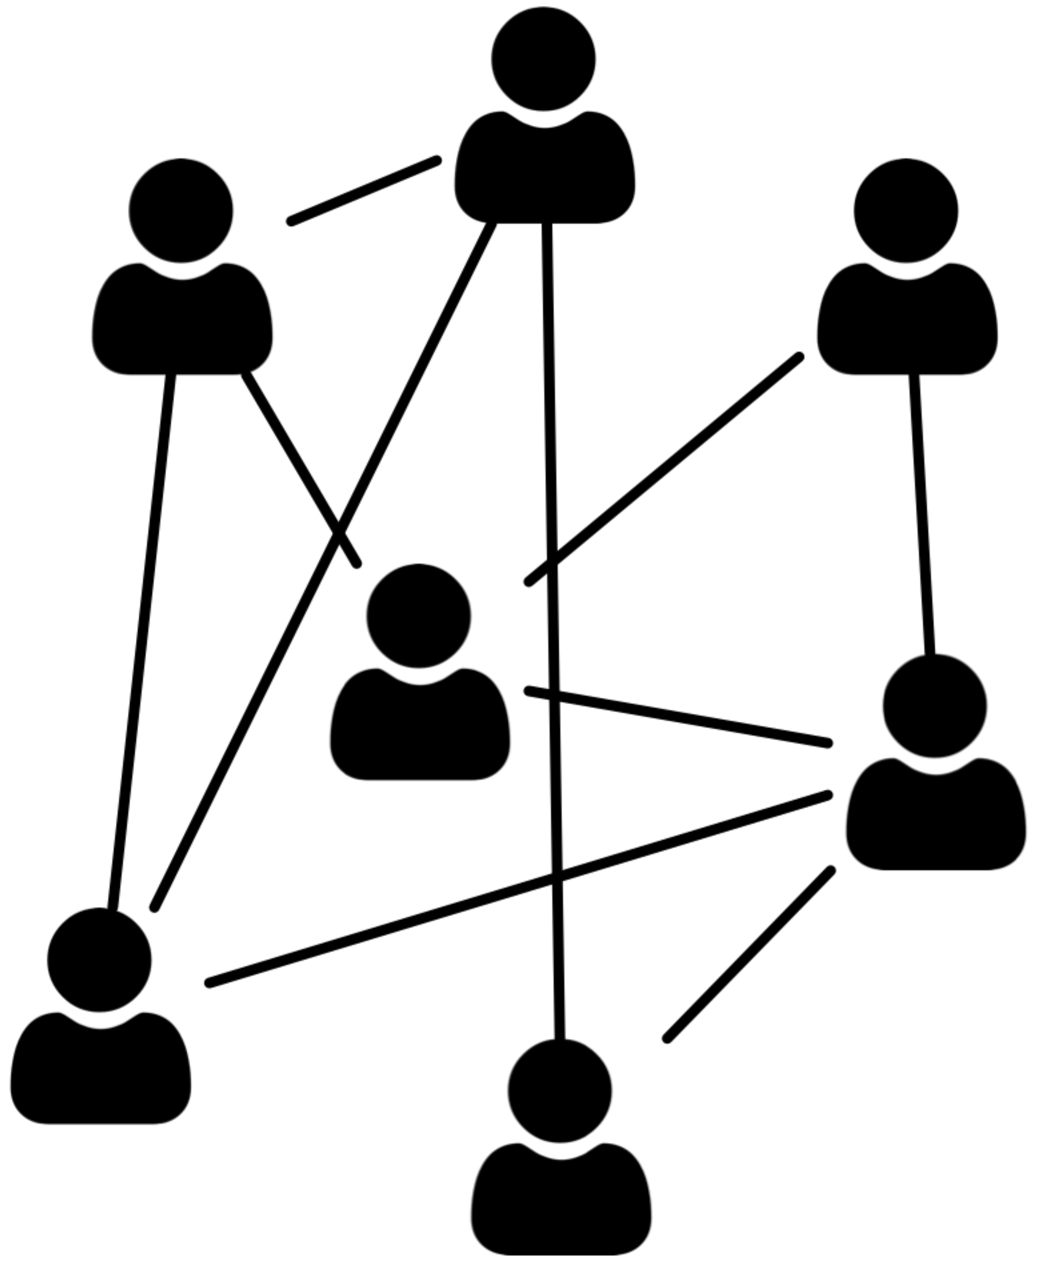

Complex strategic interactions

Actors

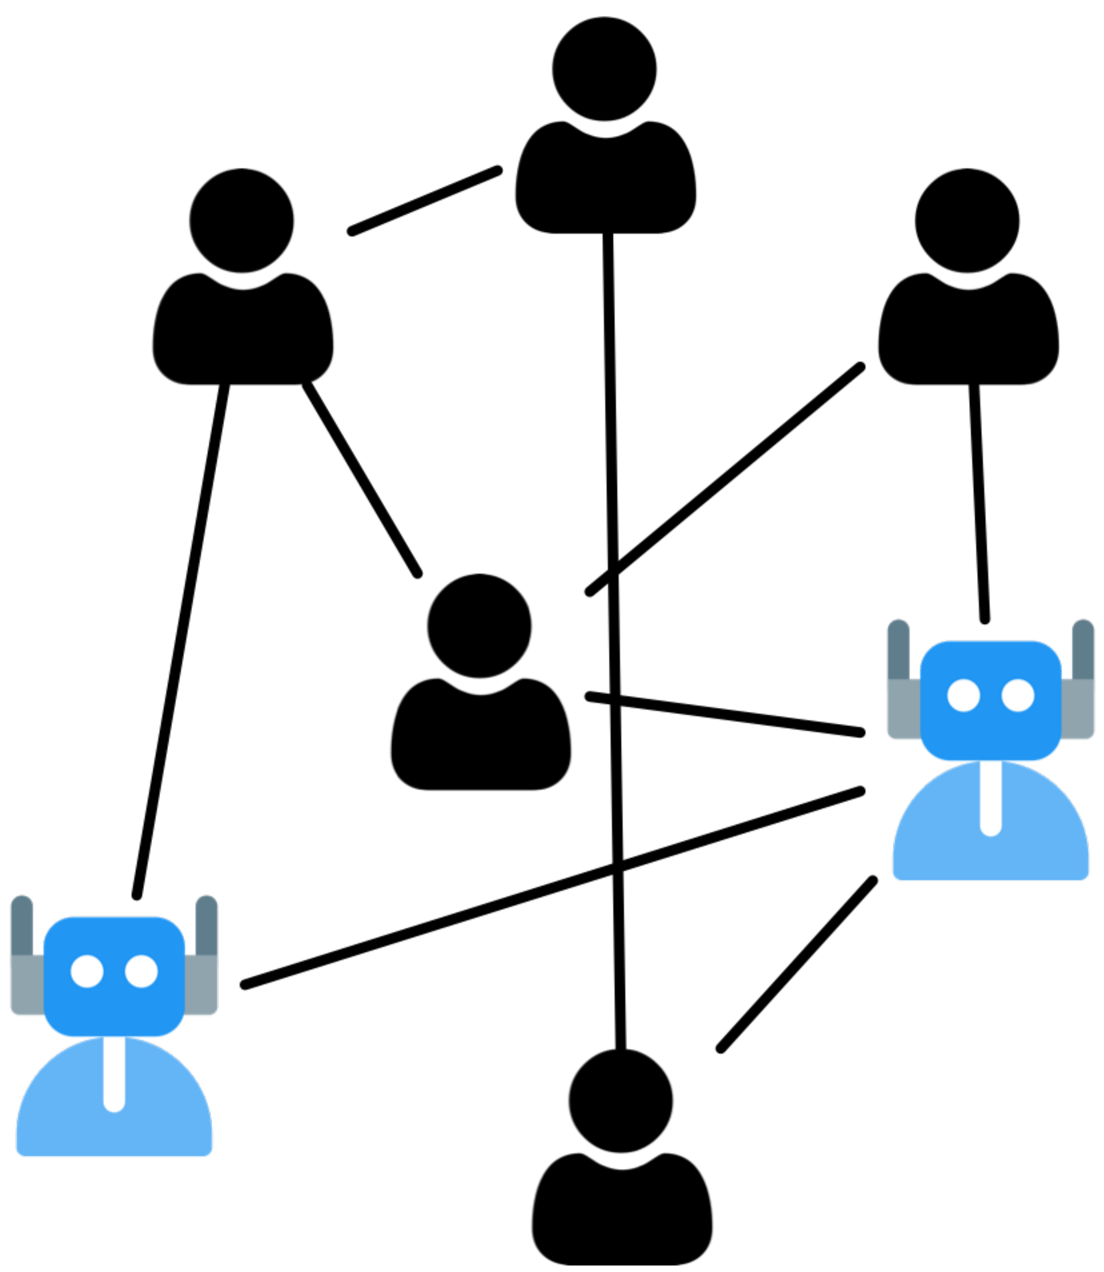

Complex strategic interactions

Actors **Game/Environment**

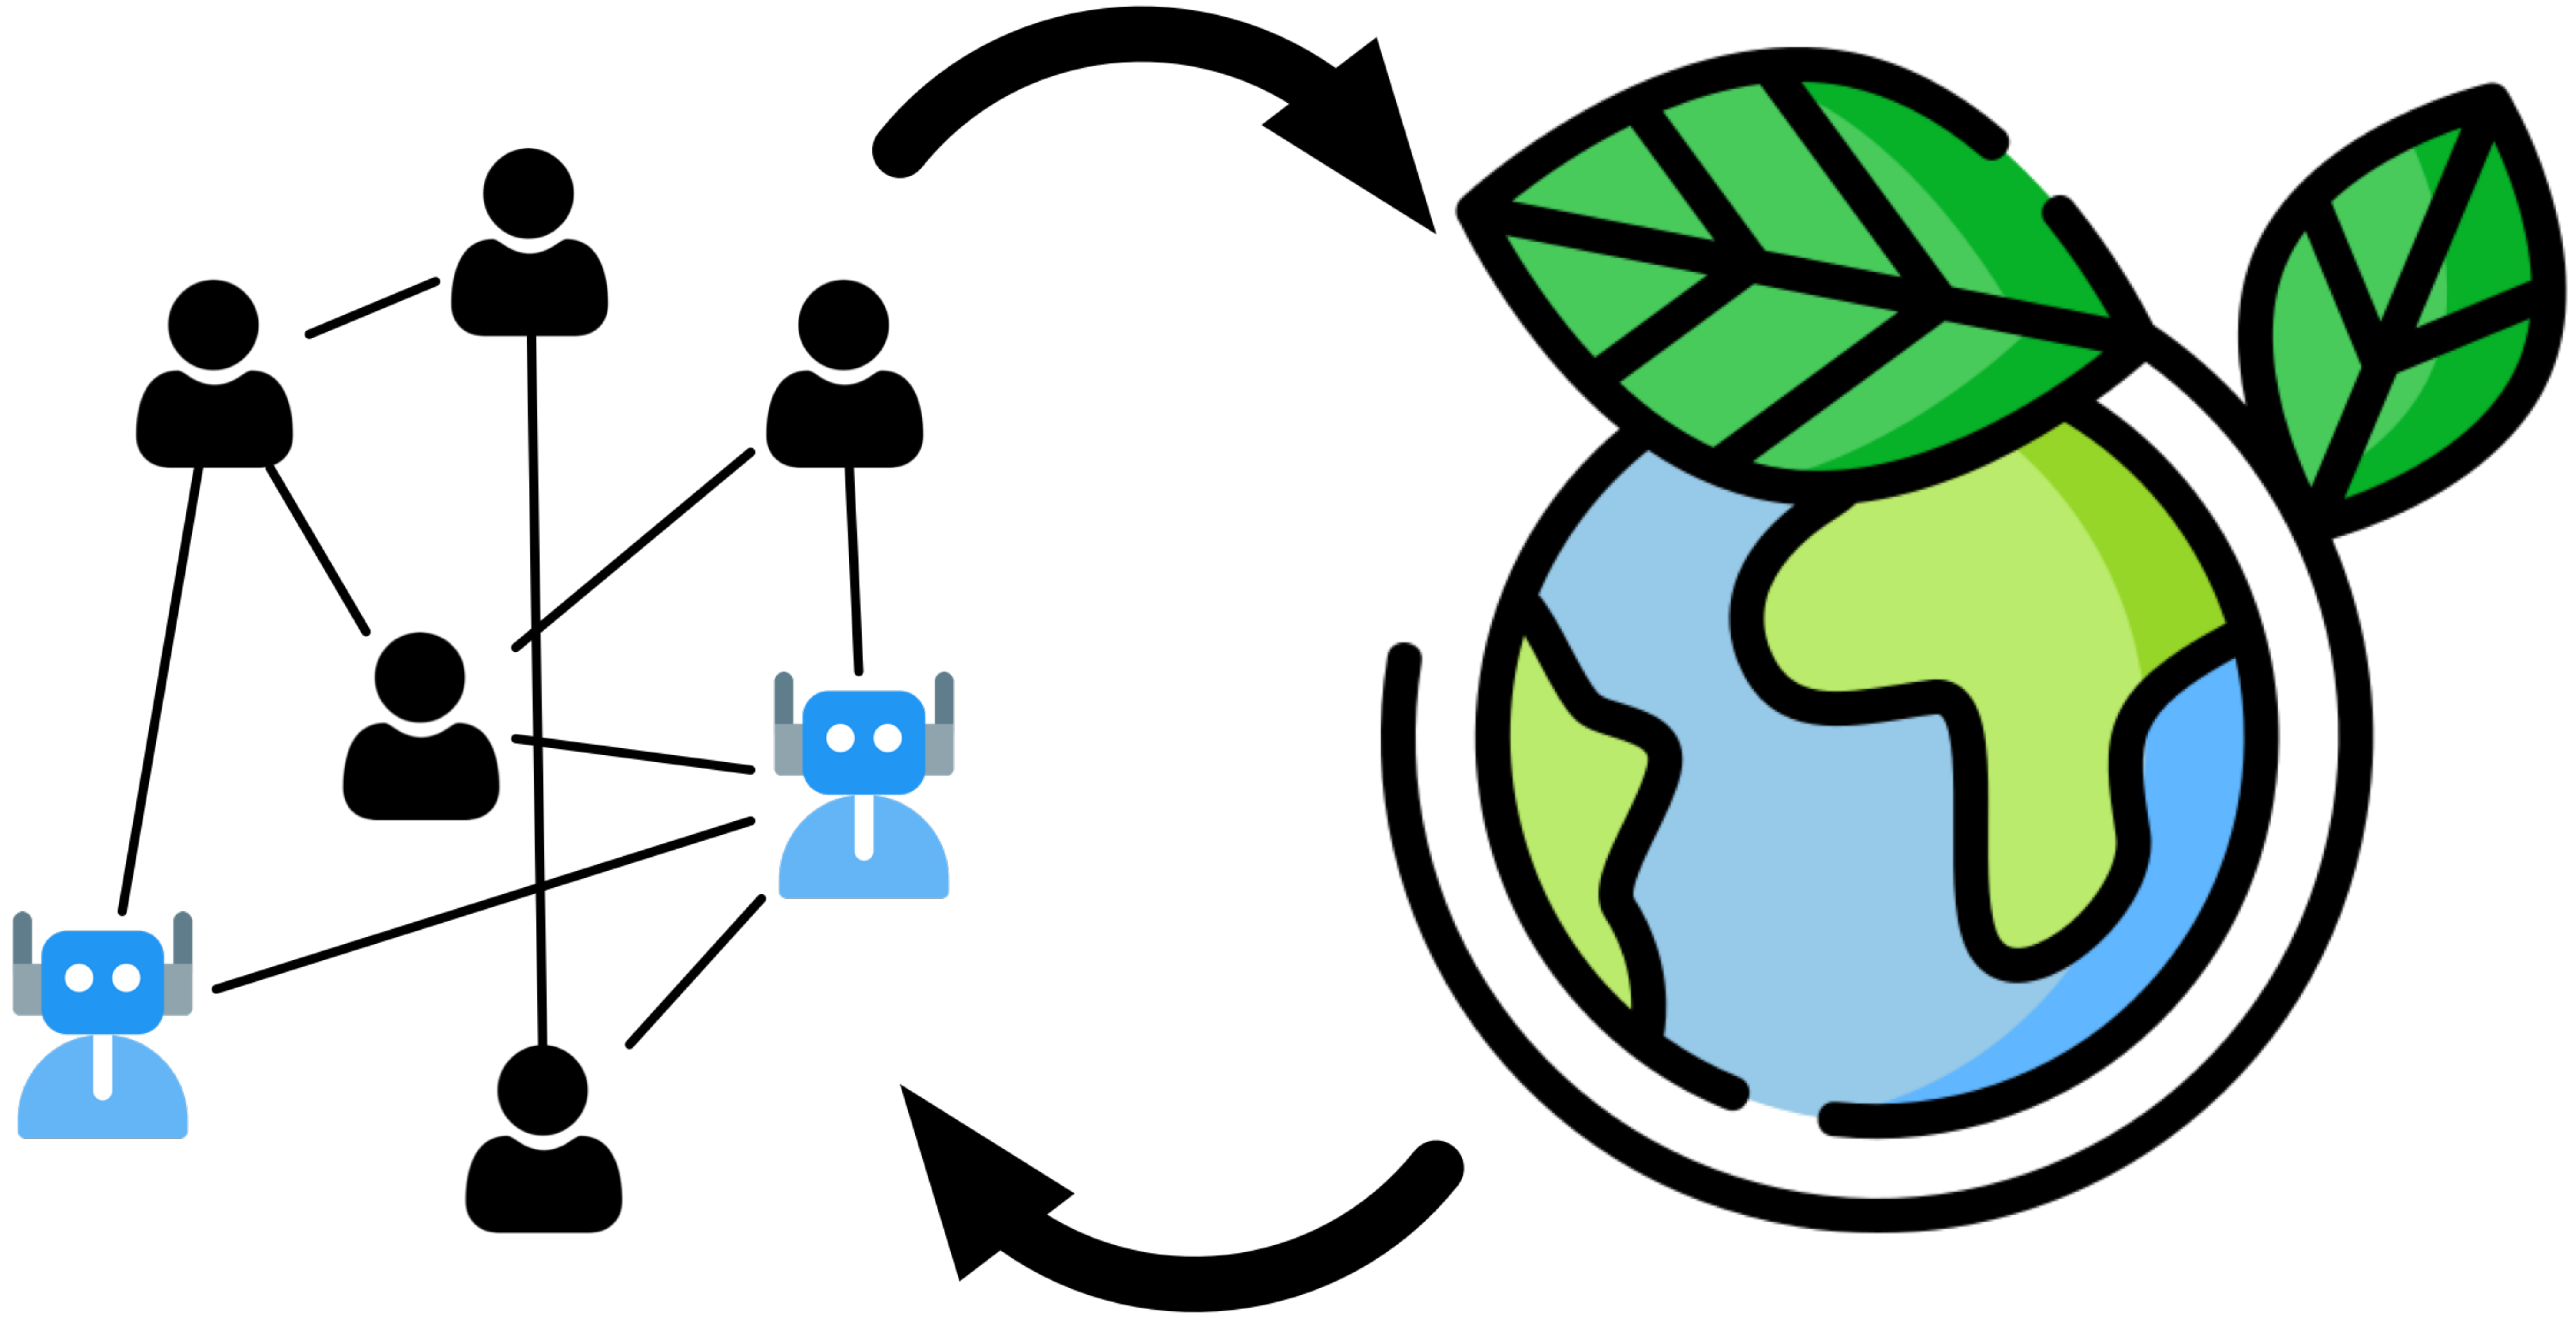

"Management by algorithm is becoming common place, and most successful corporations will delegate critical business decisions to algorithms"

Computers in Human Behavior 134 (2022) 107308

Contents lists available at ScienceDirect

Computers in Human Behavior

journal homepage: www.elsevier.com/locate/comphumbeh

Rise of the machines: Delegating decisions to autonomous AI☆

Cindy Candrian *, Anne Scherer *

URPP Social Networks, Faculty of Business, Economics and Informatics, University of Zurich, Switzerland

ARTICLE INFO

Keywords:
Decision delegation
Artificial intelligence
Social risk
Control premium

ABSTRACT

Delegation is an important part of organizational success and can be used to overcome personal shortcomings and draw upon the expertise and abilities of others. However, delegation comes with risks and uncertainties, as it entails a transfer of power and loss of control. Indeed, research has documented that people tend to under-delegate to other humans, often leading to poor decisions and ultimately negative economic consequences. Today, however, people are faced with a new delegation choice: Artificial Intelligence (AI). Fueled by Big Data, AI is rapidly becoming more intelligent and frequently outperforming human forecasters and decision-makers. Given this evolution of computational autonomy, researchers need to revisit the hows and whys of decision delegation and clarify not only whether people are willing to cede control to AI agents but also whether AI can reduce the under-delegation that is especially pronounced when people are faced with decisions that spur a high desire for control. By linking research on decision delegation, social risk, and control premium to the emerging field of trust in AI, we propose and find that people prefer to delegate decisions to AI as compared to human agents, especially when decisions entail losses (Studies 1–3). Results further illuminate the underlying psychological process involved (Study 1 and 2) and show that process transparency increases delegation to humans but not to AI (Study 3). These findings have important implications for research on trust in AI and the applicability of autonomous AI systems for managers and decision makers.

1. Introduction

Artificial Intelligence (AI) is reshaping the world and presents great opportunities for individuals and businesses. To date, most AI systems that are in wide use directly or indirectly operate under the responsibility of humans, who are in control of the analytical process or outcome. However, a growing number of AI systems go beyond acting as human proxies and operate in a truly autonomous manner. These systems are designed and empowered to make their own decisions fueled by the vast amounts of data they receive, analyze, and interpret (ServiceNow, 2020). Such autonomous AI systems have the capability to surpass human intelligence across various industries and business functions, making them a powerful force for competitive advantage (Schrage, 2017). This technological progress creates entirely new opportunities for humans to delegate decisions to algorithms and artificial agents that no longer require human supervision or direction (Goldbach et al., 2019).

In business practice, Management by Algorithm (MBA) is becoming more commonplace, and many predict that the most successful corporations will be those who delegate critical business decisions to smart algorithms (Schrage, 2017). Autonomous AI can determine entire marketing and capex strategies, identify competitors and target segments, personalize products and prices to customers, and customize communications to individualized preferences (Huang & Rust, 2021). For example, Renaissance Technologies, along with other investment funds, are relying on autonomous algorithms to analyze a situation, author a strategy, and execute it (Schrage, 2017). On an individual level, AI can automate bidding in online auctions (Adomavicius et al., 2009), trading in financial markets (Hendershott et al., 2011), and purchase decisions for customers, as well as automate and augment sales processes and frontline employee tasks (Grewal et al., 2020). Reliance on such new technologies can affect users' judgments and decisions, influence the magnitude of behavioral biases (Dowling et al., 2020; Herrmann et al., 2015), and thus substantially change and even improve decision making, business strategies (Davenport et al., 2020), and market outcomes (Herrmann et al., 2015). Given the vast applicability and the huge potential, some claim that organizations need to clarify when talented humans must defer to algorithmic judgment and delegate

☆ This work was supported by the URPP Social Networks.

Whose interests does the AI agent represent?

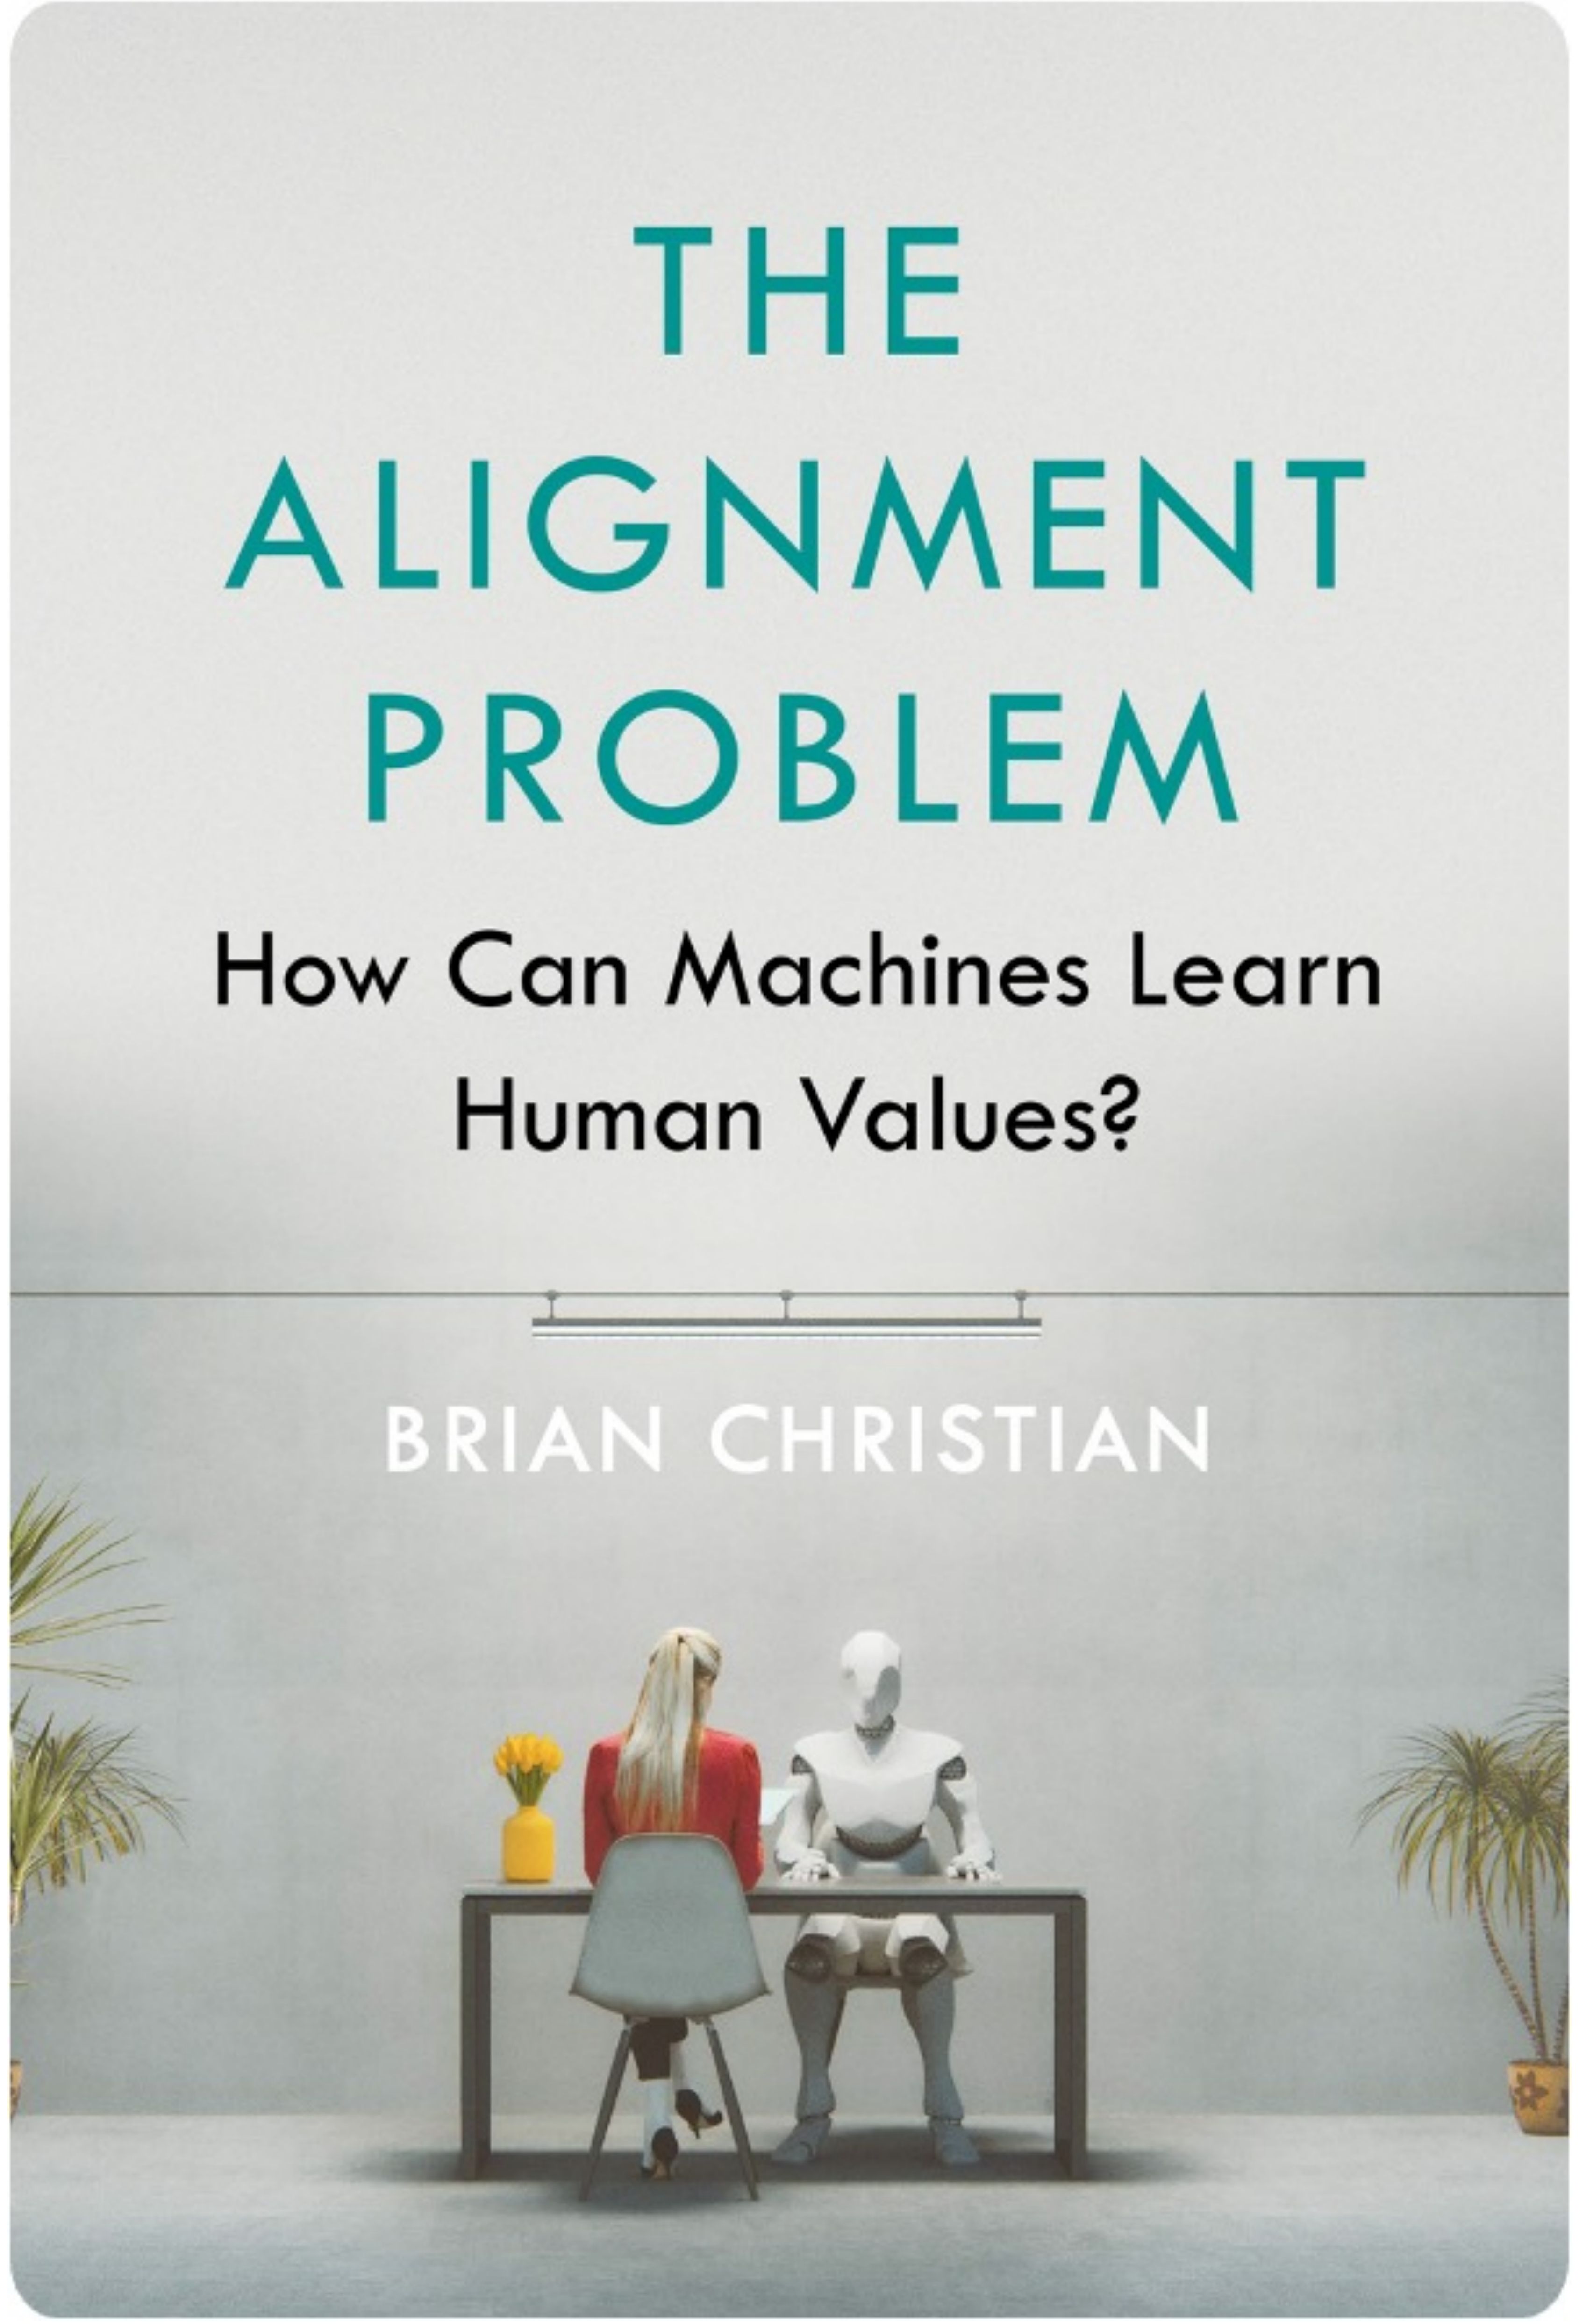

Can humans and AI cooperate?

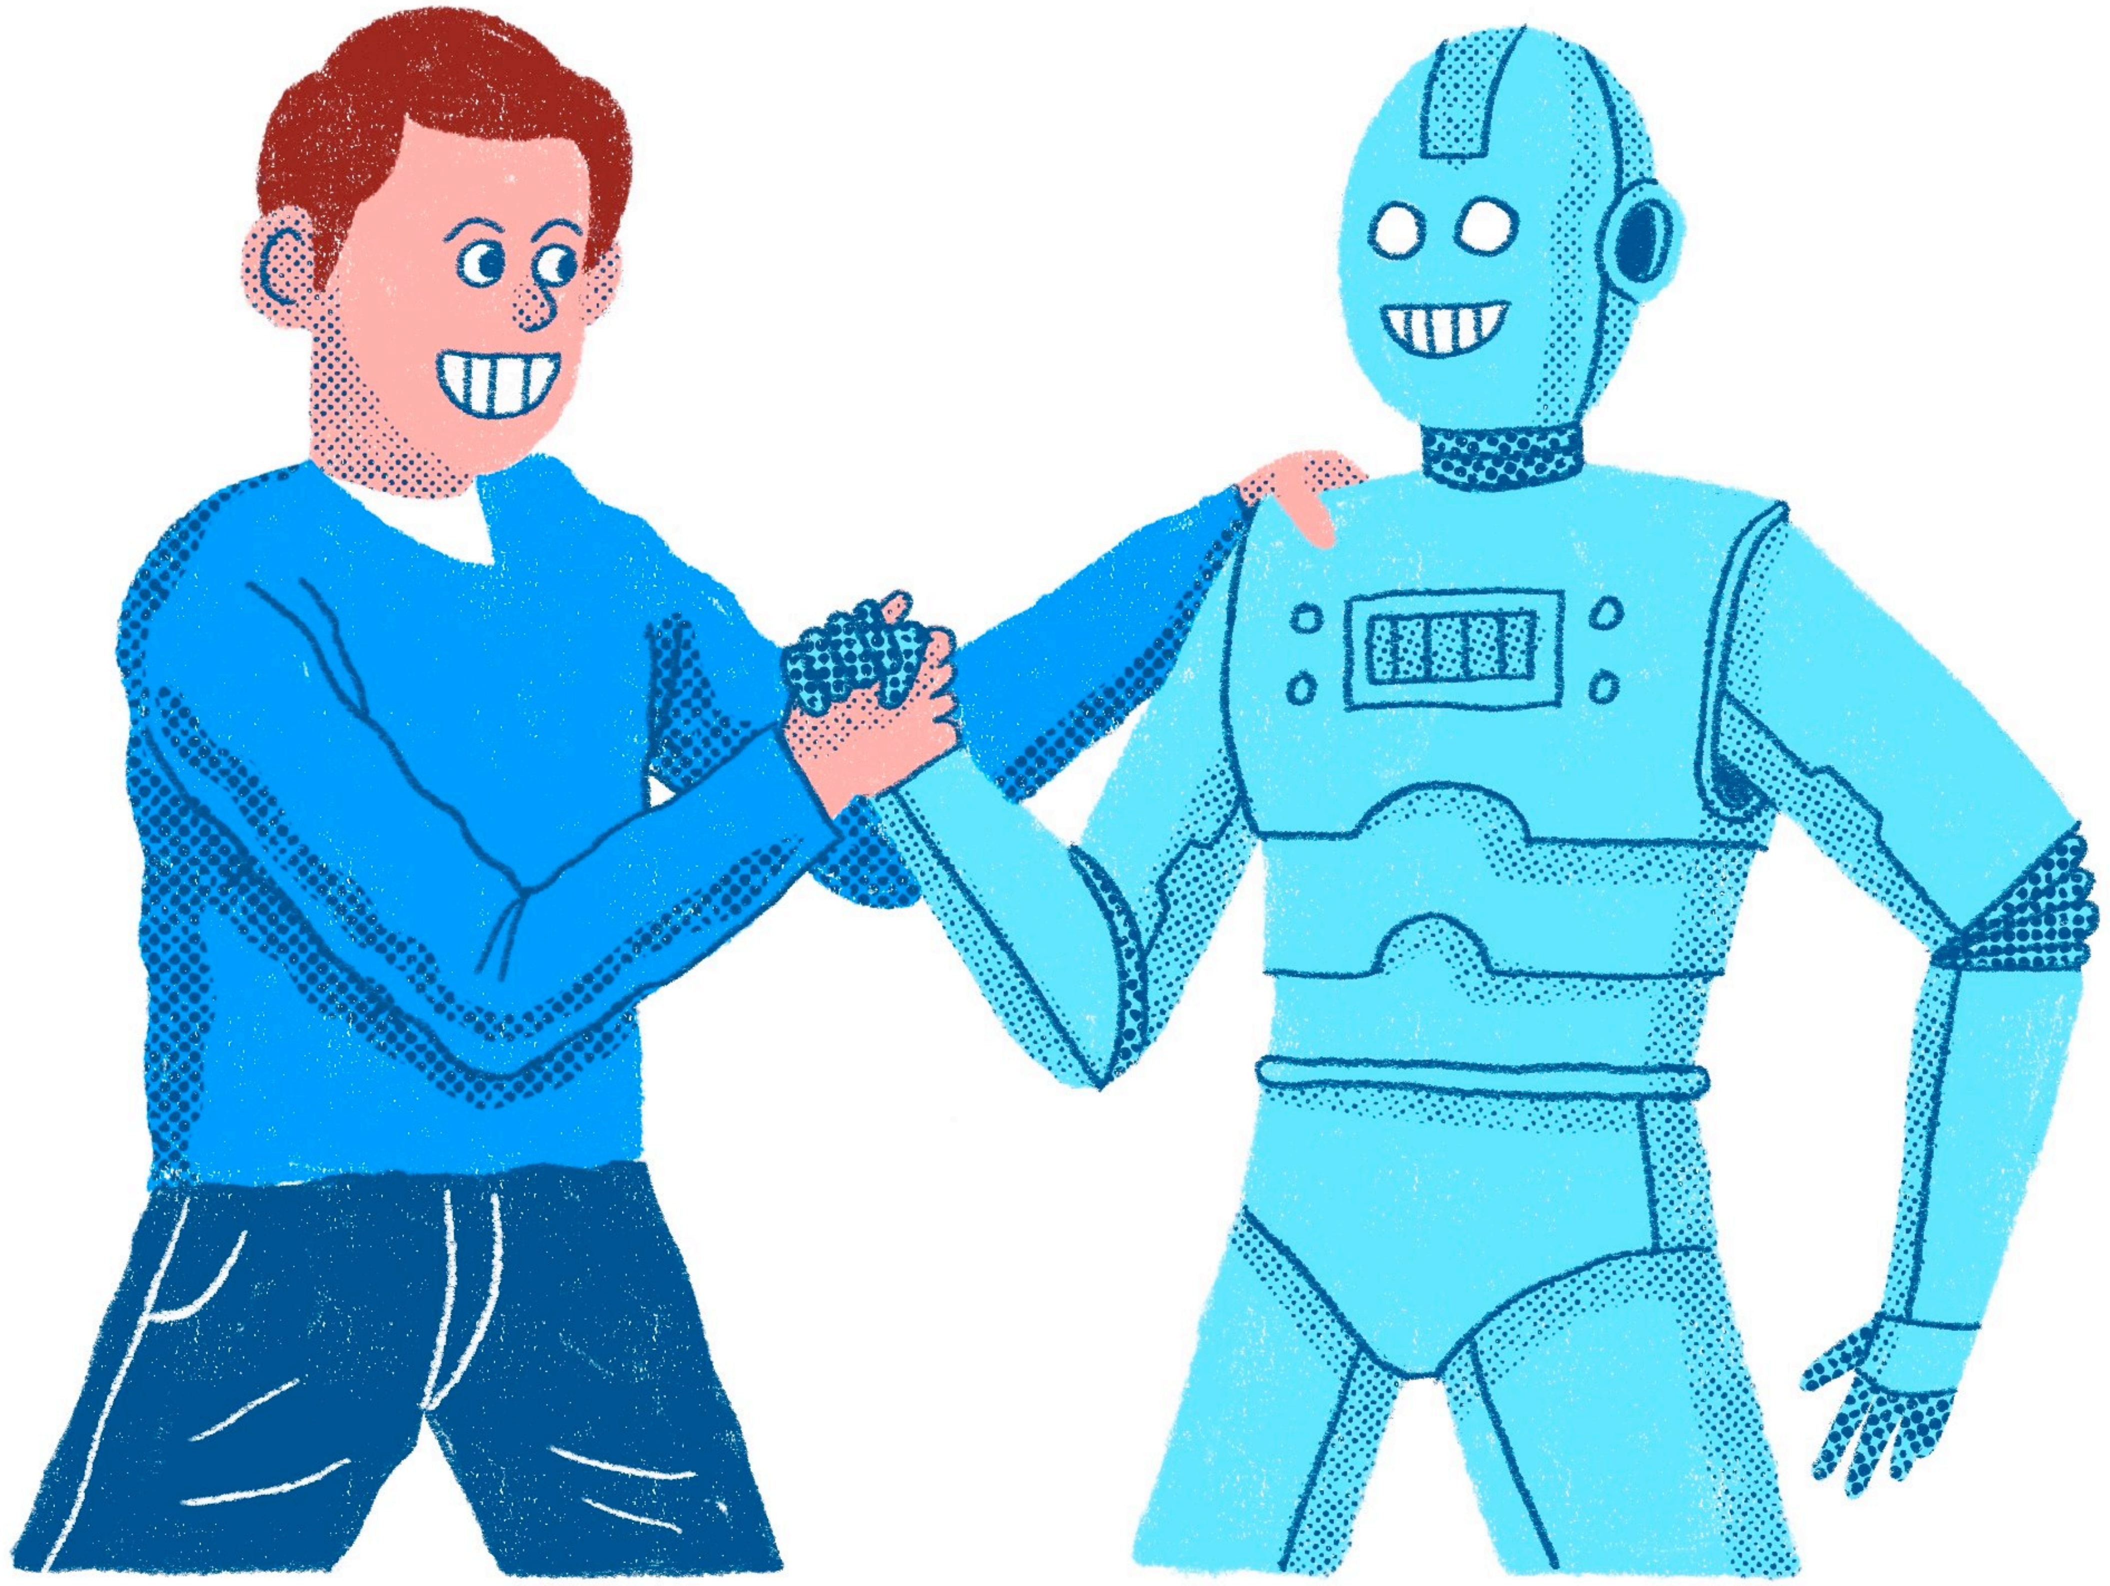

The objective of **Cooperative AI** is to create AI agents that can cooperate with each other and with humans.

Why do we need cooperative AI?

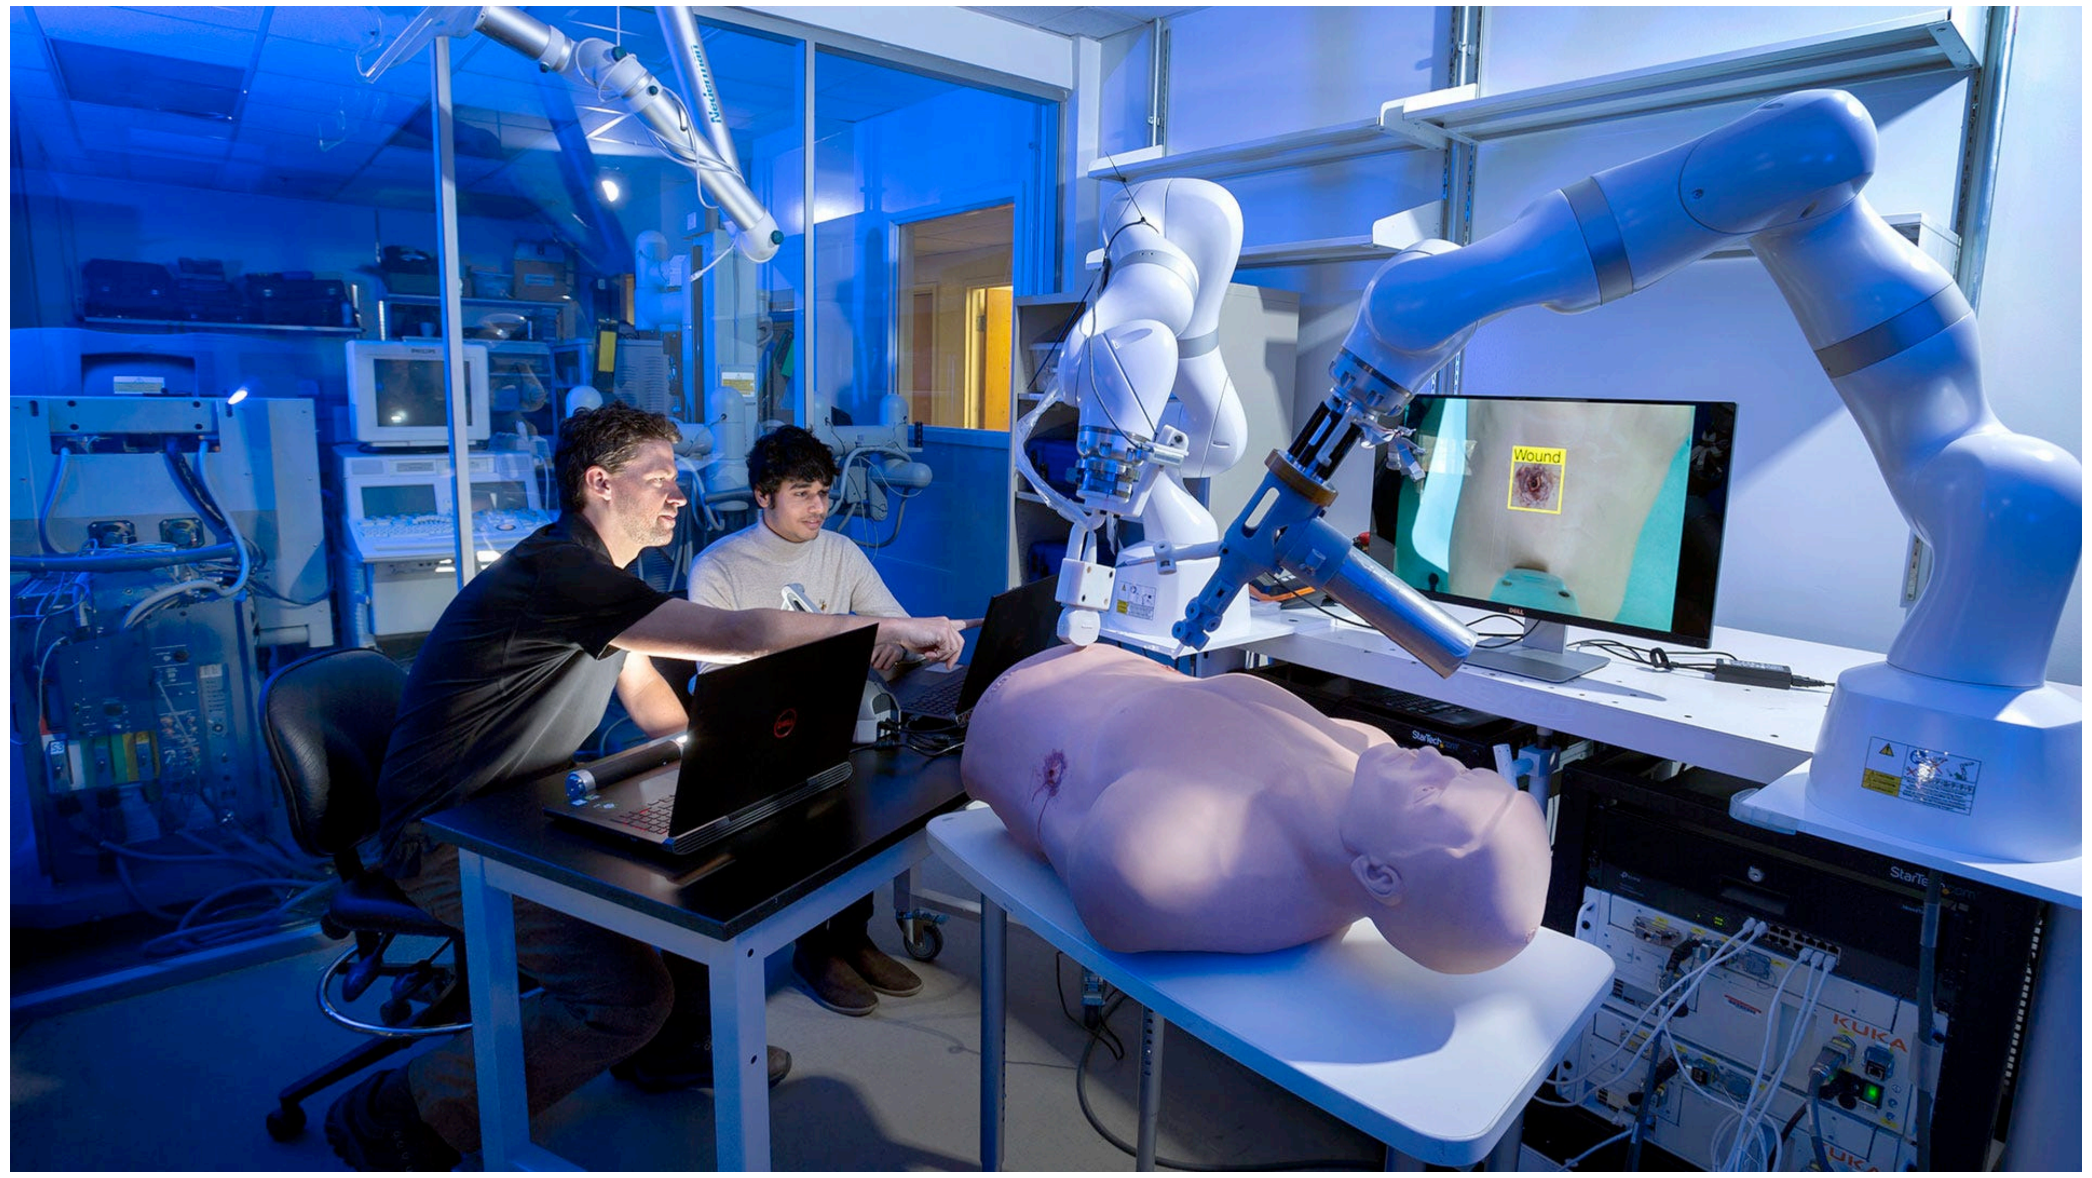

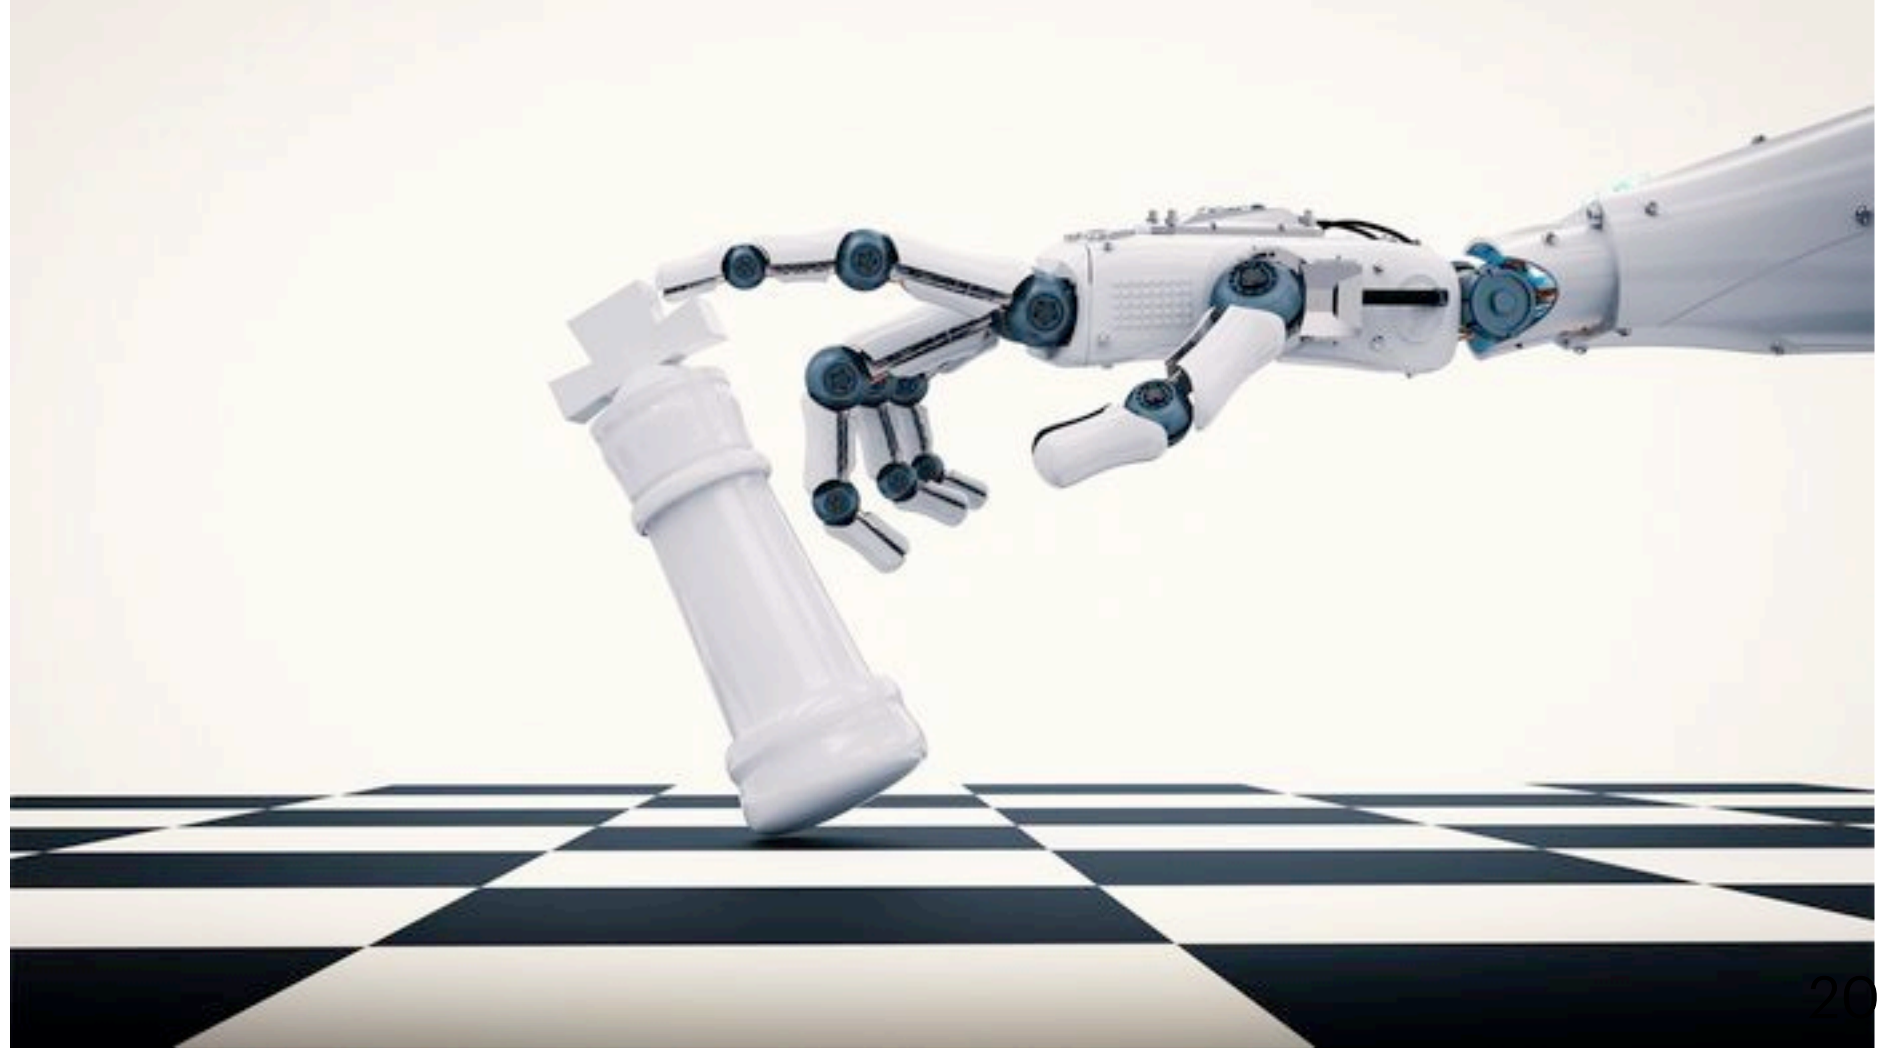

Why do we need cooperative AI?

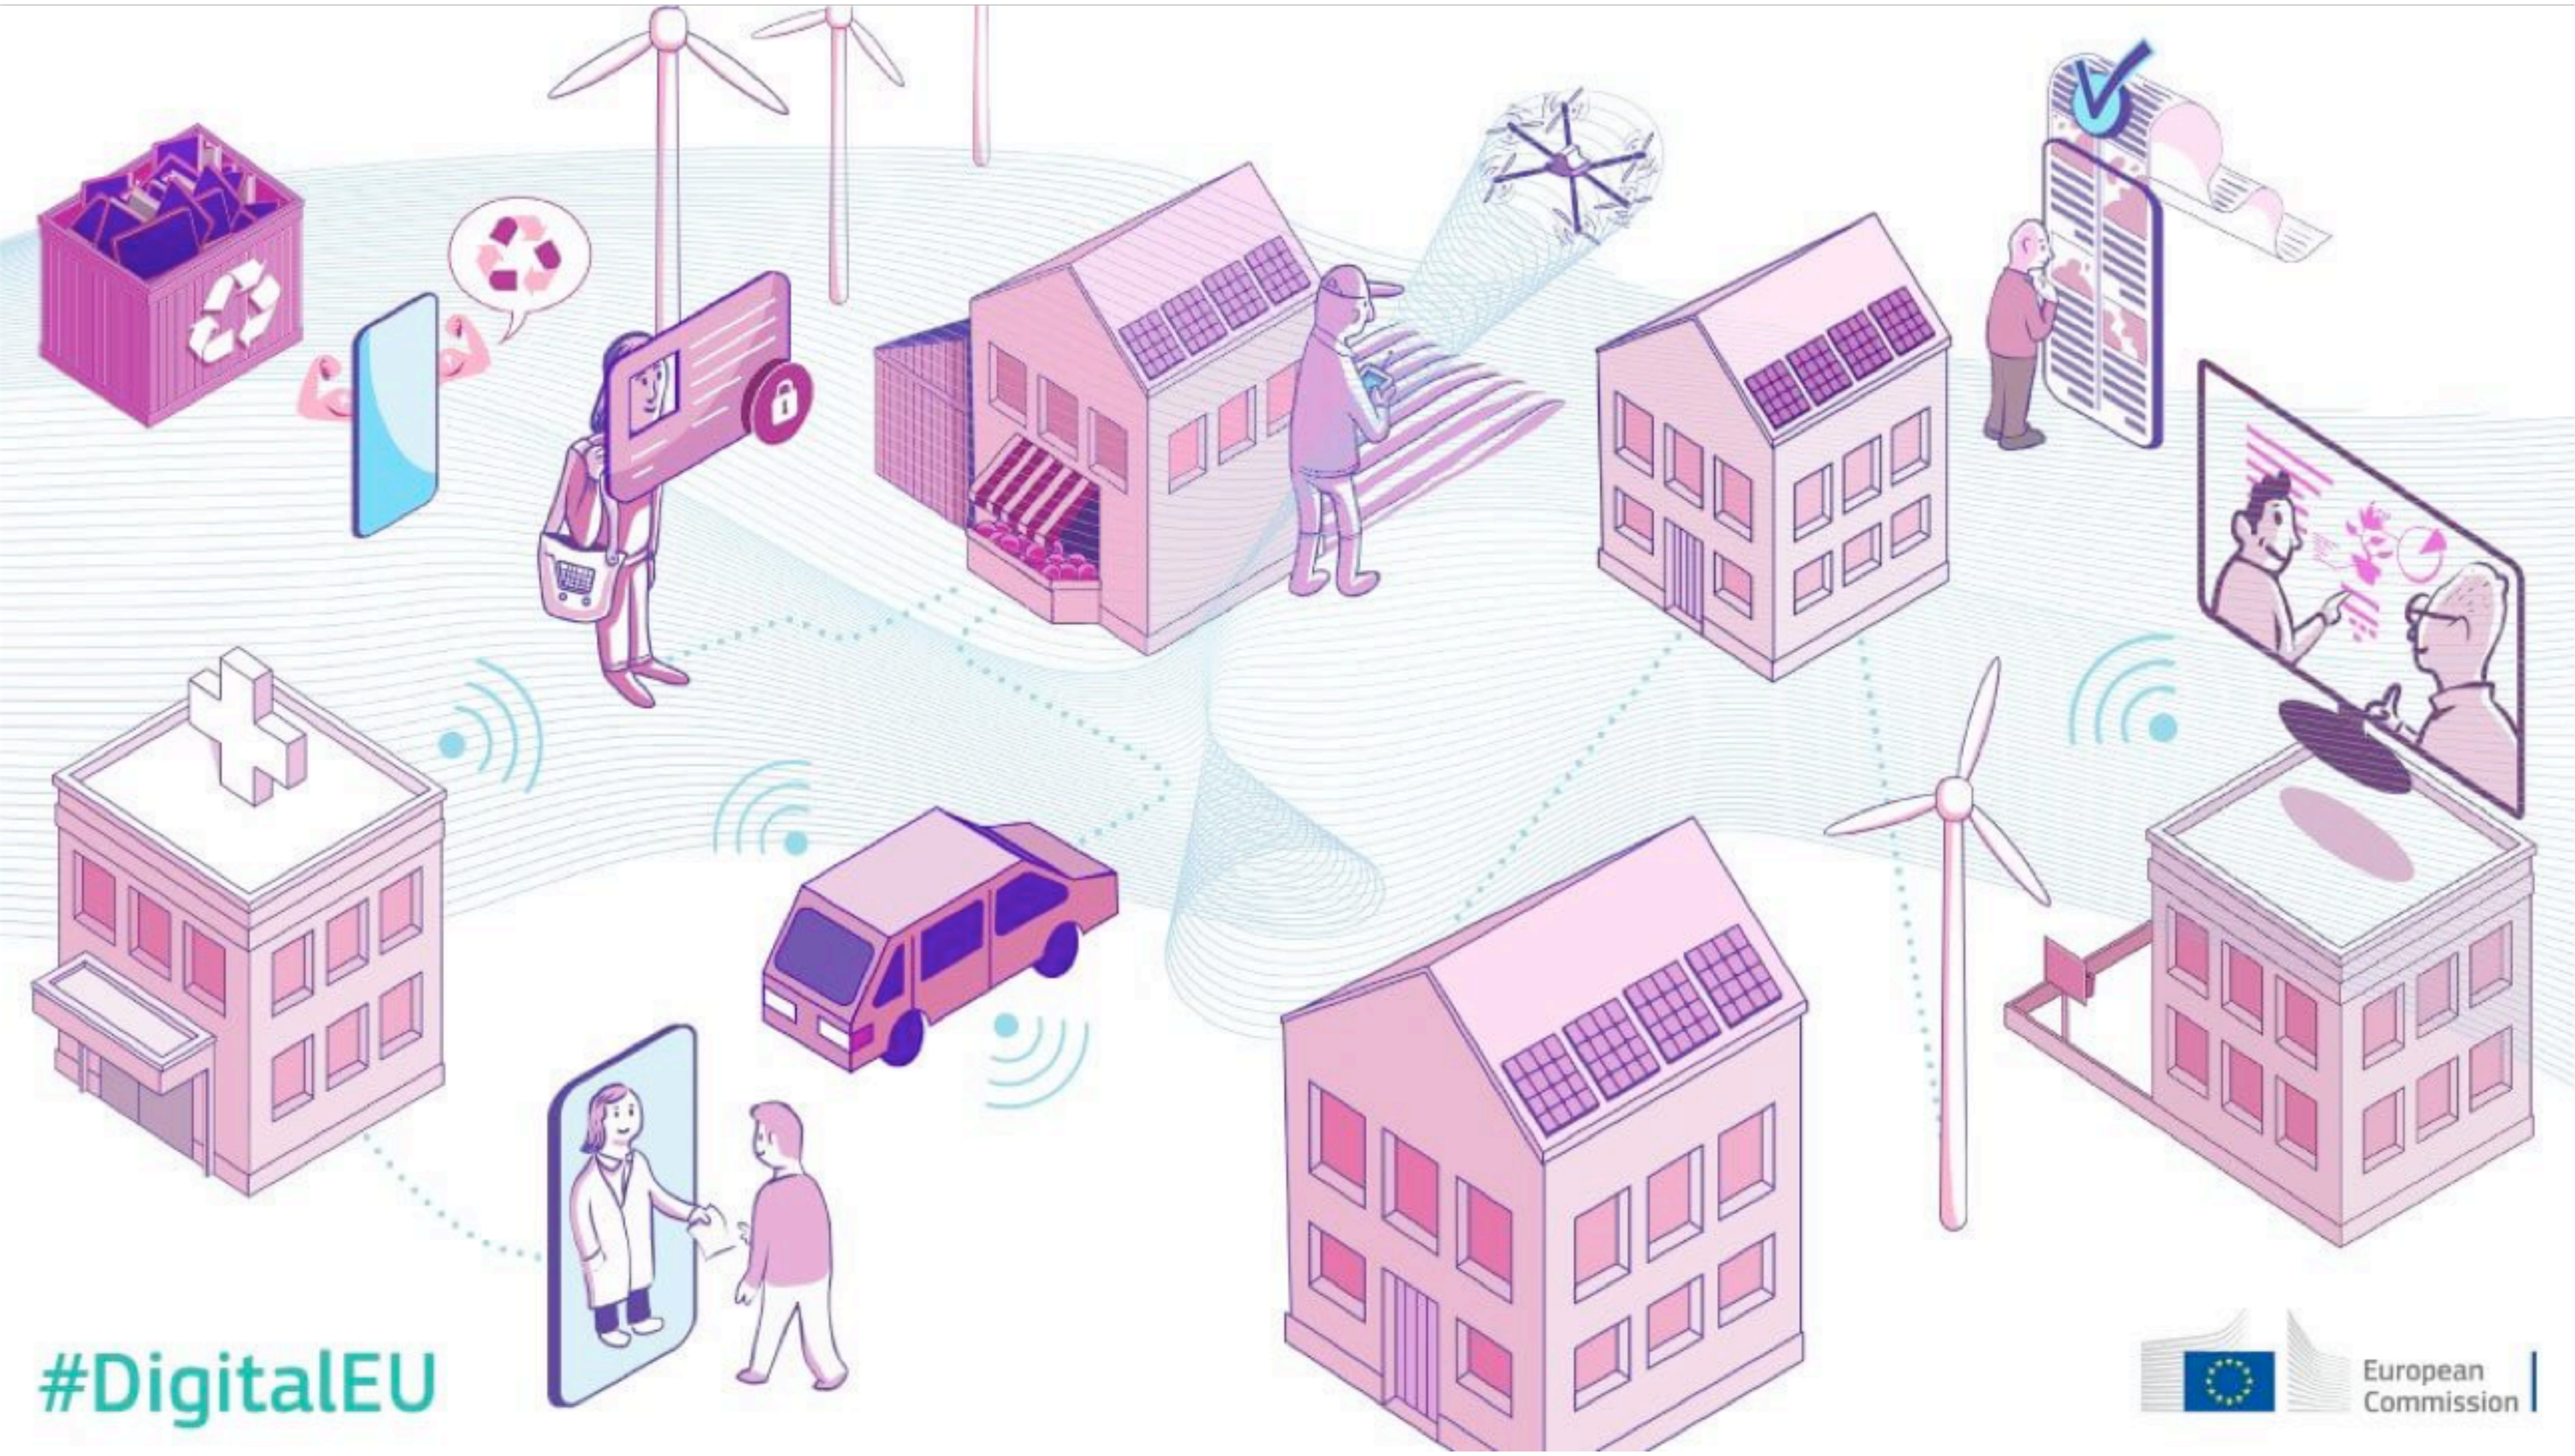

Why do we need cooperative AI?

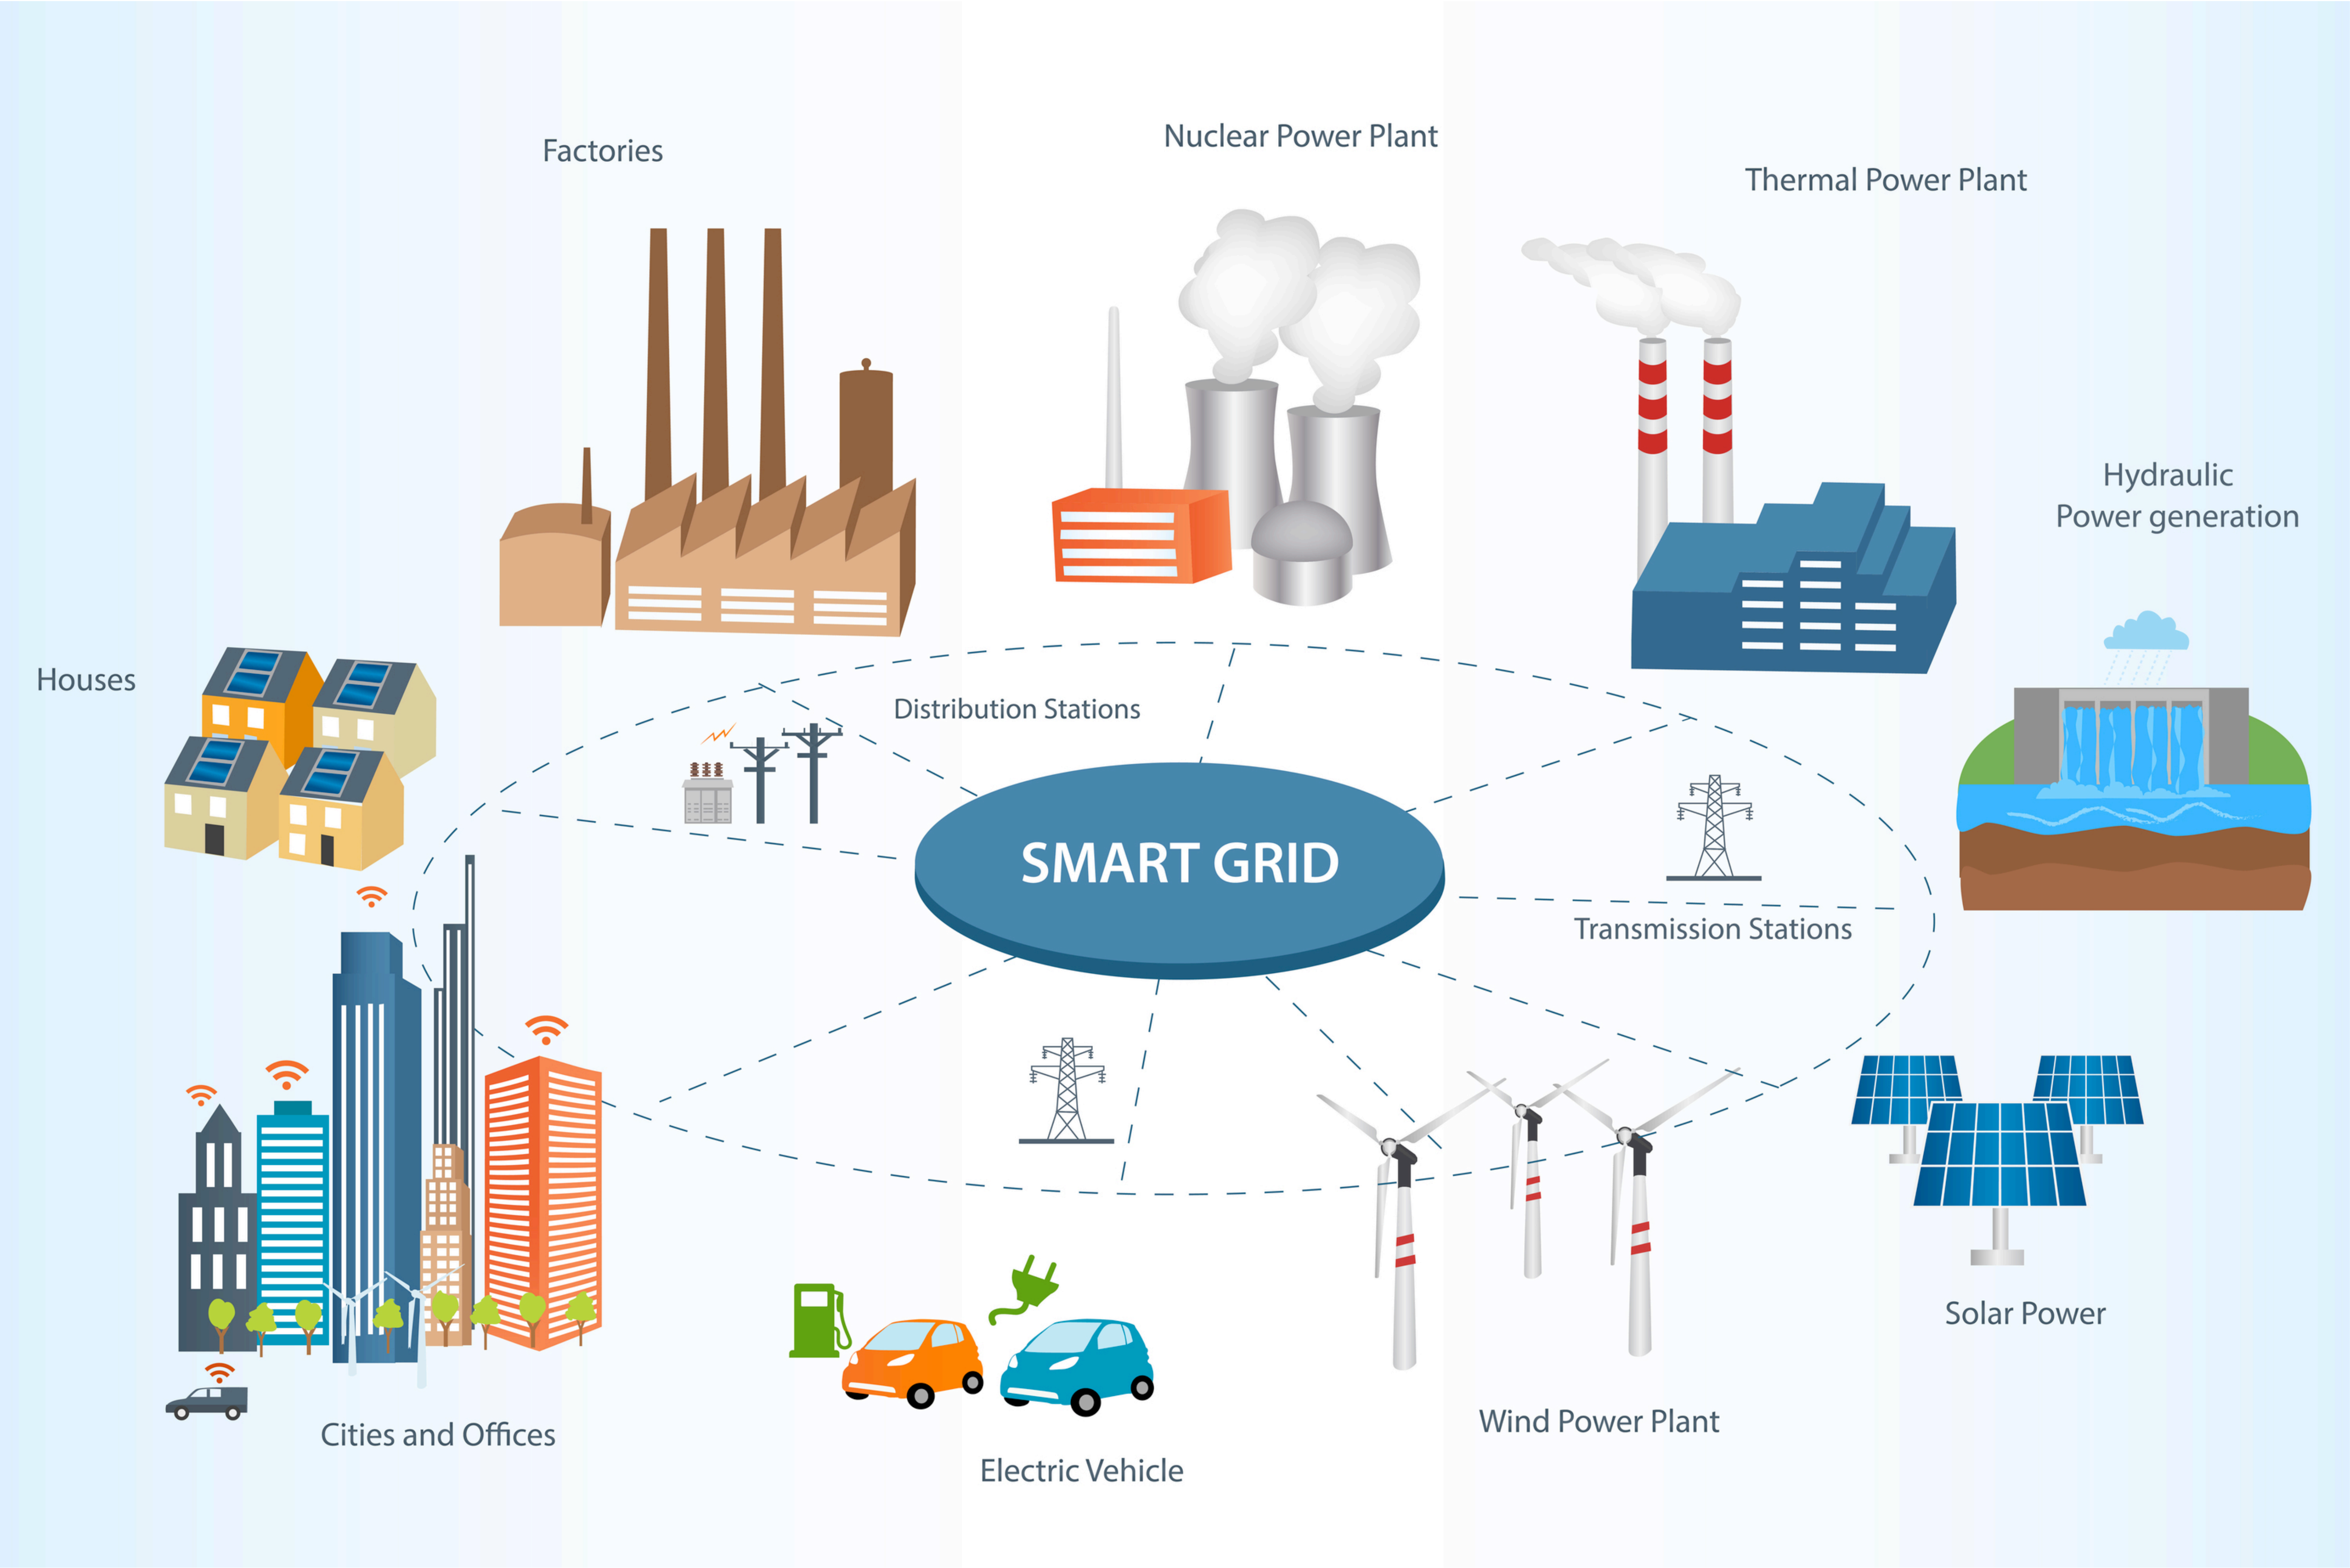

Why do we need cooperative AI?

Why do we need cooperative AI?

COMMENT | 04 May 2021

Cooperative AI: machines must learn to find common ground

To help humanity solve fundamental problems of cooperation, scientists need to reconceive artificial intelligence as deeply social.

Allan Dafoe, Yoram Bachrach, Gillian Hadfield, Eric Horvitz, Kate Larson & Thore Graepel

Artificial-intelligence assistants and recommendation algorithms interact with billions of people every day, influencing lives in myriad ways, yet they still have little understanding of

Human societies are complex (adaptive) systems

Complex systems are systems composed of many elements with various (non-linear) dependencies

Human societies are complex adaptive systems

We need a complex systems approach to Cooperative AI

- The complex system community has vast experience as approaching complex social problems.
- Cooperative AI IS a social problem
- Social Dynamics of AI : psychological and economical cues
- Collective intelligence -> effect on norm evolution
- Behavioral attacks in hybrid populations

Workshop on Evolutionary Dynamics in social, cooperative and hybrid AI (EDAI)

EDAI 2024

Evolutionary Dynamics in social, cooperative and hybrid AI Workshop at ECAI 2024, Santiago de Compostela, Spain

EDAI 2024: Evolutionary Dynamics in social, cooperative and hybrid AI

19.10 or 20.10, 2024, Santiago de Compostela, Spain

- News
- Description
- Important Dates
- Submission Details
- Accepted Papers
- Program
- Organization

https://edai-workshop.github.io/2024/

Part 2: Introduction to Game Theory

*“Game theory studies (strategic) decision-making where the outcome depends on the decisions of other **agents** involved in the **interaction** “*

Nodes/vertices

Links/edges

Networks

Game

“Game theory studies (strategic) decision-making where the outcome depends on the decisions of other ***agents*** *involved in the* ***interaction*** *“*

Nodes/vertices

Links/edges

Networks

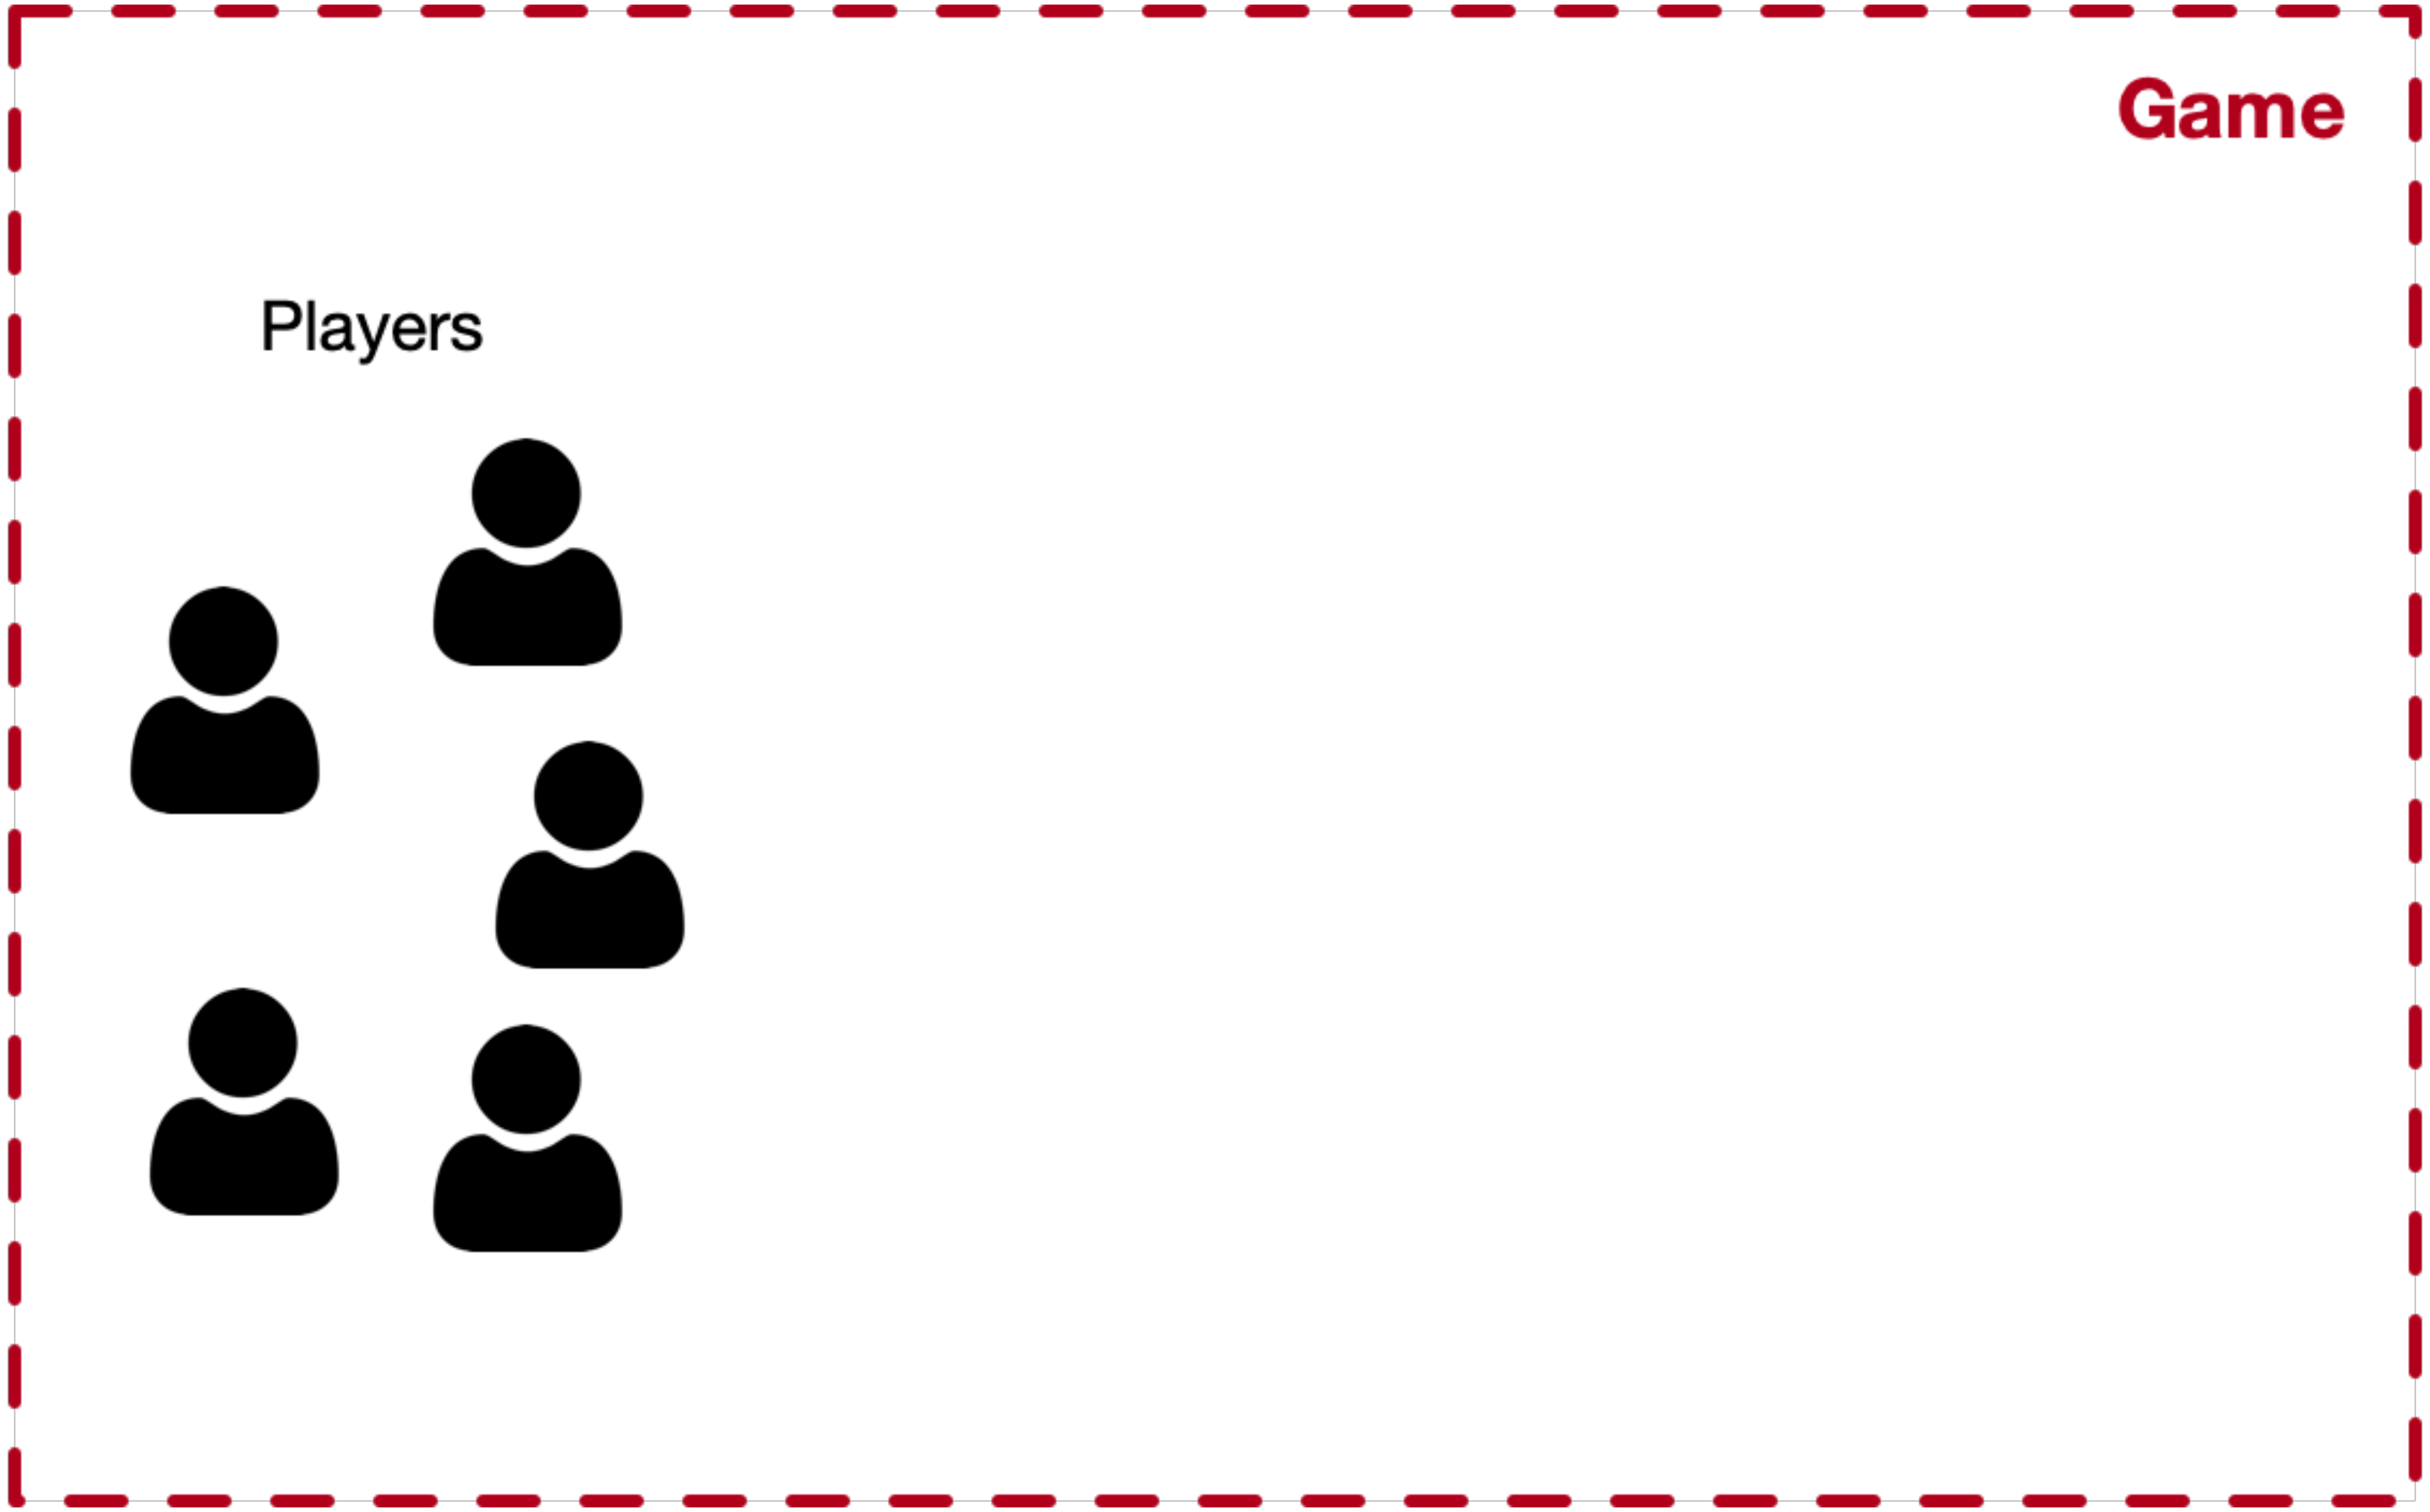

Game
Players

Game

Players

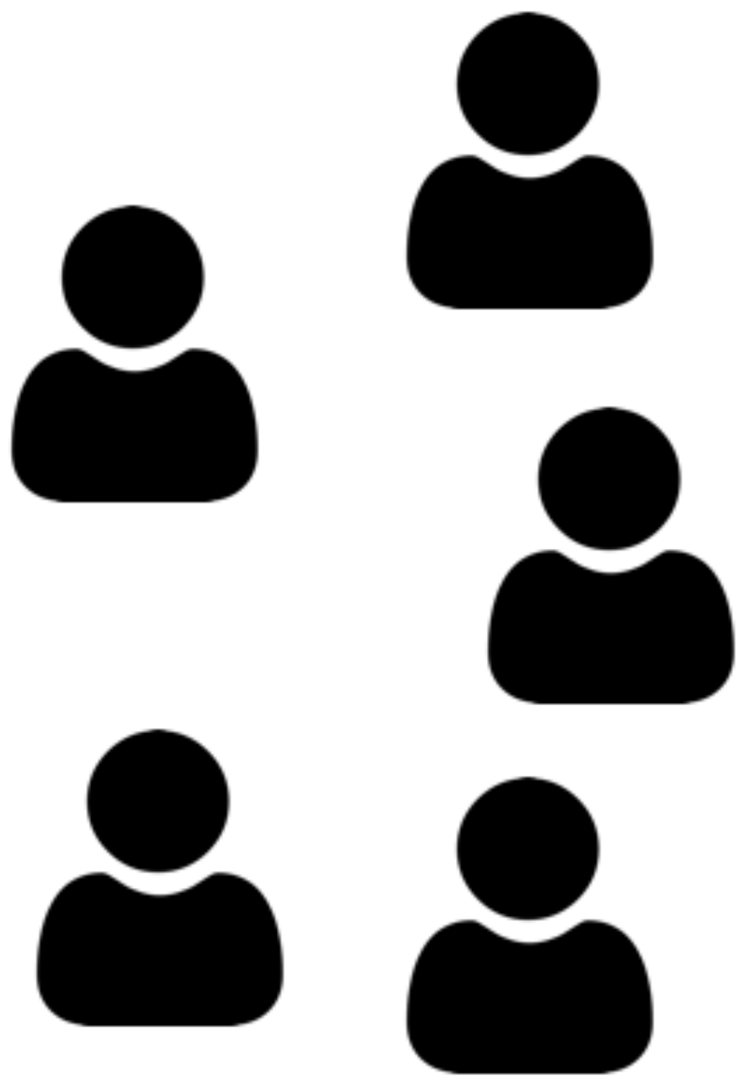

Strategic setting

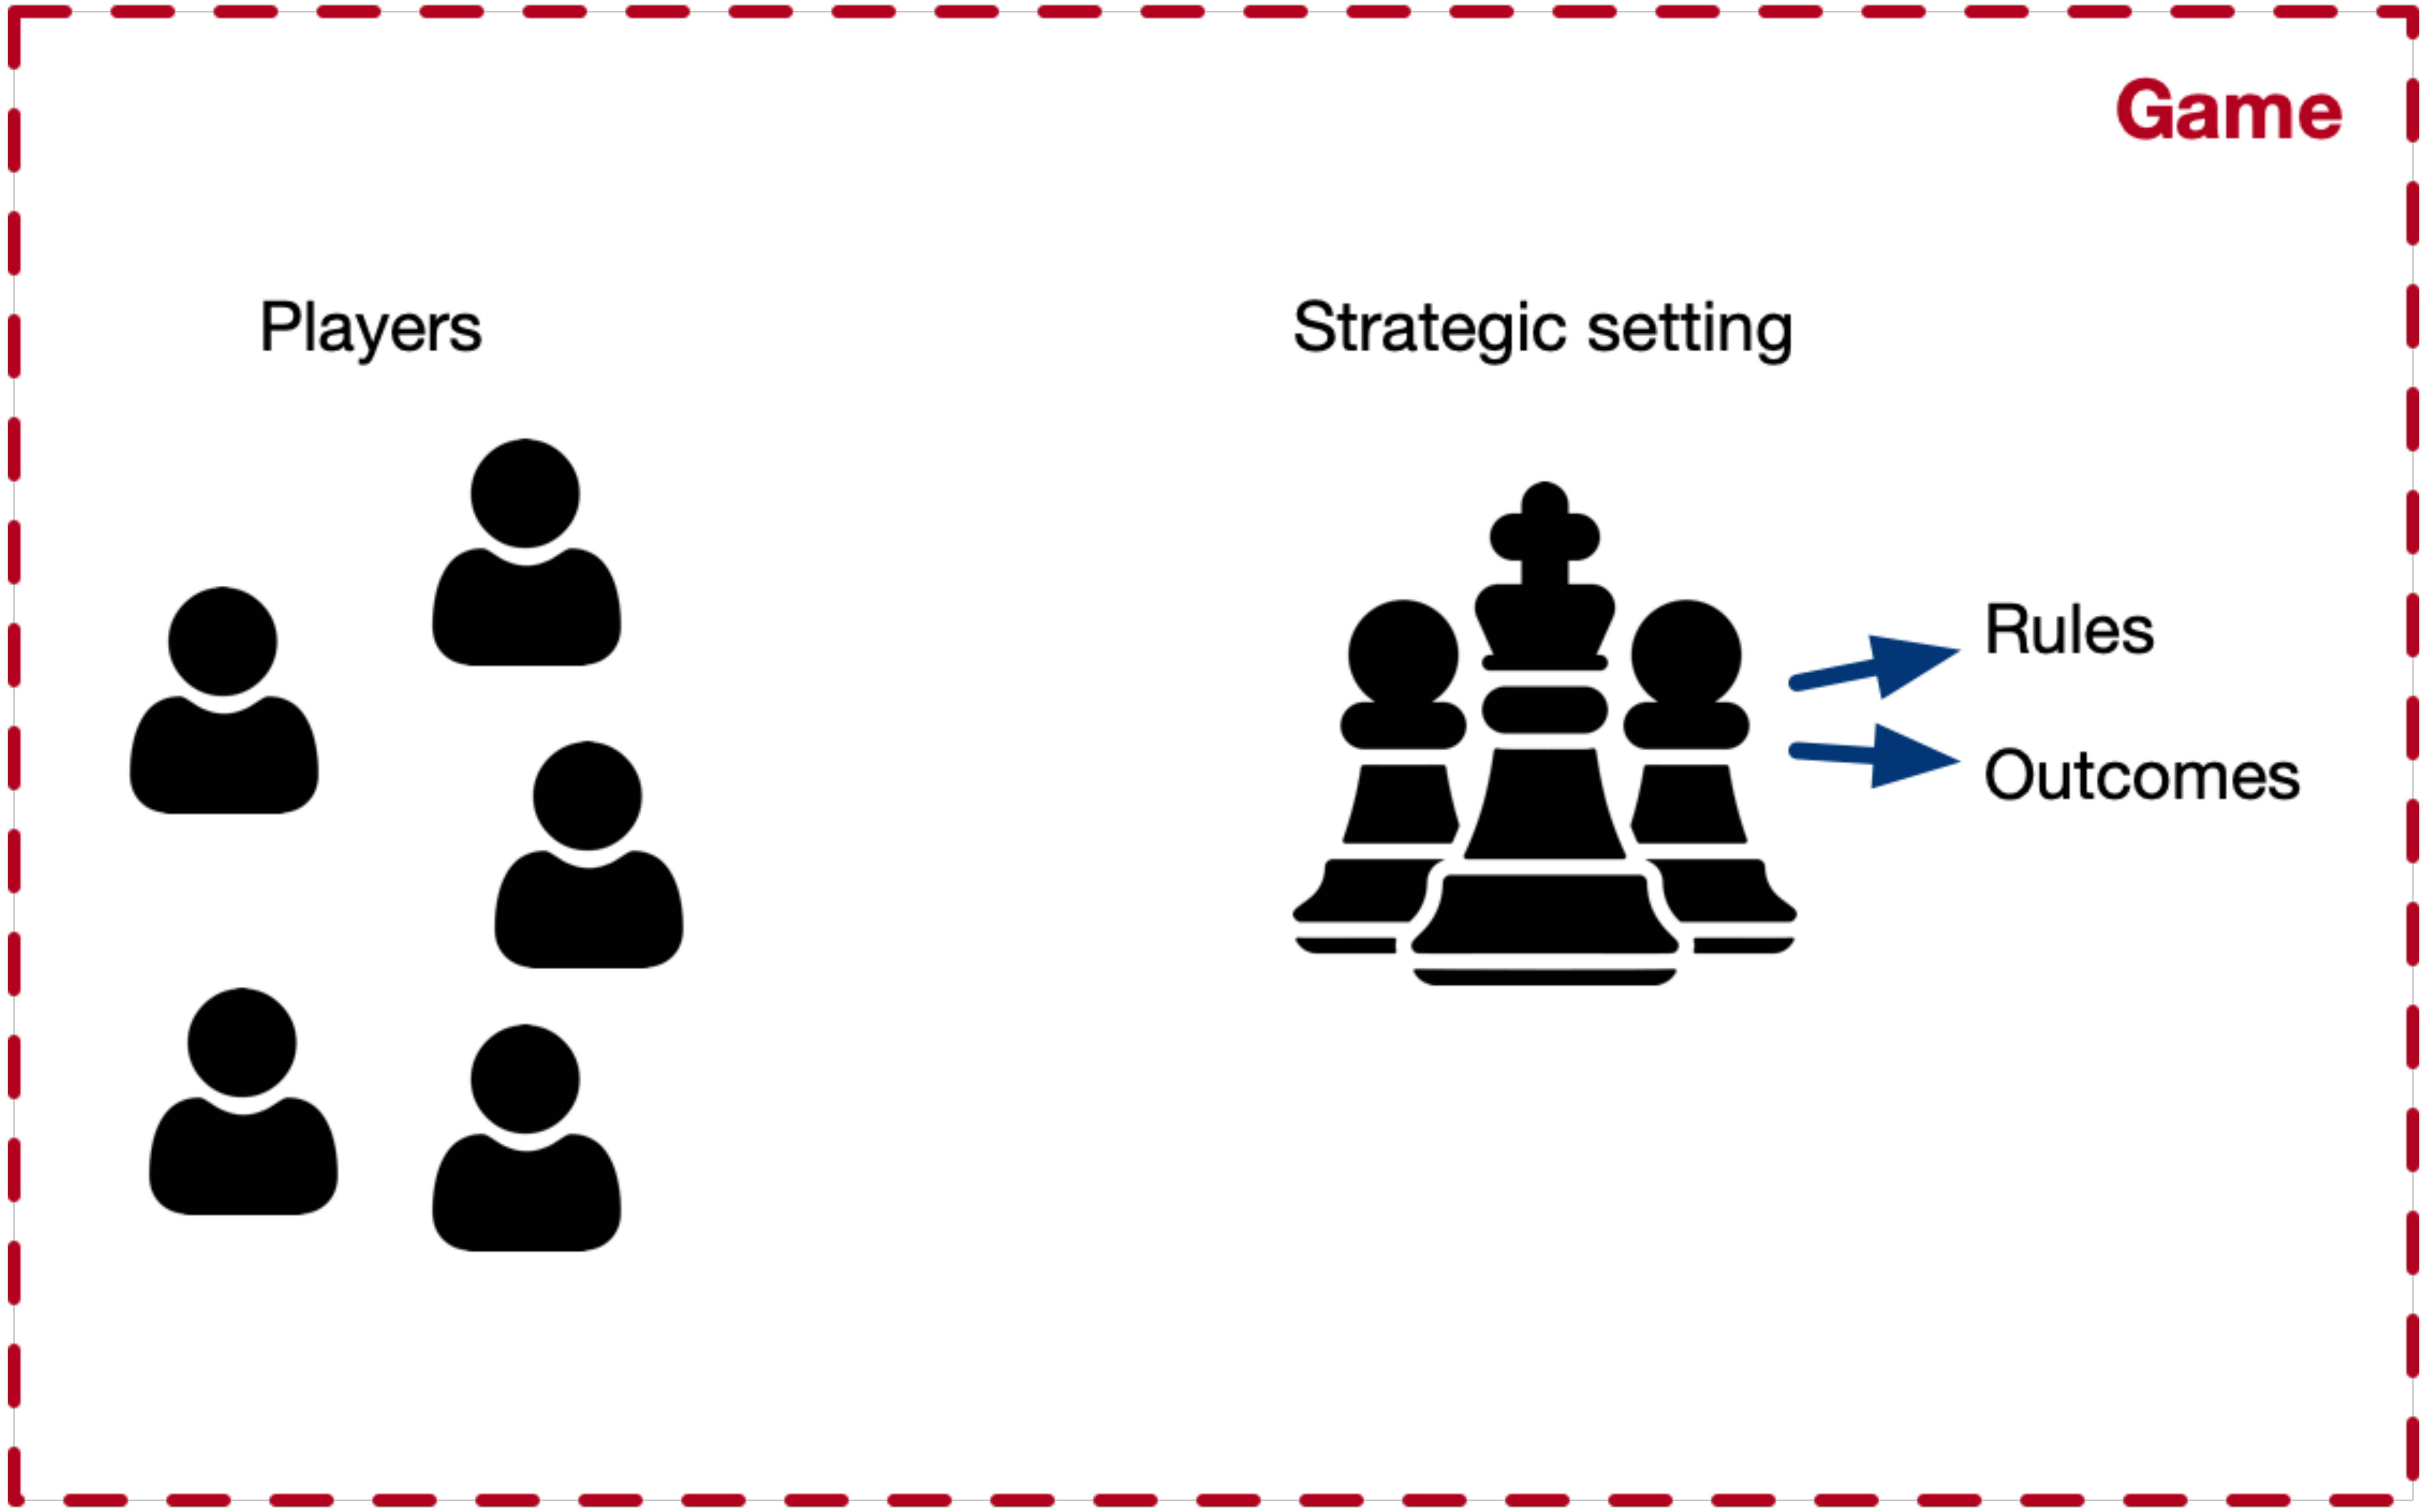
Game
Players
Strategic setting
Rules
Outcomes

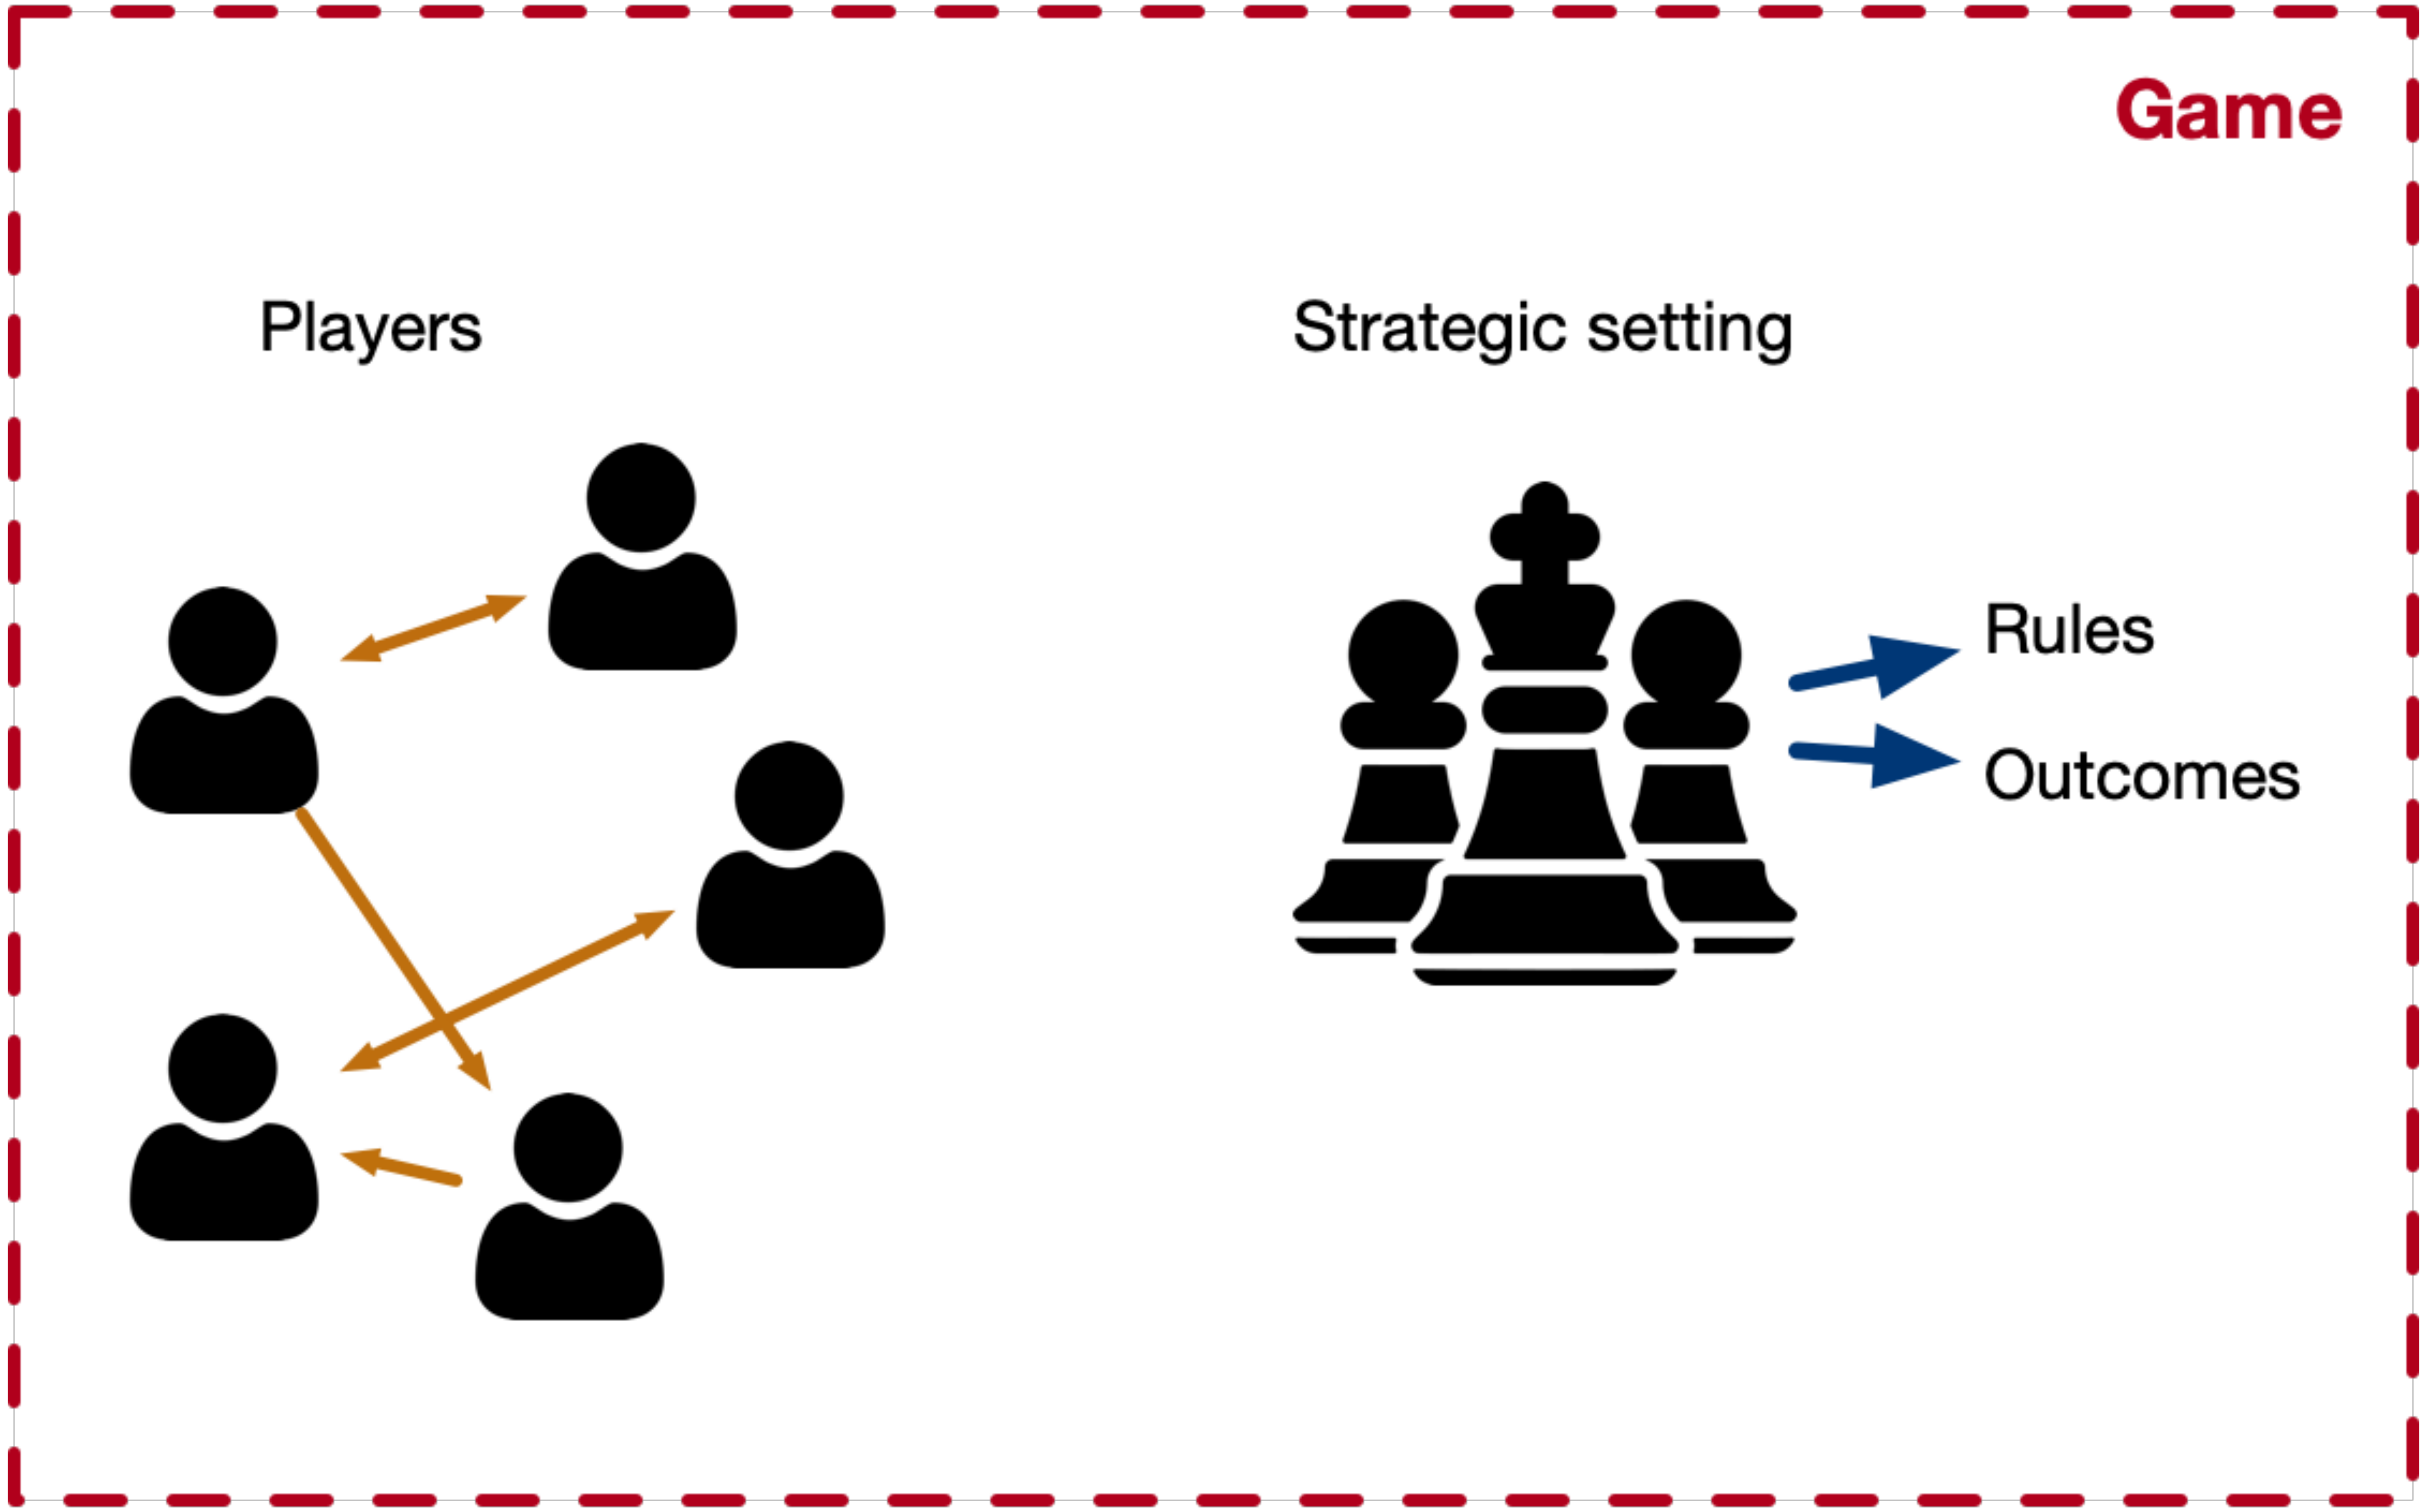
Game
Players
Strategic setting
Rules
Outcomes

A **Game** defines the set of **actions** a player can take, and their **consequences**

A **Game** defines the set of **actions** a player can take, and their **consequences**

A player's **strategy** is the combination of those actions

Some important definitions

Action

The set of actions refers to the available options that a player has at a given moment in a strategic interaction.

Strategy

A strategy represents **how a player chooses among the available actions** in a setting where the outcome depends on the actions of all involved participants. In other words, a strategy consists of an assignment of action for any situation in the game (e.g., an algorithm).

Some important definitions

Pure strategy

If this assignment is **deterministic**, we commonly refer to it as a pure strategy. Pure strategies are a particular case of a wider set of probabilistic assignments between actions and game situations known as Mixed strategies.

Mixed strategy

Probabilistic strategies are known as mixed strategies and can also be represented by a probability of choosing a given pure strategy at each game situation.

Strategy profile

A strategy profile defines **the set of strategies adopted by all players**.

Players have **preferences** over the available choices and consequences!

Rationality and utility

Important: in this course we will, unless indicated, assume that utility is equivalent to expected payoff, and will abuse the notation:

$$E[u(x)] \equiv \Pi(x) \equiv u(x)$$

We will also use the following notation to represent the payoff of player i when making action a_i, given the action of all other players a_{-i}.

$$u_i(a_i, a_{-i}) \equiv \pi_i(a_i, a_{-i})$$

Finally, we will use the notation e_i to represent a strategy of player i to avoid any confusion with the state space S

Introducing game theory

YouTube video starting at 4:12

“Golden Balls is a British daytime game show which was presented by Jasper Carrott. It was broadcast on the ITV network from 18 June 2007 to 18 December 2009. It was filmed at the BBC Television Centre. Golden Balls Ltd licensed their name to Endemol for the game show and merchandise.” [Wikipedia Oct. 2020]

Sarah and Steve playing the golden balls game for 100150 pound

Players 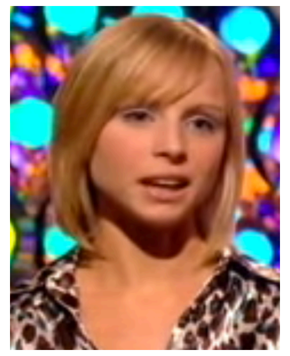 

Sarah Steve

Actions ∈ {split, steal}

Preferences over actions:

Both prefer 100150, over 50075, over 0

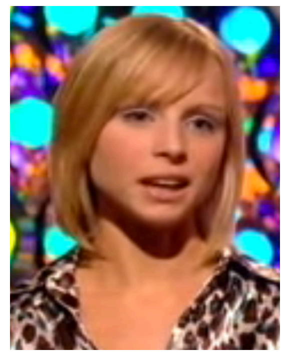

(steal, split) > (split, split) > (split, steal)=(split,split)

(steal, split) > (split, split) > (split,steal)=(split,split)

We call this a **symmetric** game

Normal form of the game

Sarah

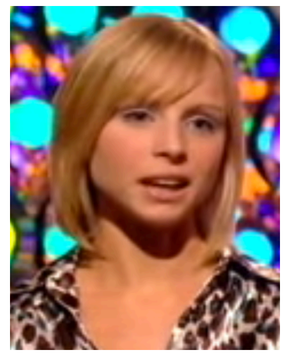

Steve

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

The simultaneous choice of both players is a **strategy profile**, e.g. (Split, Steal)

What should they do ?

Strict **Dominance**

In a strategic game player i's strategy e_i'' **strictly** dominates strategy e_i' if

$u_i(e_i'', e_{-i}) > u_i(e_i', e_{-i})$ for every list e_{-i} of the other player's strategies

Weak **Dominance**

In a strategic game player i's strategy e_i'' **weakly** dominates strategy e_i' if

$u_i(e_i'', e_{-i}) \geq u_i(e_i', e_{-i})$ for *every* list e_{-i} of the other player's actions <u>and</u>

$u_i(e_i'', e_{-i}) > u_i(e_i', e_{-i})$ for *some* list e_{-i} of the other player's actions

Normal form
of the game

Sarah

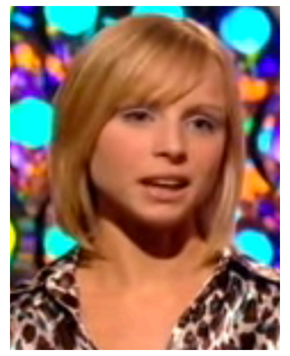

Steve

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Solution concepts

Solution concepts ?

Principles according to which one can identify interesting subsets of outcomes of a game [see book Leyton-Brown and Shoham]

The Nash equilibrium is one of the most famous and important, yet others exist:

We'll provide later some additional solution concepts for games that are expressed in normal form (note there are more) and (if time allows) games expressed in extensive form

Correlated equilibria

Evolutionary Stable Strategy

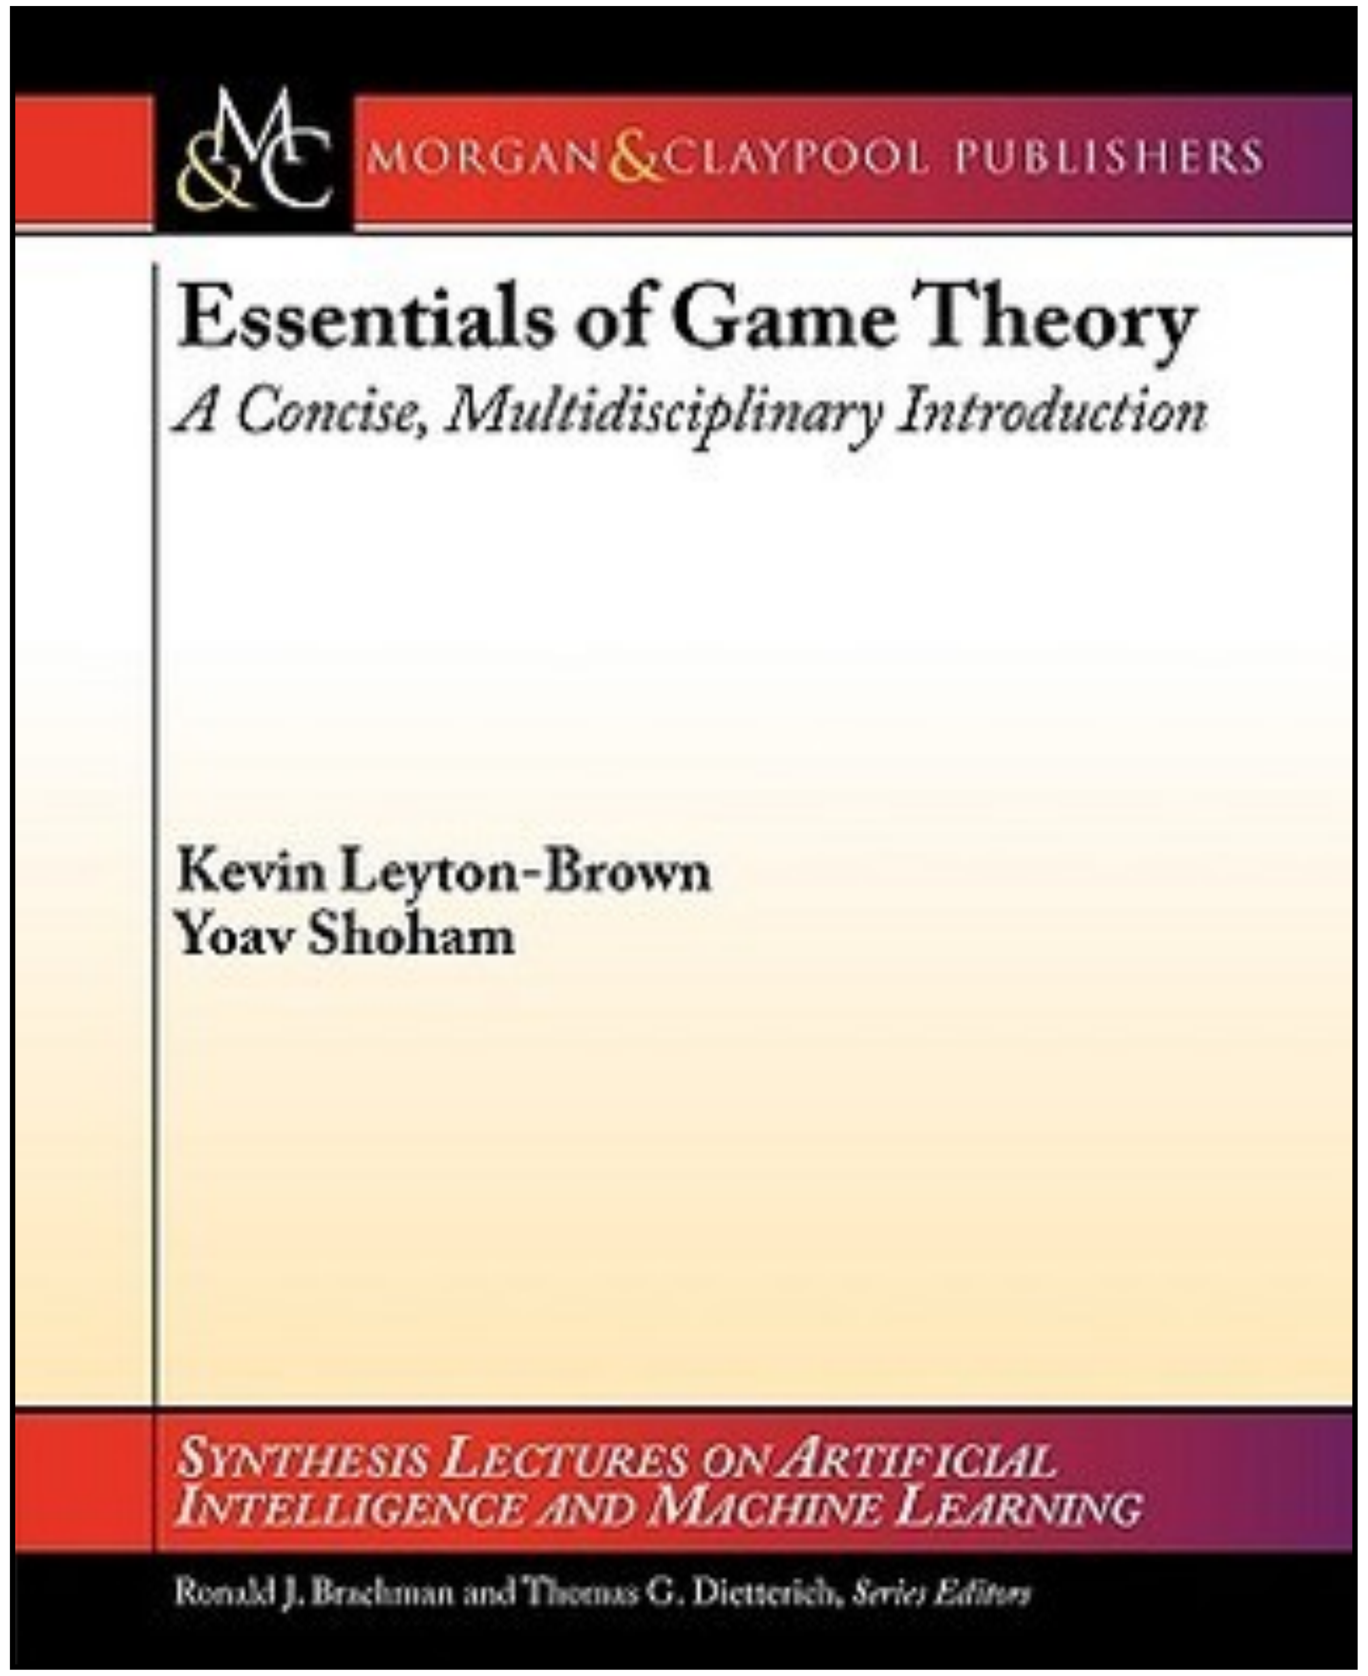

https://www.gtessentials.org/toc.html

Solution concepts; the Nash equilibrium

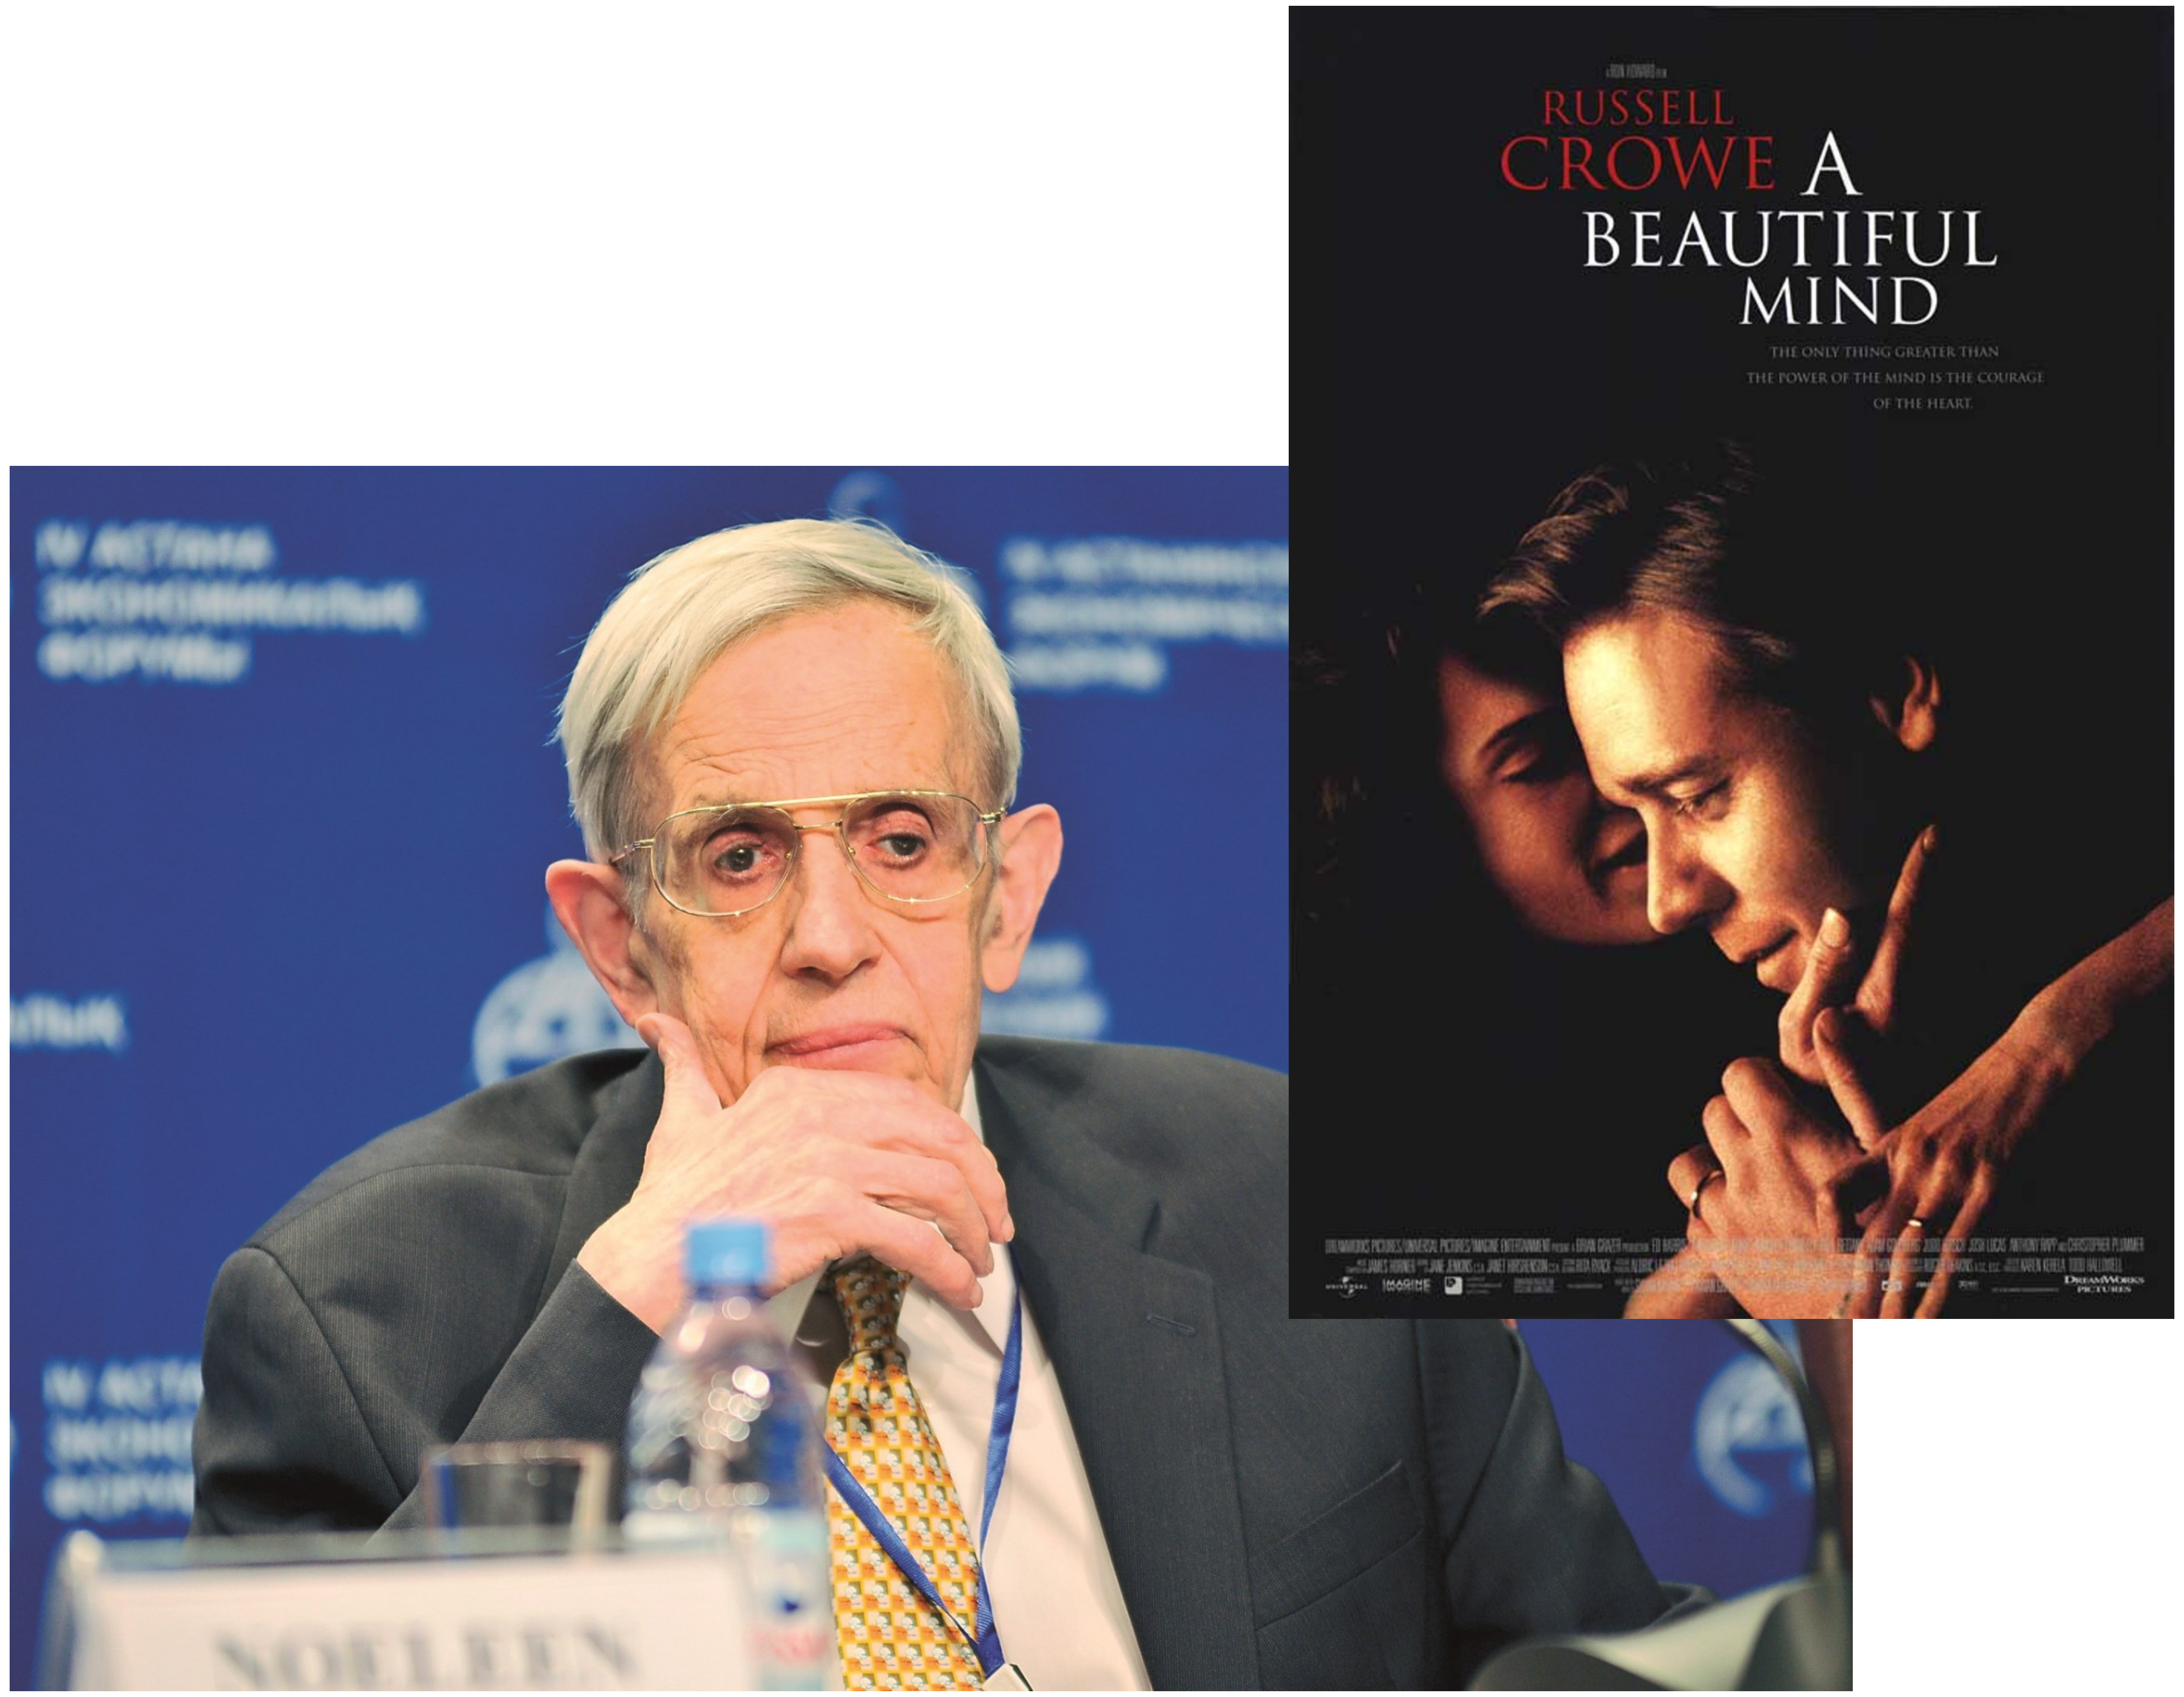

The **Nash** equilibrium

Strategy profile from which no player can increase their utility by deviating unilaterally

Normal form of the game

Sarah

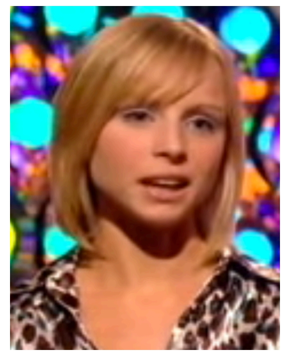

Steve

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Nash Equilibrium

A strategy profile $e^* = (e_1^*, \ldots, e_i^*, \ldots e_N^*)$ in a group of N players is said to be a **Nash equilibrium** if there is no other e such that a single player's change in strategy e_i^* increases her/his personal payoff π_i^*.

- This happens when each equilibrium strategy is a best response to the other ($e_i \in BR(e_{-i}), \forall i$), i.e., strategy e_i^* maximises the expected utility $u_i(e_i^*, e_{-i}^*)$ of player i assuming that the other players adopt strategies $e_{-i}^* = e^* \setminus - \{e_i^*\}$ for all i.

- The equilibrium is strict if $u(e_i^*, e_{-i}^*) > u(e_i, e_{-i})$.

Solution concepts; the Nash equilibrium

Finding the **Nash** equilibrium

A strategy profile e^* is a Nash equilibrium if and only if every player's *i* strategy is a **best response** (B_i) **to the other player's strategies** e_{-i}

e_i^* is in $B_i(e_{-i}^*)$ for every player *i*

A best response is defined as:

$$B_i(e_{-i}) = \{e_i \in E_i : u_i(e_i, e_{-i}) \geq u_i(e_i', e_{-i}) \forall e_i' \in E_i\}$$

Normal form of the game

Sarah

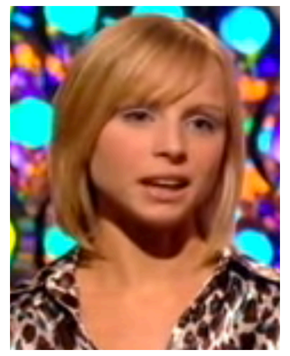

Steve

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Nash equilibria of the game

Strictly dominated strategies can never be part of a NE

Pareto optimality

Pareto optimality refers to an strategic situation in which it is impossible to improve the payoff of one player without worsening the payoff of another player. Formally, in a group of N individuals that adopt a strategy profile $e* = (e_1^*, \ldots, e_i^*, \ldots e_N^*)$, $e*$ is Pareto optimal (or Pareto efficient) if there is no other strategy profile $e = (e_1, \ldots, e_i, \ldots e_N)$ such that:

- $u_i(e) \geq u_i(e*), \forall i \in \{1,...,N\}$
- $u_j(e) > u_j(e*)$, for at least one $j \in \{1,...,N\}$

The notion of optimality in games

The notion of **optimality** in games

Are some outcomes of the game better than others?

Difficult to answer as one cannot rank the interests of players, but …

Pareto **dominance**

The strategy profile e dominates the strategy profile $e^{'}$ if for all players i, $u_i(e) \geq u_i(e^{'})$, and there is some player j for which $u_j(e) > u_j(e^{'})$

This provides a partial ordering over profiles

(Split, Split) vs. (Steal, Steal)?

(Steal, Split) vs. (Steal, Steal)?

(Split, Split) vs. (Steal, Split)?

Normal form
of the game

Sarah

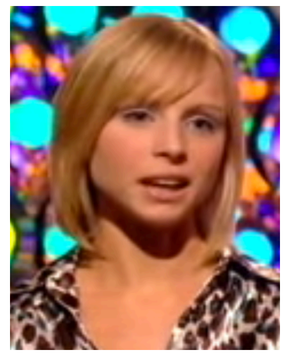

Steve

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

The notion of optimality in games

The notion of **optimality** in games

Are some outcomes of the game better than others?

Difficult to answer if the cannot rank interests of players, but …

Pareto **dominance**

The strategy profile e dominates the strategy profile e' if for all players i, $u_i(e) \geq u_i(e')$, and there is some player j for which $u_j(e) > u_j(e')$

This provides a partial ordering over profiles

Pareto **optimality**

The strategy profile a is Pareto Optimal (efficient) if there is no other strategy profile e' that Pareto dominates e

Normal form
of the game

Sarah

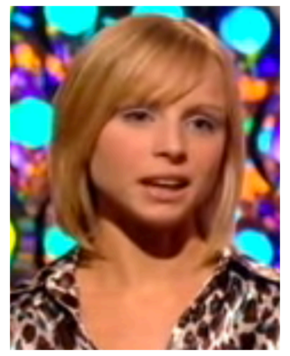

Pareto optimal Solutions

Steve

	SPLIT (Sarah)	STEAL (Sarah)
SPLIT (Steve)	50075£ 50075£	100150£ 0£
STEAL (Steve)	0£ 100150£	0£ 0£

Often limited to the analysis of the NE (here the NE is not Pareto optimal)

The problem of cooperation

Golden balls game

Sarah

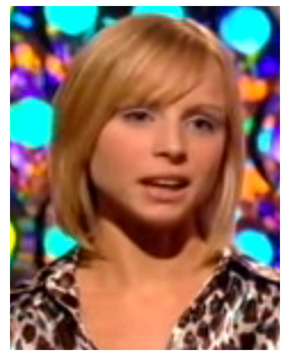

Cooperation

Steve	Sarah: SPLIT	Sarah: STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Steve

Defection

Prisoners Dilemma, T>R, P>S

Greed and fear

	C	D
C	R R	T S
D	S T	P P

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

The problem of cooperation

Golden balls game

Sarah

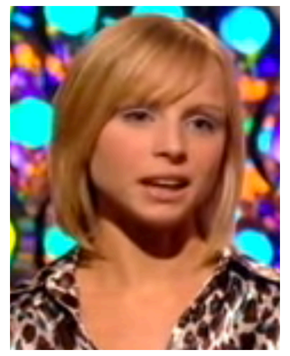

Cooperation

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Steve

Defection

Stag hunt, R>T & P>S

Only fear

	C	D
C	R R	T S
D	S T	P P

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

The problem of cooperation

Golden balls game

Sarah

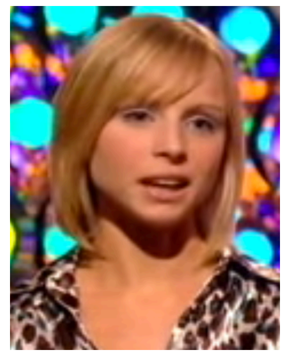

Steve

Cooperation

	SPLIT	STEAL
SPLIT	50075£ 50075£	100150£ 0£
STEAL	0£ 100150£	0£ 0£

Defection

Snow drift, S>P & T>R

Only greed

	C	D
C	R R	**T** **S**
D	**S** **T**	P P

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

	C	D
C	3 / 3	7 / 0
D	0 / 7	1 / 1

NE

	C	D
C	7 / 7	3 / 0
D	0 / 3	1 / 1

NE

NE

	C	D
C	3 / 3	7 / 1
D	1 / 7	0 / 0

NE

NE

Note: D strictly dominates C

More equilibria

“The Big Bang Theory is an American television sitcom created by Chuck Lorre and Bill Prady, both of whom served as executive producers and head writers on the series, along with Steven Molaro. It aired on CBS from September 24, 2007, to May 16, 2019, running for 12 seasons and 279 episodes.” [Wikipedia Oct. 2020]

This Fragment has Sheldon and Raj playing the game Rock-Paper-Scissors-Lizard-Spock to settle a dispute about what to watch on TV (“The Lizard-Spock Expansion” episode, Nov 2008). Game invented by Sam Kass and Karen Bryla (http://www.samkass.com/theories/RPSSL.html)

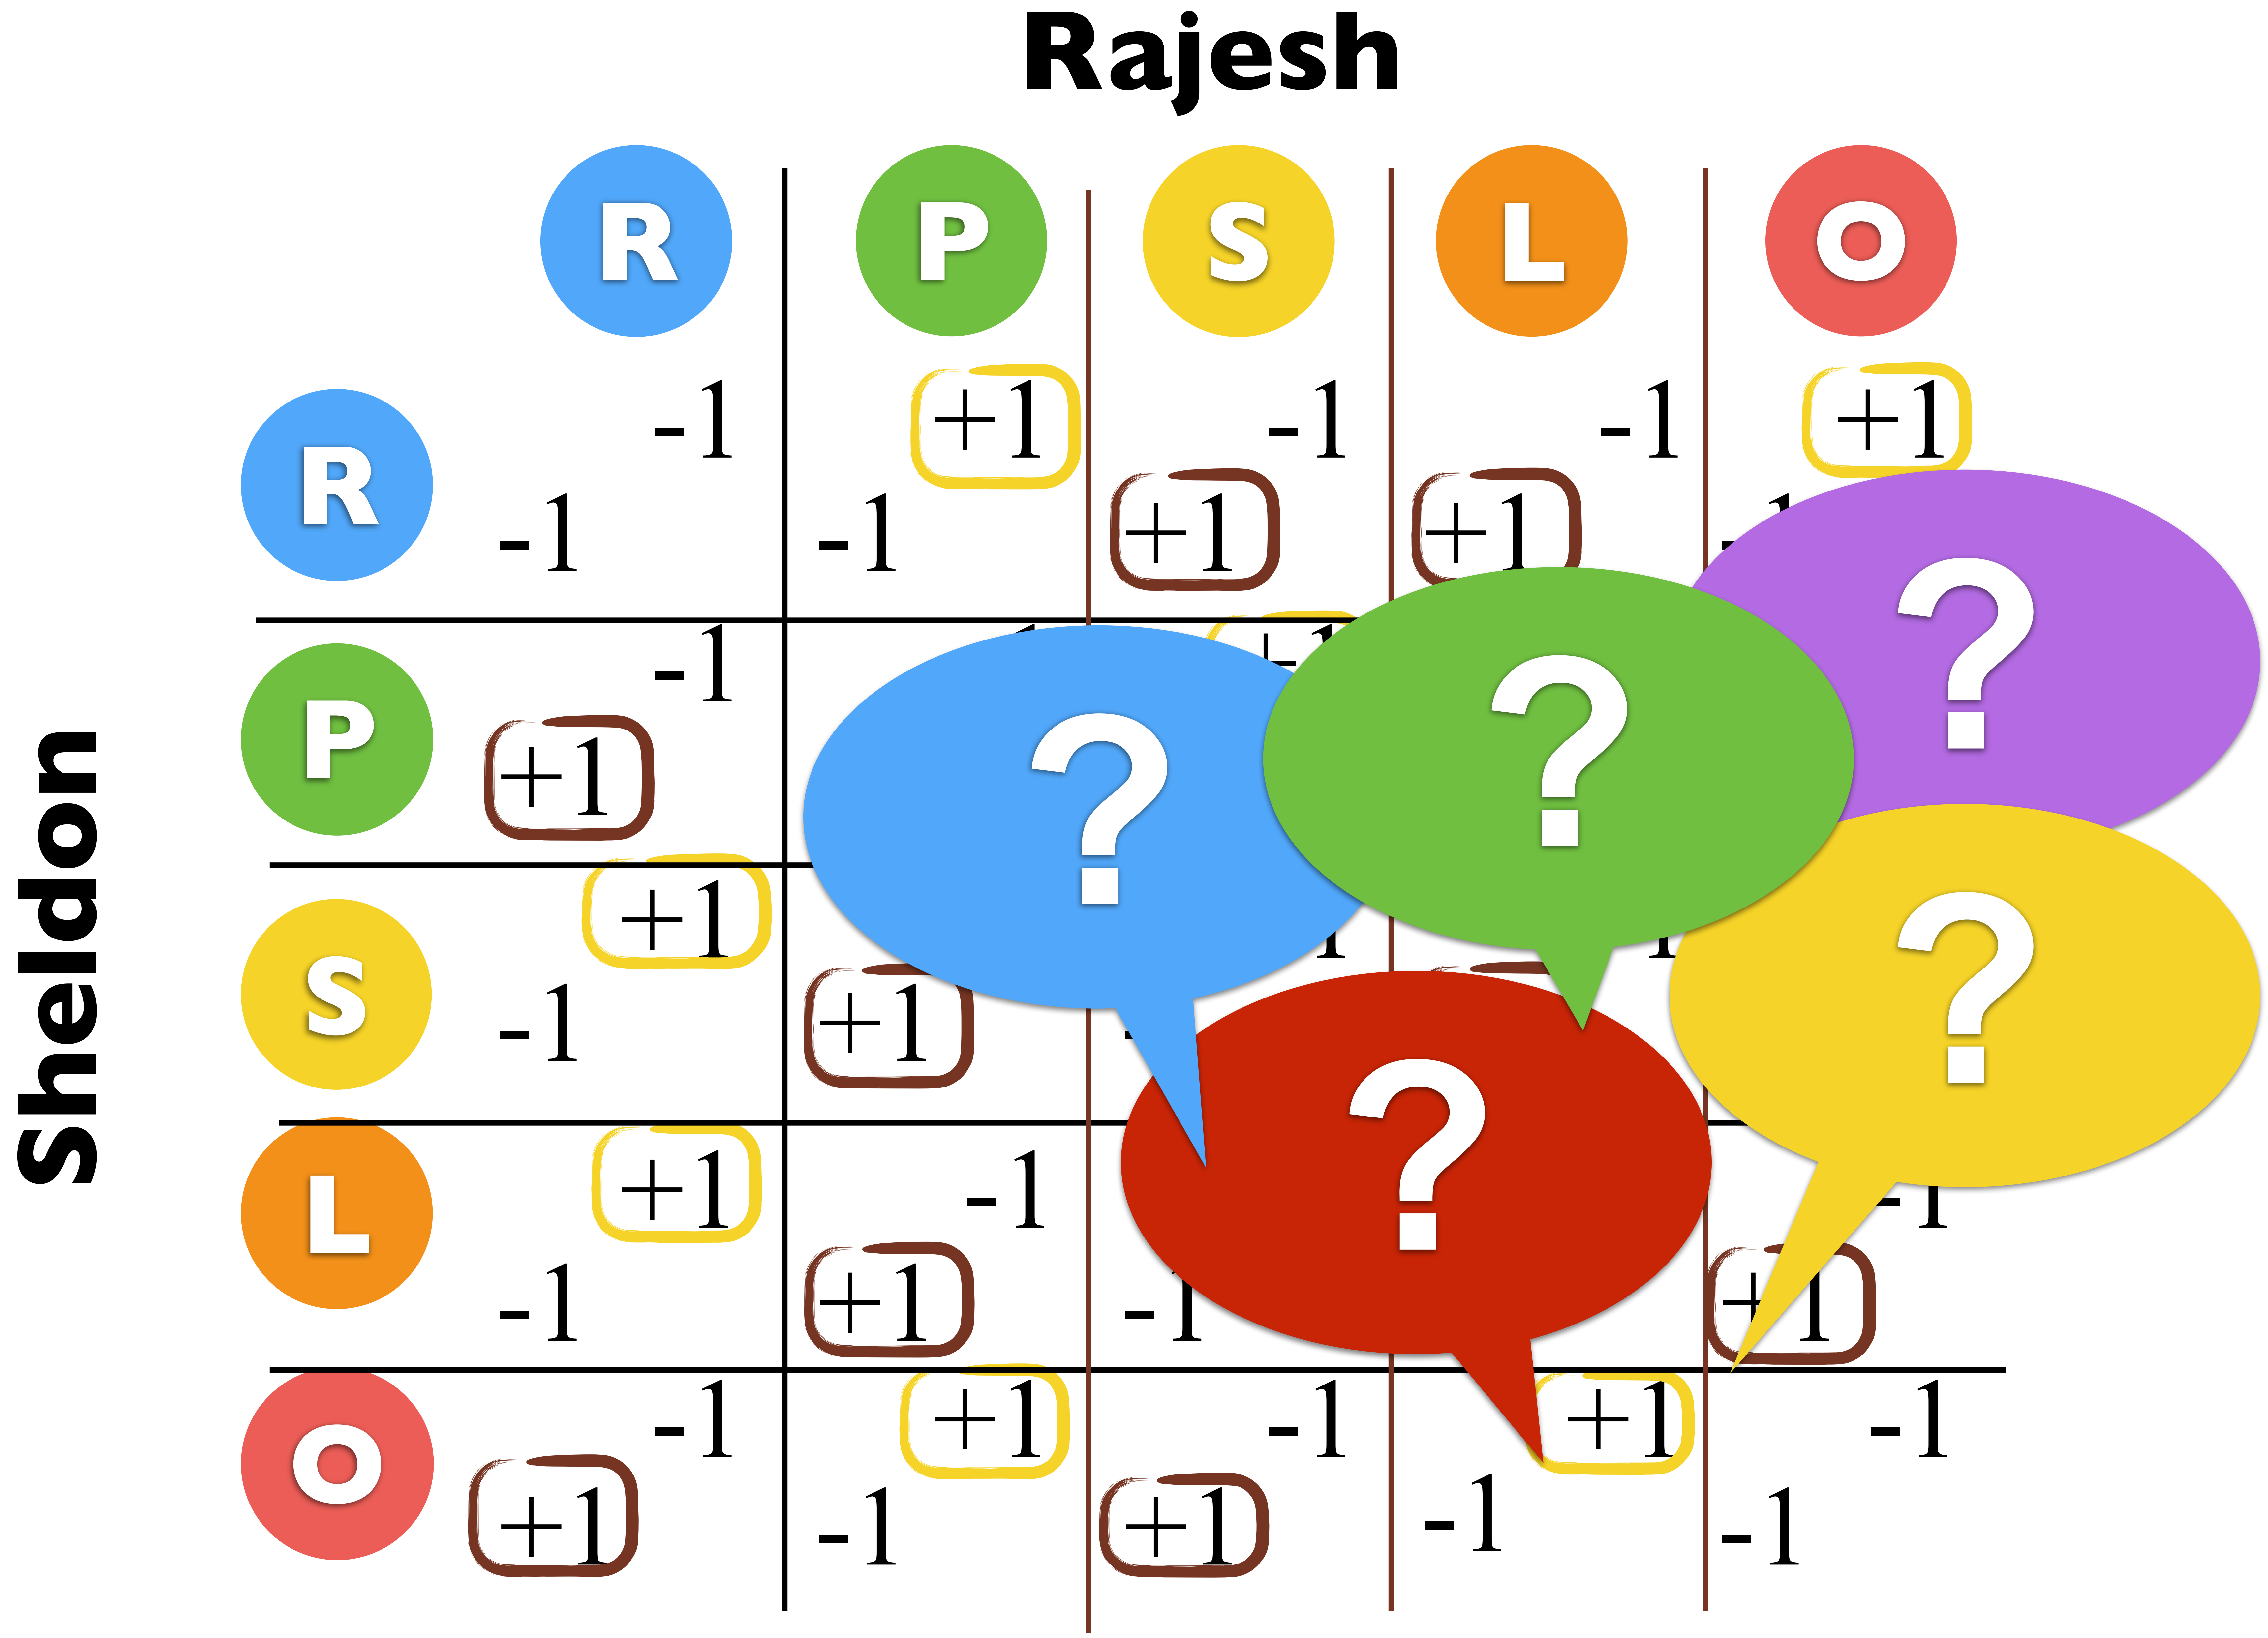

We call this a **zero-sum** game

Every strategic game in which each player has a finite number of actions has at least one Nash equilibrium [Nash 1951]

		Original	
		H	T
Imitator	H	+1, -1	-1, +1
	T	-1, +1	+1, -1

A **mixed** strategy profile

$$e = ((10\,\%\,H, 90\,\%\,T); (70\,\%\,H, 30\,\%\,T))$$

Mixed strategy Nash equilibrium

e^* is a mixed Nash equilibrium if and only if for every player i and for every **mixed strategy** e_i the **expected payoff to** i in e^* is at least as large as the expected payoff to i in (e_i^*, e_{-i}^*) according to the payoff function.

	Original: H	Original: T
Imitator: H	+1, -1	-1, +1
Imitator: T	-1, +1	+1, -1

A **mixed** strategy profile

$$e = ((10\,\%\,H, 90\,\%\,T); (70\,\%\,H, 30\,\%\,T))$$

Mixed NE

A mixed strategy Nash equilibrium is a strategy profile $e^* = (e_1^*, e_2^*, \ldots, e_n^*)$ such that for each player i, the mixed strategy e_i^* maximises the player's expected payoff, assuming the strategies of the other players are fixed. That is:

$$\Pi(e_i^*, e_{-i}^*) \geq \Pi(e_i, e_{-i}^*)$$

Original

Imitator \ Original	q H	T $1-q$
p H	+1, -1	-1, +1
T $1-p$	-1, +1	+1, -1

Expected payoffs for imitator (I) …

$$\Pi_I = p\underbrace{(q\ \pi_I(H,H) + (1-q)\ \pi_I(H,T))}_{\Pi_I(q\,|\,H)} +$$
$$(1-p)\underbrace{(q\ \pi_I(T,H) + (1-q)\ \pi_I(T,T)}_{\Pi_I(q\,|\,T)}$$

$$\Pi_I = p\ \Pi_I(q\,|\,H) + (1-p)\ \Pi(q\,|\,T)$$

… and original (O)

$$\Pi_O = q(p\ \pi_O(H,H) + (1-p)\ \pi_O(H,T))+$$
$$(1-q)(p\ \pi_O(T,H) + (1-p)\ \pi_O(T,T)$$

$$\Pi_O = q\ \Pi_O(p\,|\,H) + (1-q)\ \Pi_O(p\,|\,T)$$

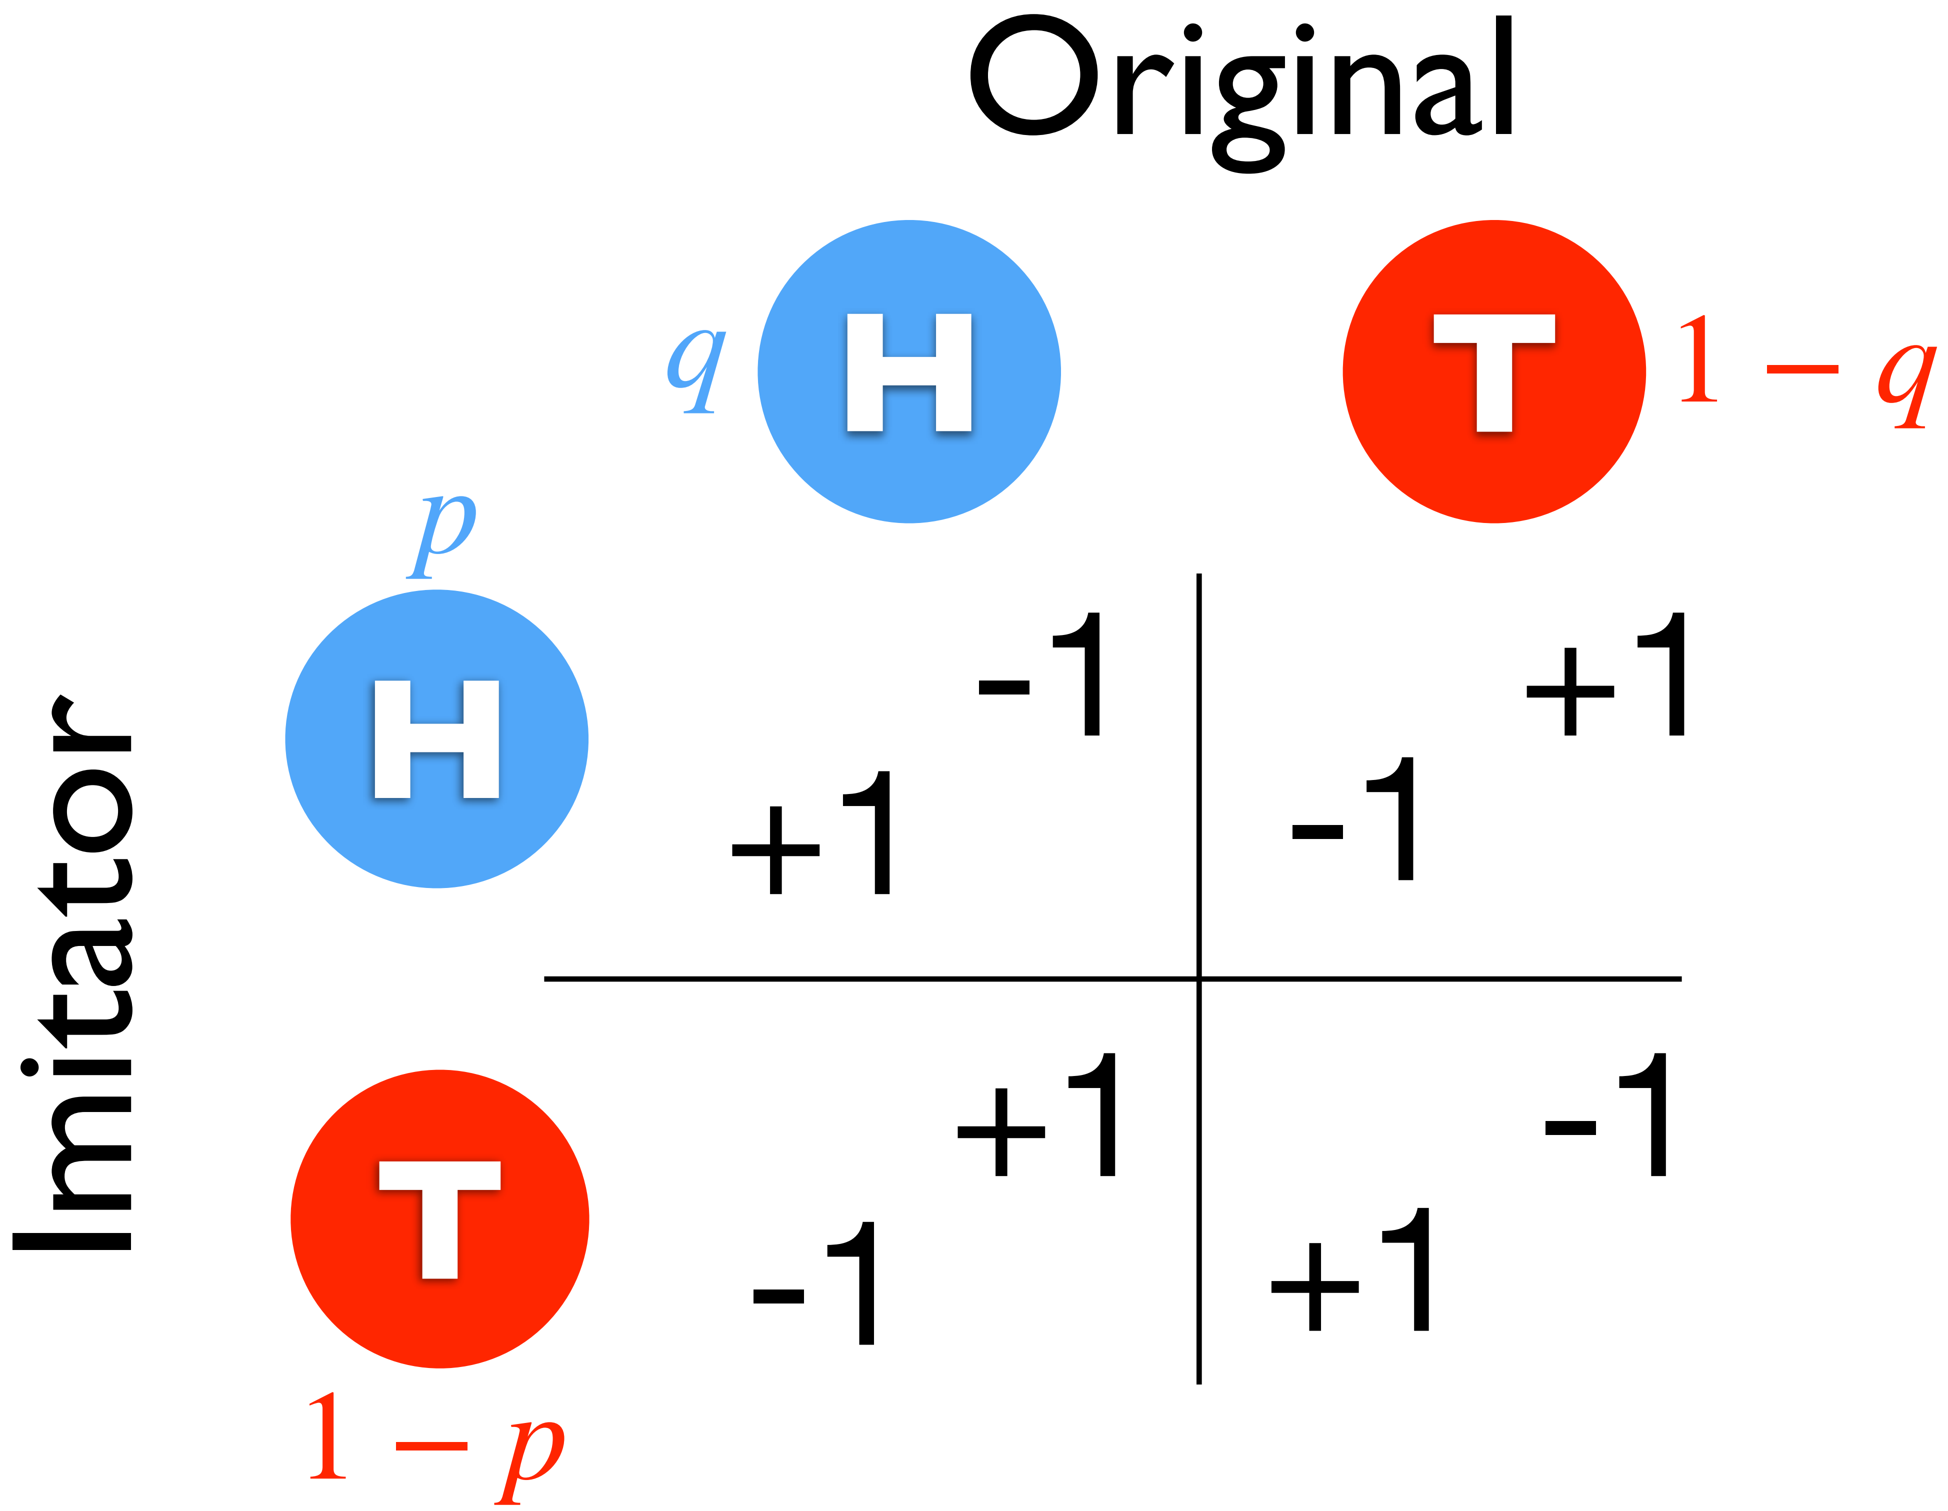

Finding the mixed Nash equilibrium

The mixed strategy profile e^* is a Nash equilibrium if and only if e_i^* is in $B_i(e_{-i}^*)$ for every player i

What is the set $B_{imitator}$ for player "Imitator" ?

H $\Pi_I(q\,|\,H) > \Pi_I(q\,|\,T)$

T $\Pi_I(q\,|\,H) < \Pi_I(q\,|\,T)$

H T $\Pi_I(q\,|\,H) = \Pi_I(q\,|\,T)$

Equivalent for $B_{original}$

Original

Imitator	q H	T $1 - q$
p H	+1, -1	-1, +1
T $1 - p$	-1, +1	+1, -1

What is the set $B_{imitator}$ for player "Imitator" ?

H

$$\Pi_I(q\,|\,H)) > \Pi_I(q\,|\,T)$$

$$q - (1 - q) > -q + (1 - q)$$

$$2q - 1 > 1 - 2q$$

$$q > \frac{1}{2}$$

T

$$\Pi_I(q\,|\,H) < \Pi_I(q\,|\,T)$$

$$q < \frac{1}{2}$$

H T

$$q = \frac{1}{2}$$

What is the set $B_{imitator}$ for player "Imitator" ?

Original

Imitator		q H	T $1-q$
	p H	+1, -1	-1, +1
	T $1-p$	-1, +1	+1, -1

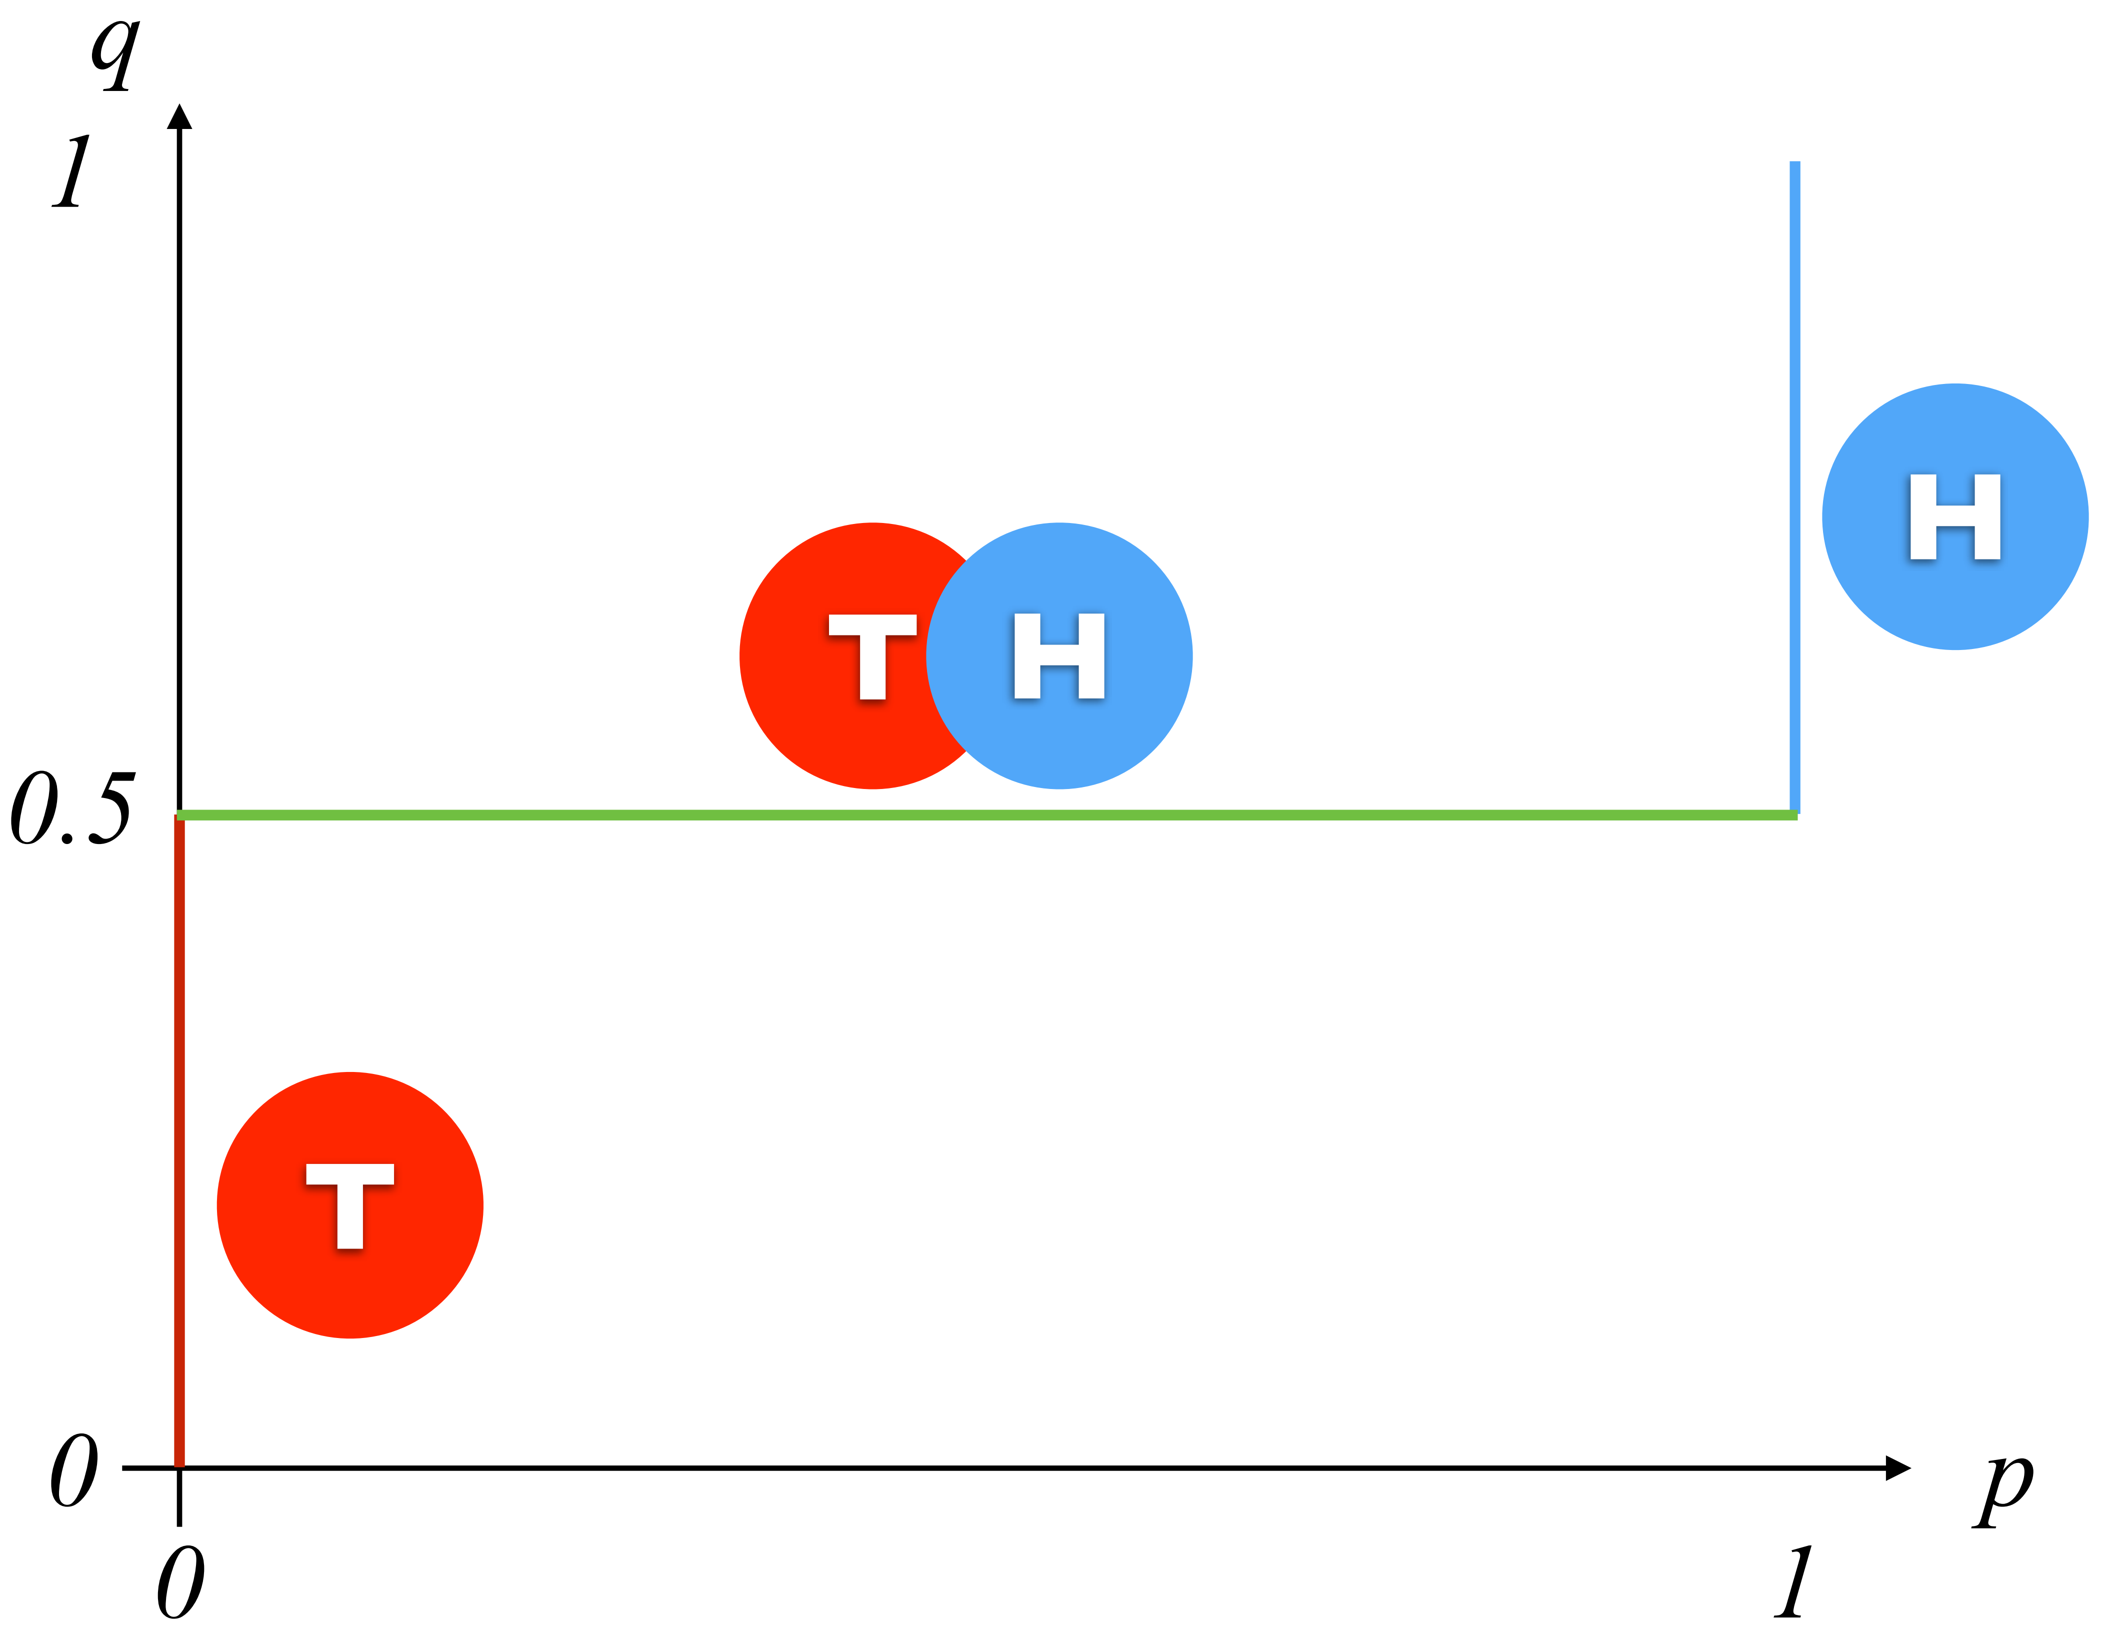

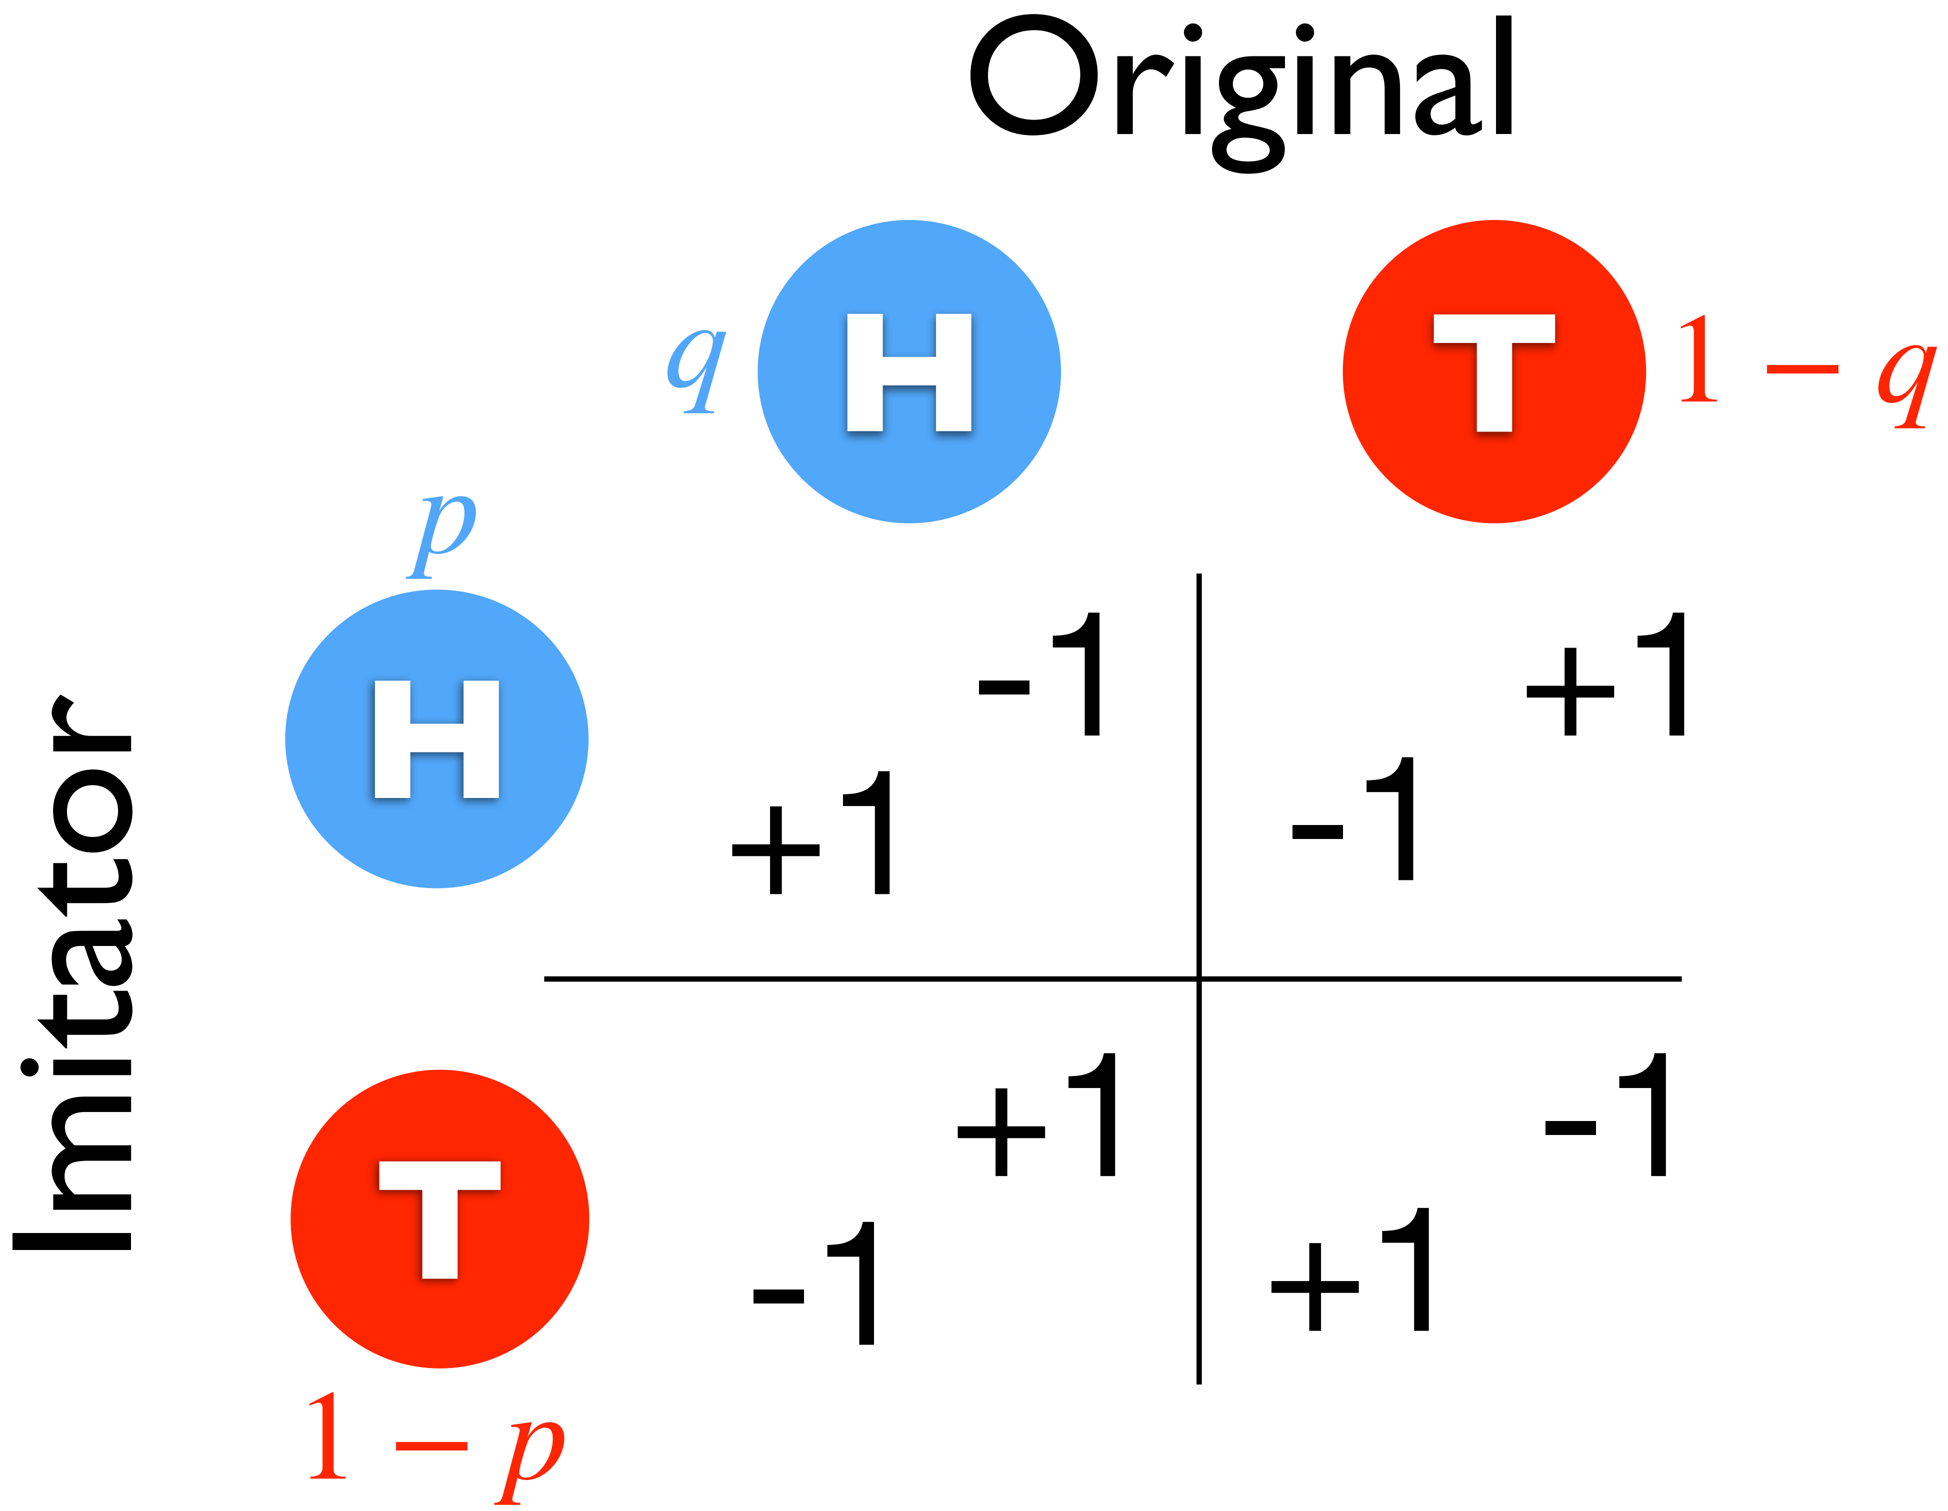

What is the set $B_{original}$ for player “Original” ?

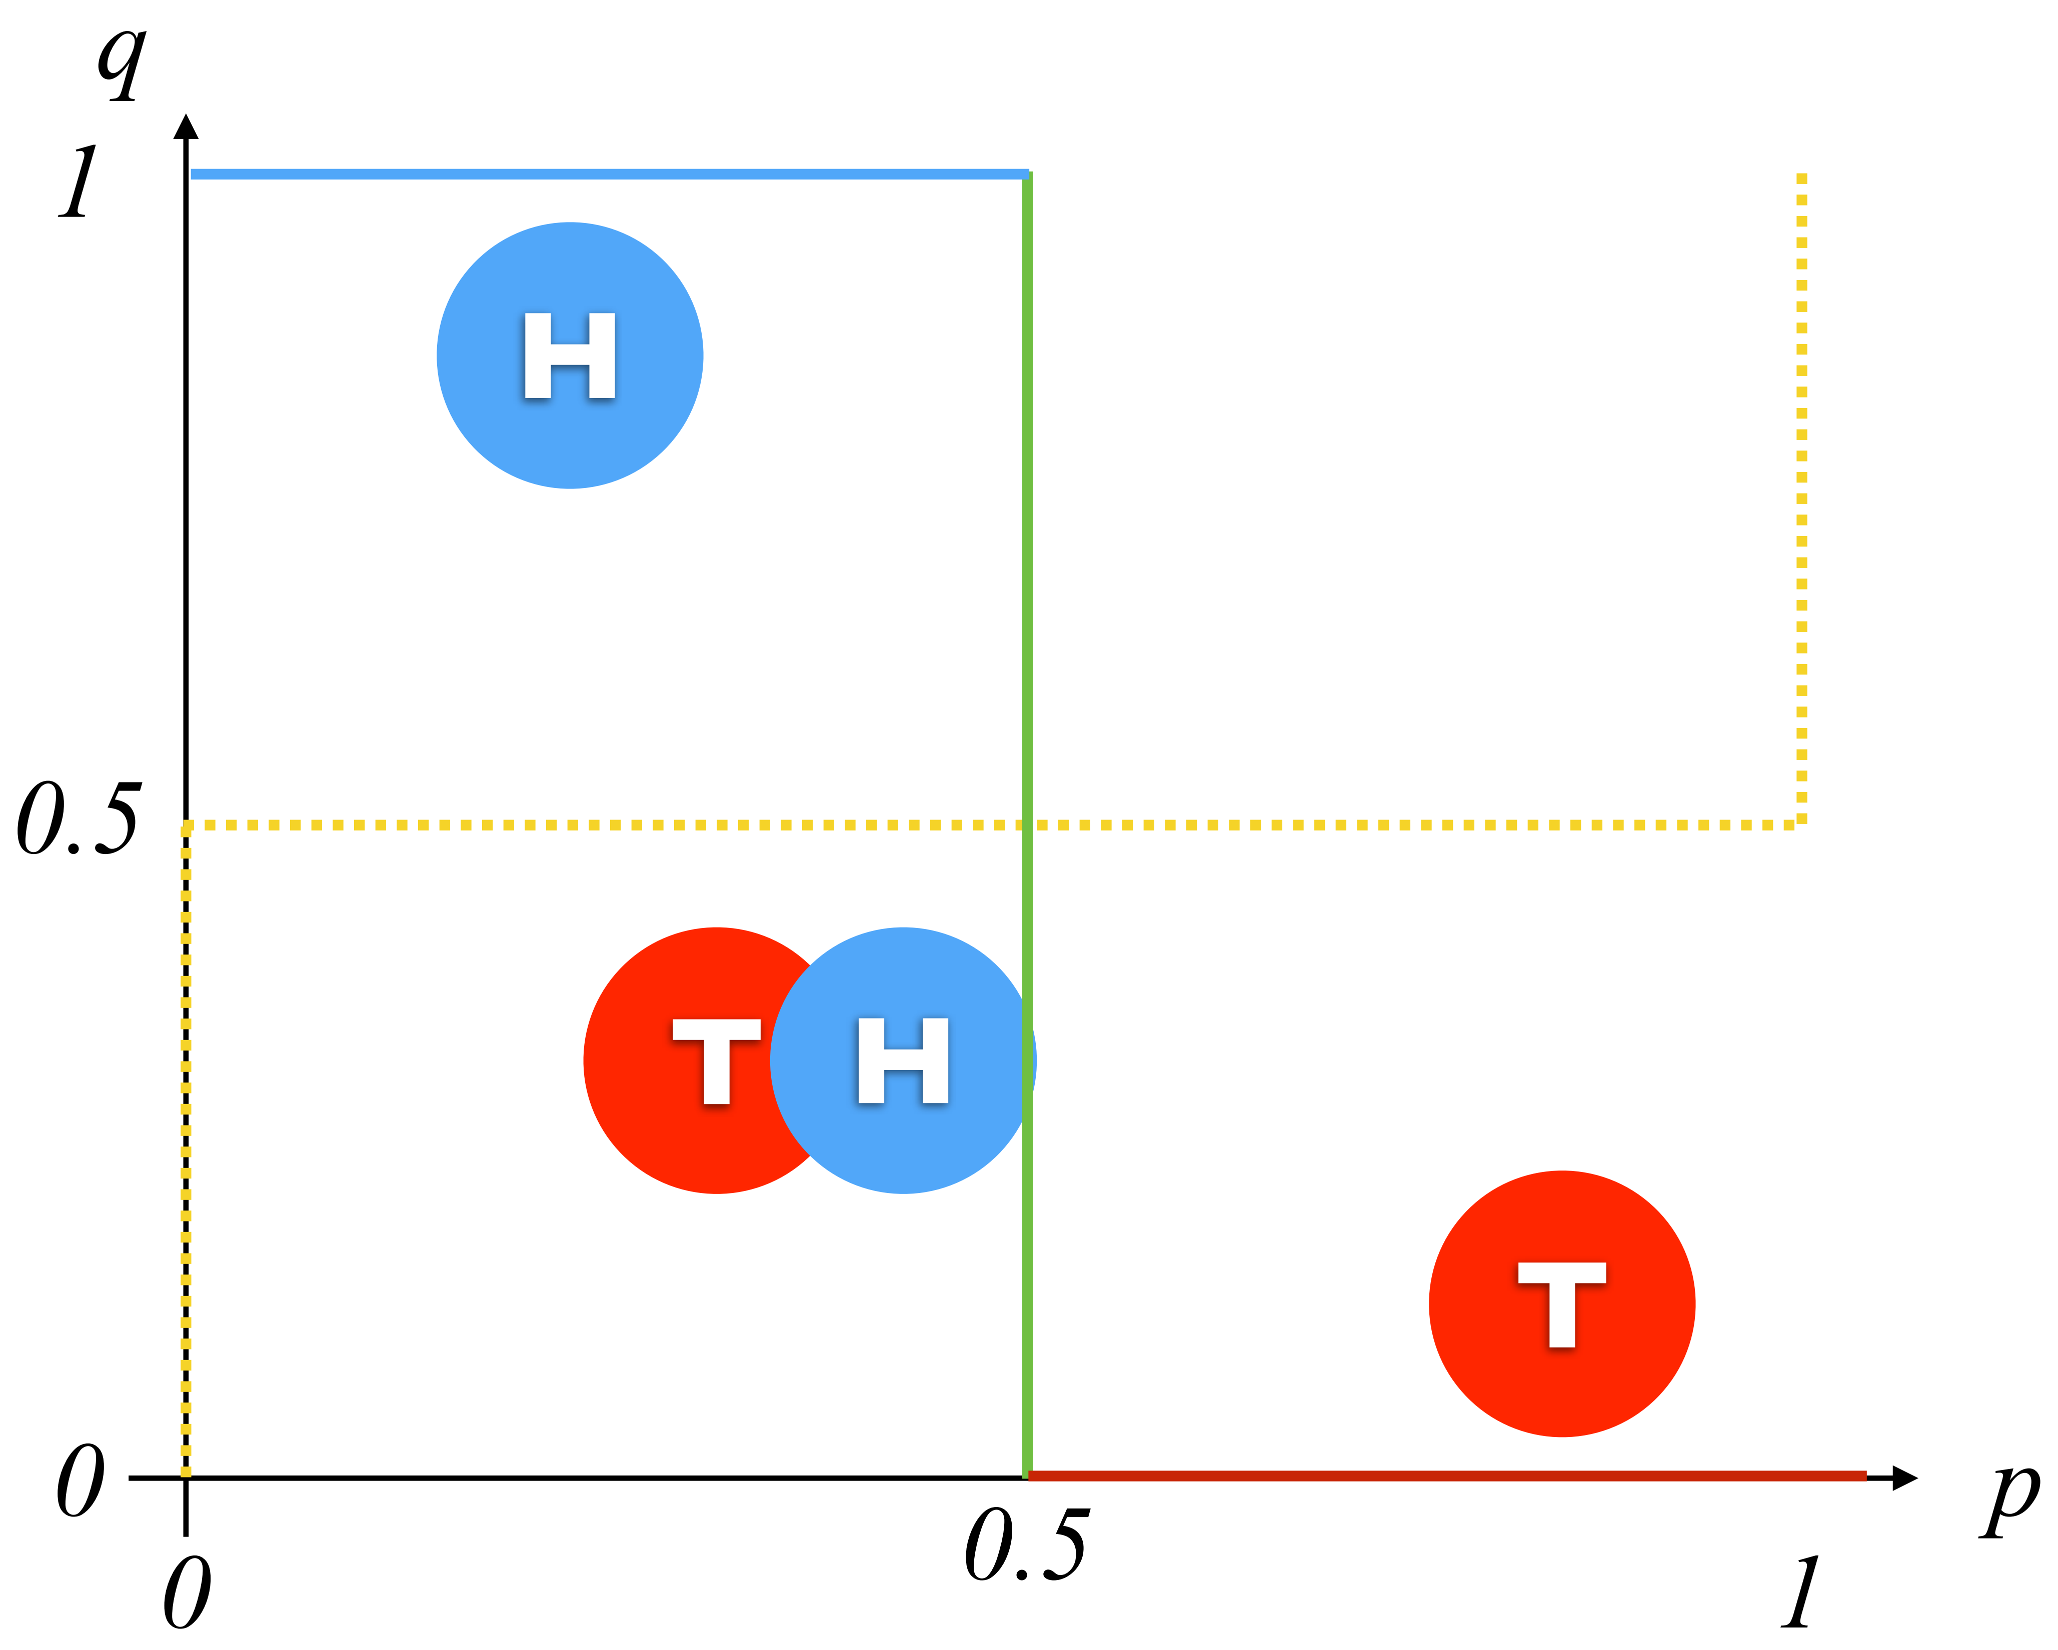

Finding the mixed strategy Nash equilibrium

Original

		H (q)	T ($1-q$)
Imitator	H (p)	+1, -1	-1, +1
	T ($1-p$)	-1, +1	+1, -1

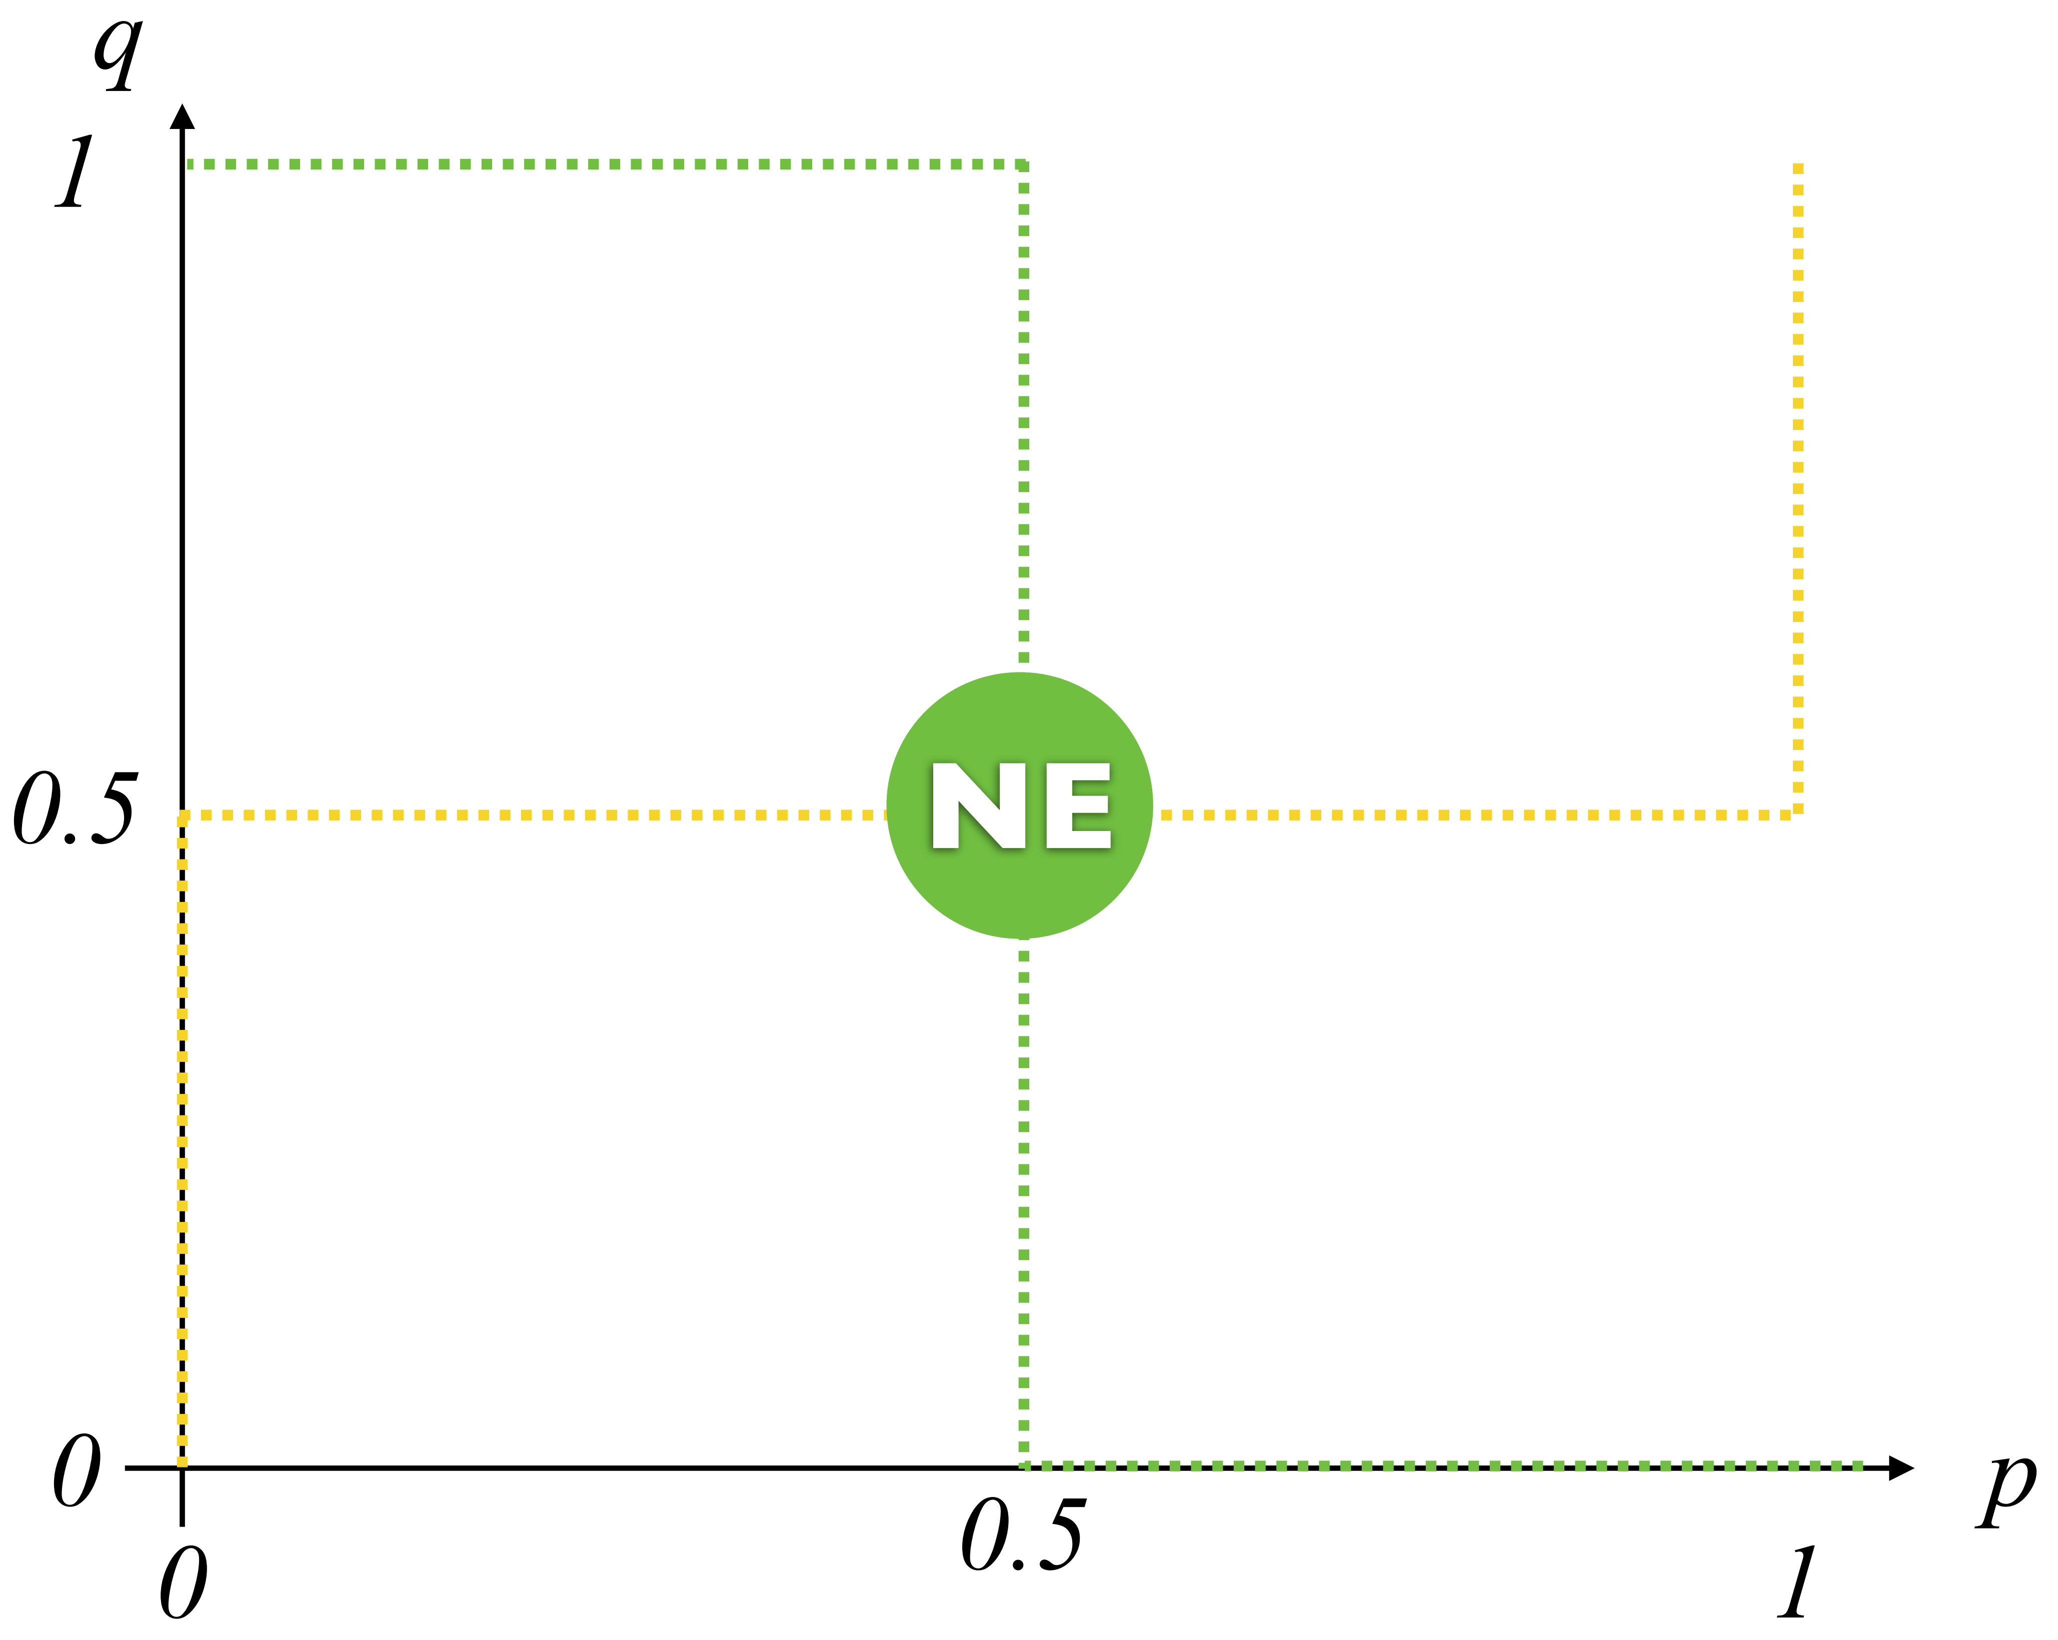

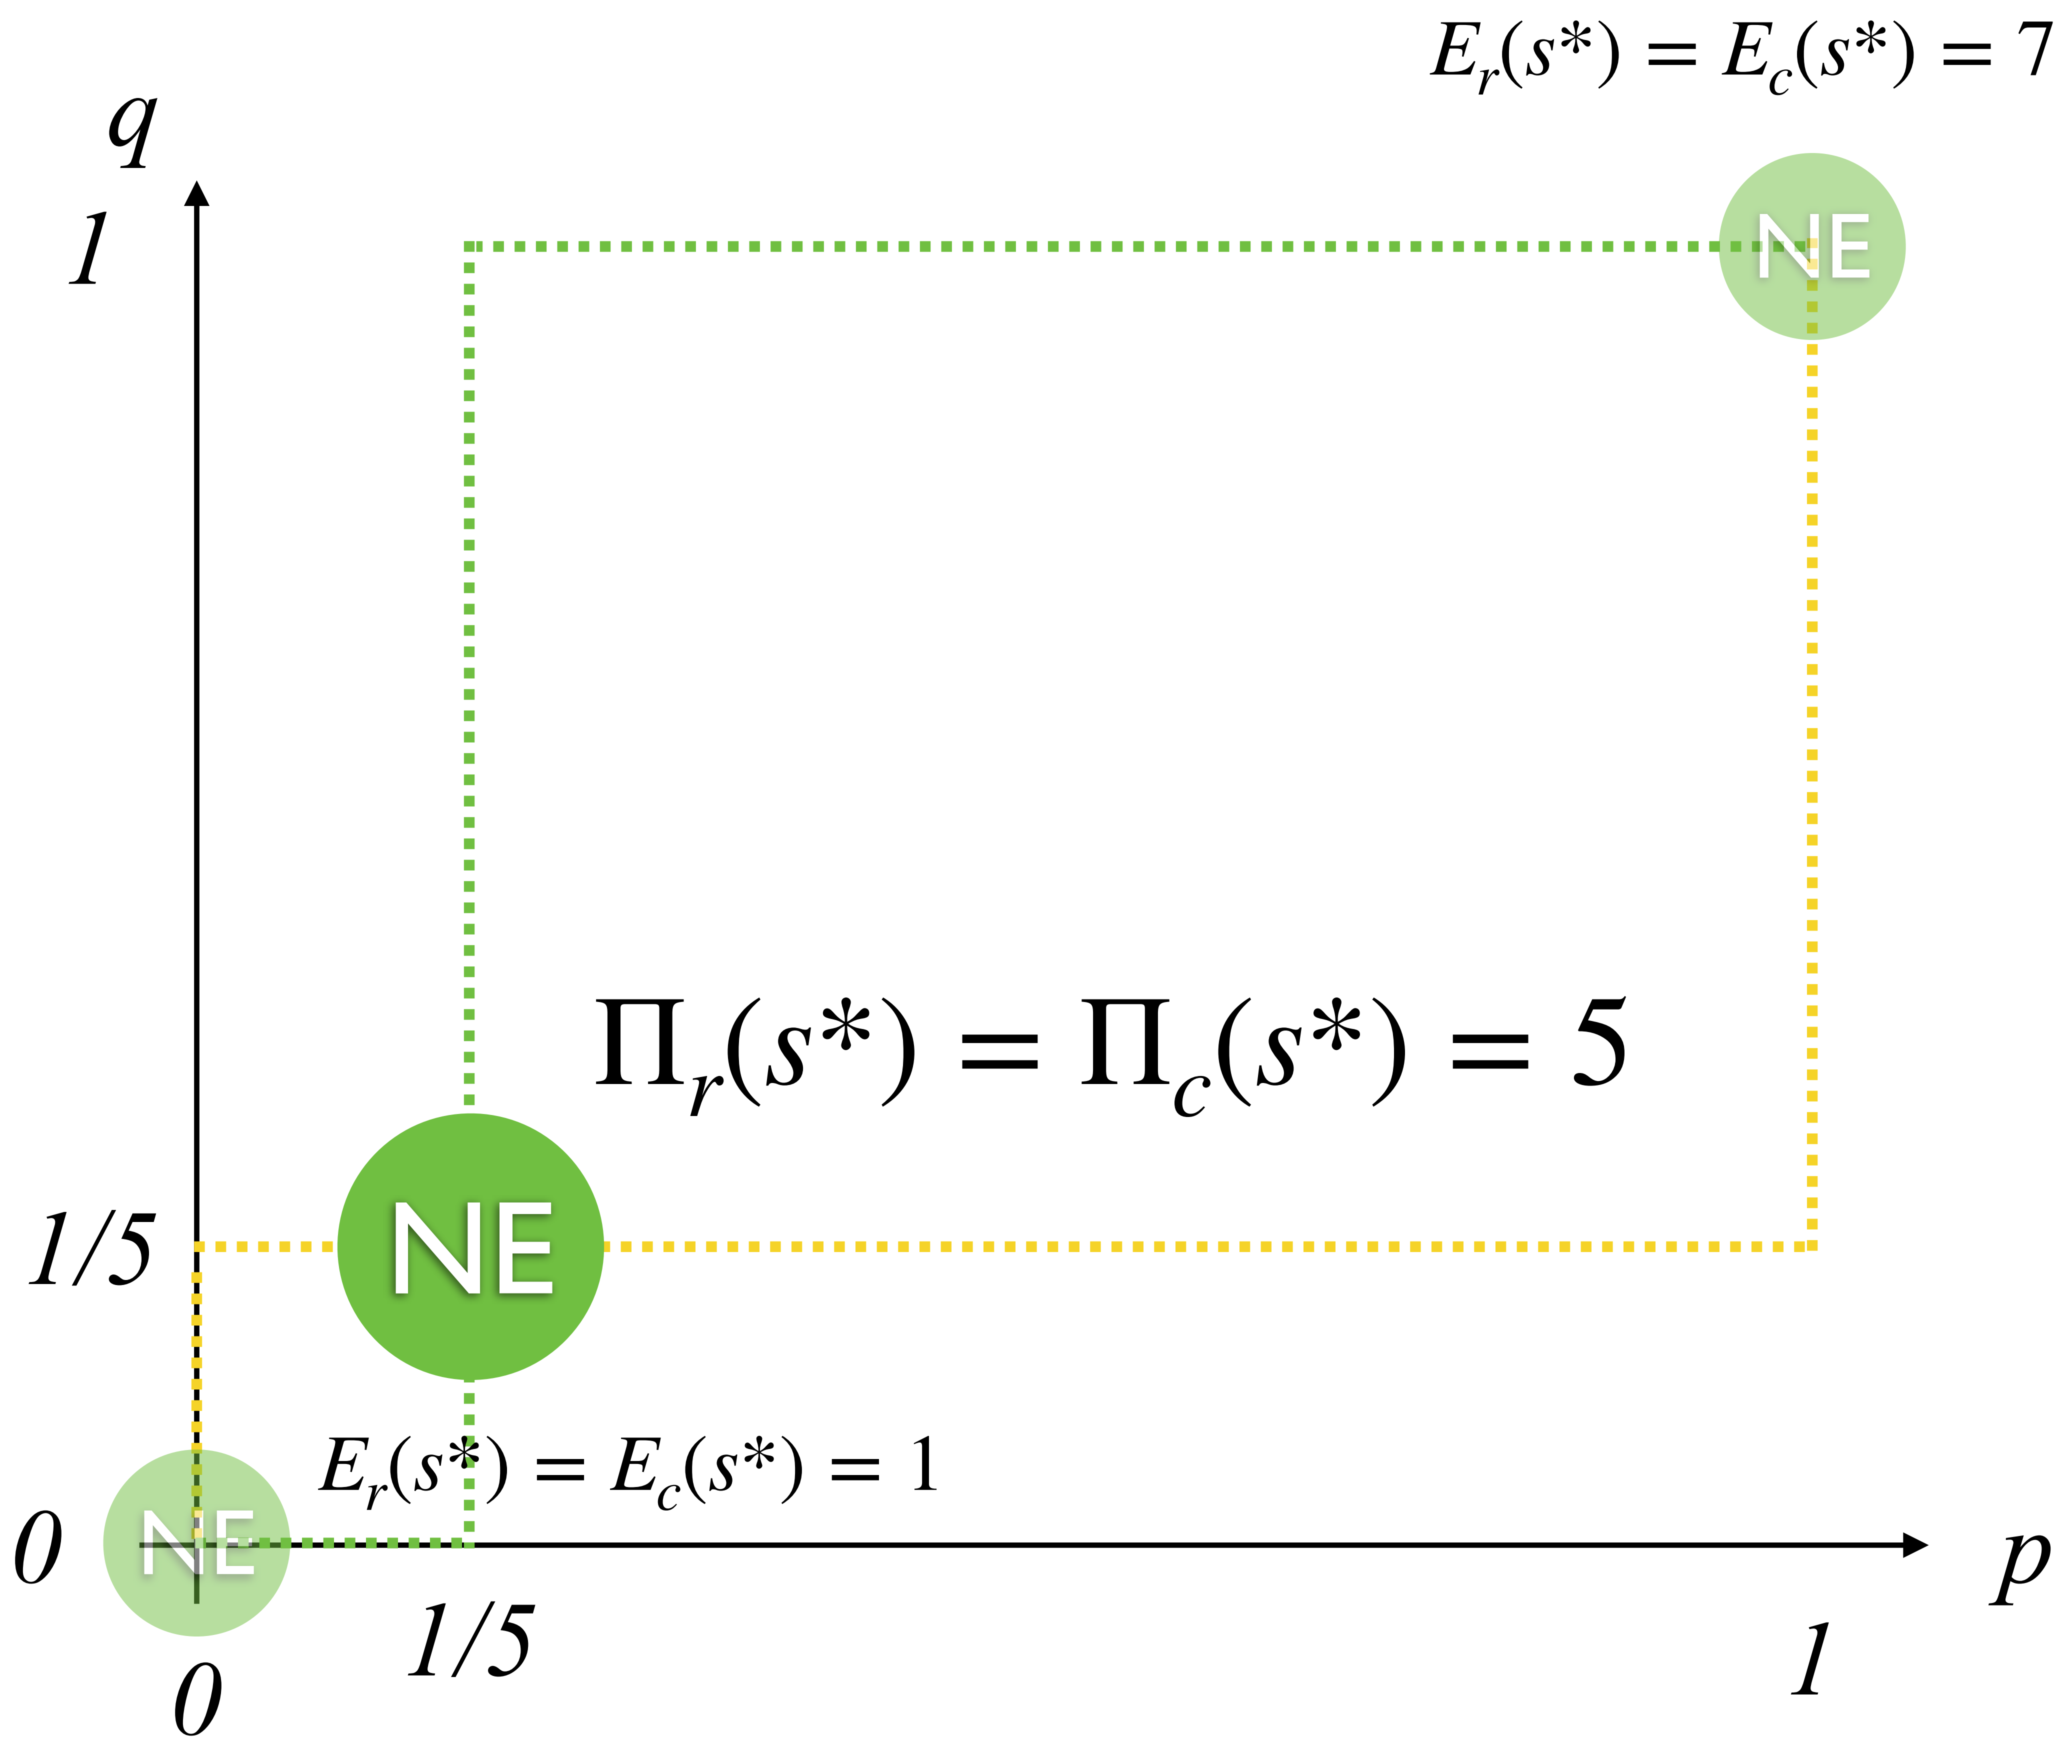

Stag hunt, R>T & P>S

Only fear

	C	D
C	7, 7	0, 3
D	3, 0	1, 1

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

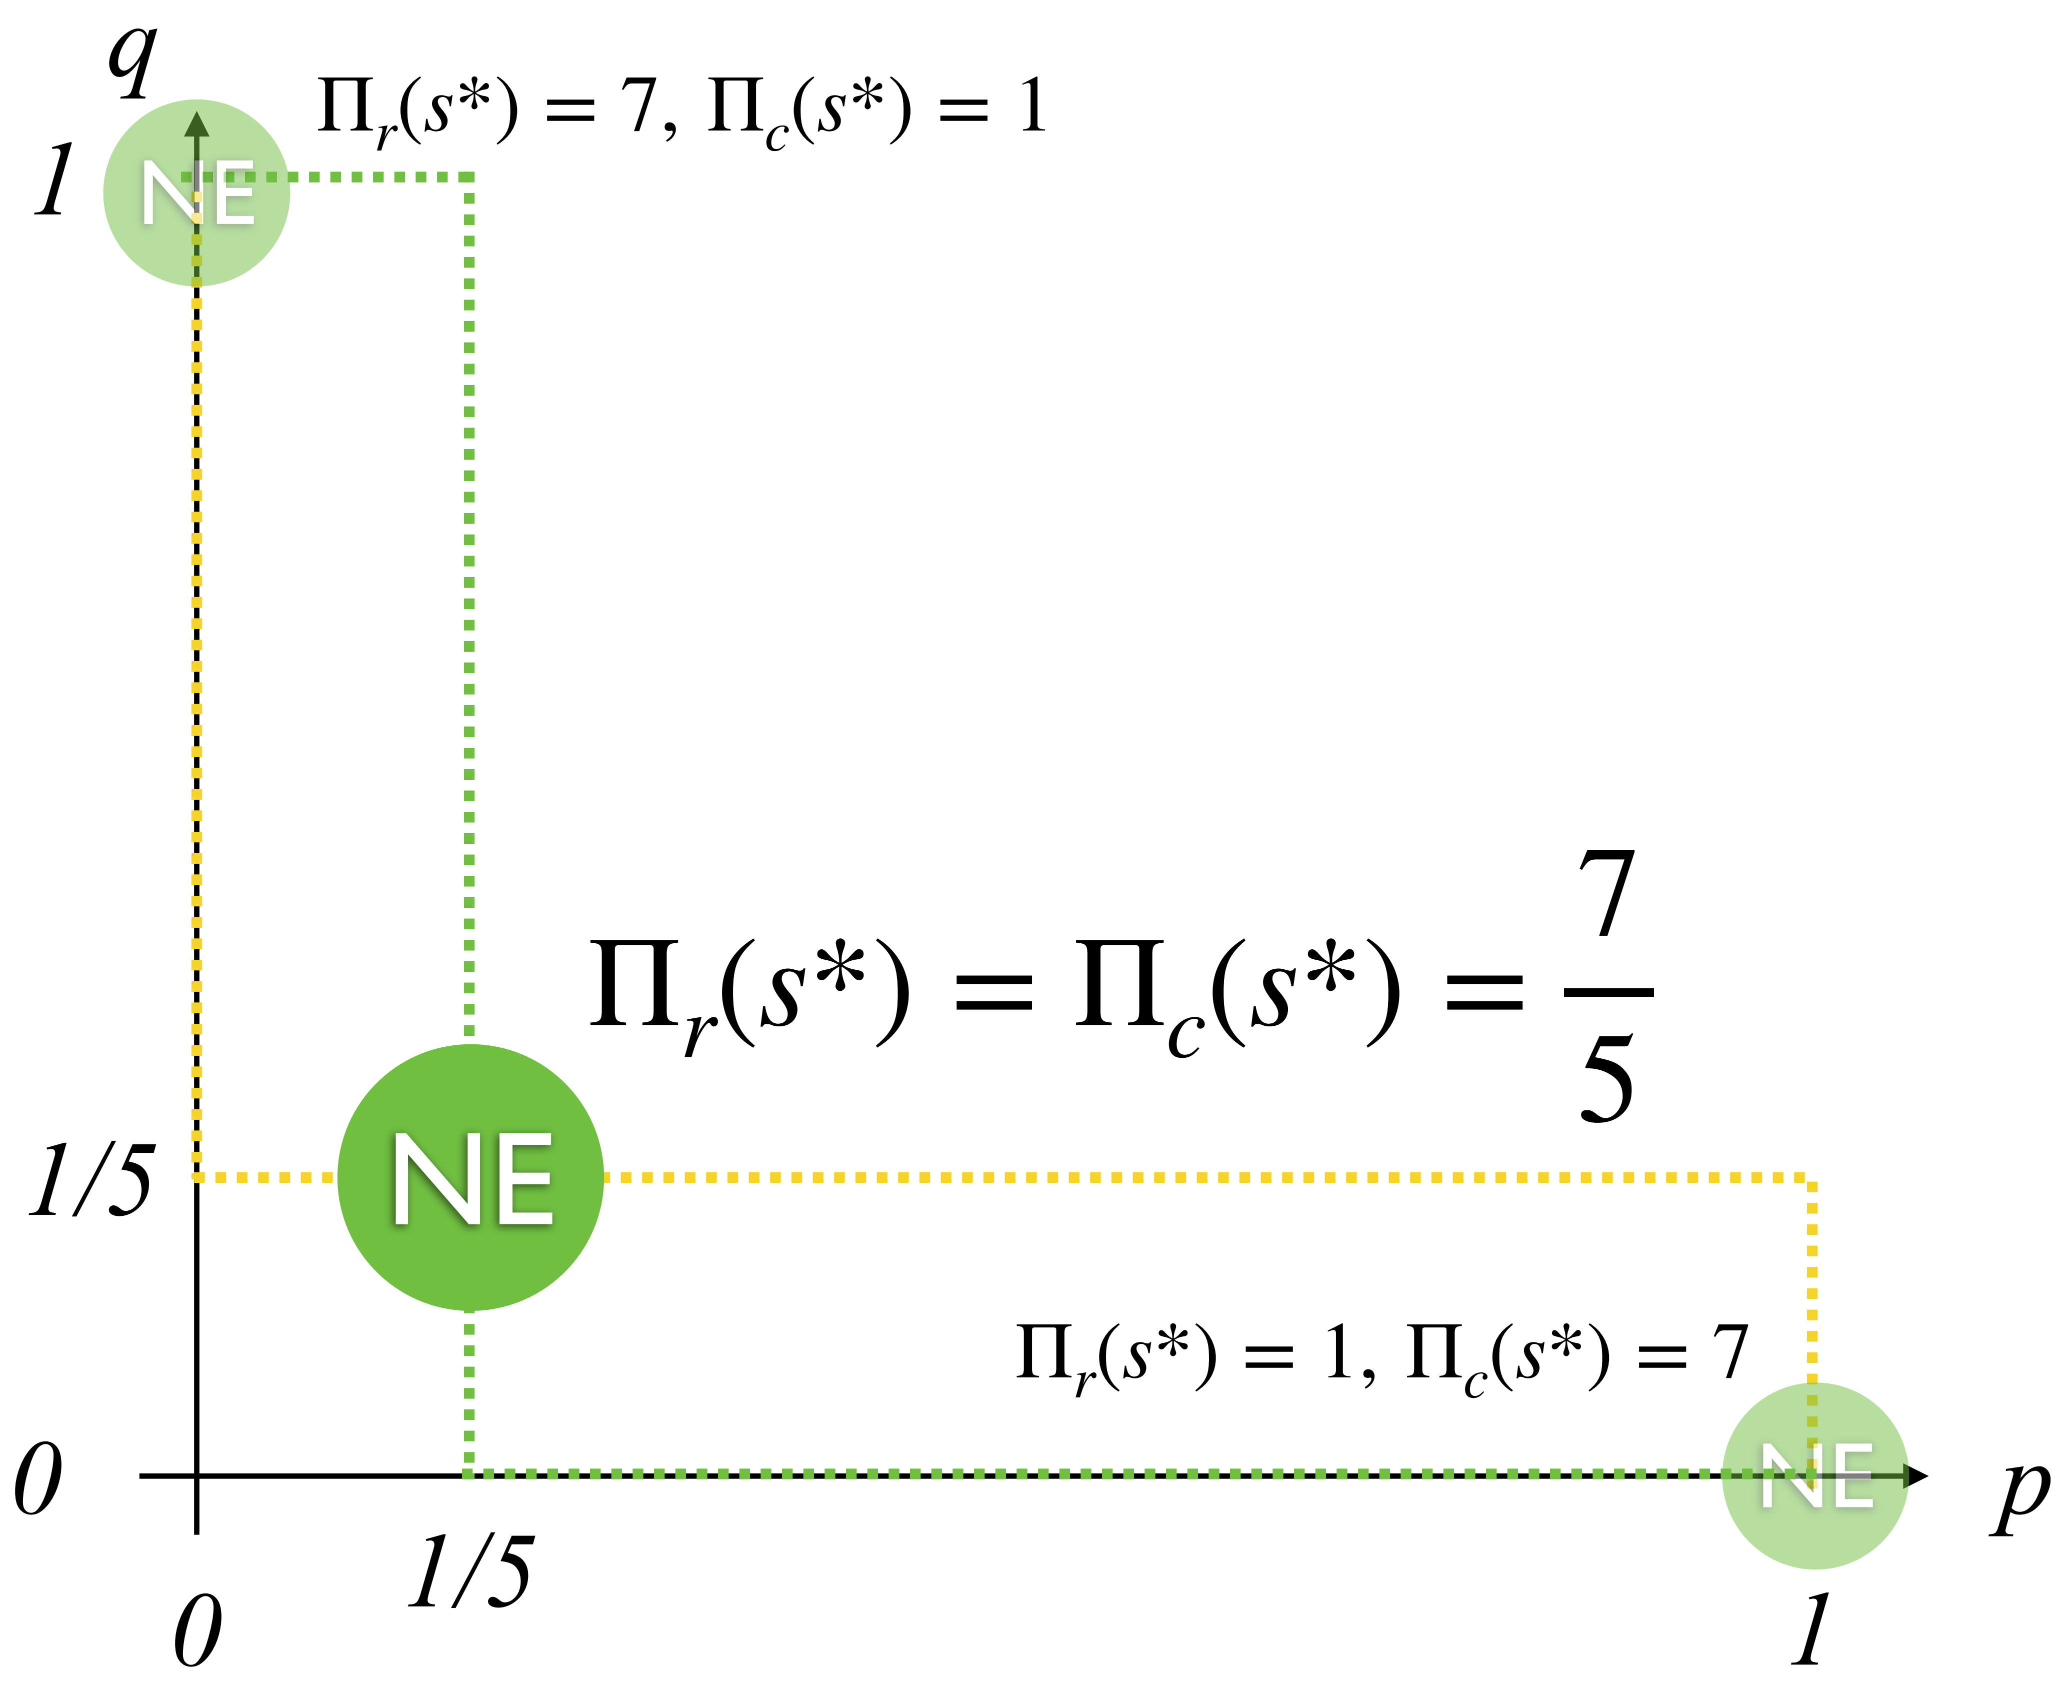

Snow drift, S>P & T>R

Only greed

	C	D
C	3, 3	1, 7
D	7, 1	0, 0

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

More general algorithms to identify mixed NE

A combinatorial optimisation problem

Remember: The mixed strategy profile e^* is a Nash equilibrium if and only if e_i^* is in $B_i(e_{-i}^*)$ for every player i

A mixed strategy is a best response if and only if **all pure strategies in its support are best responses**

Finding the NE is thus equivalent to find the pure strategies that are in the support

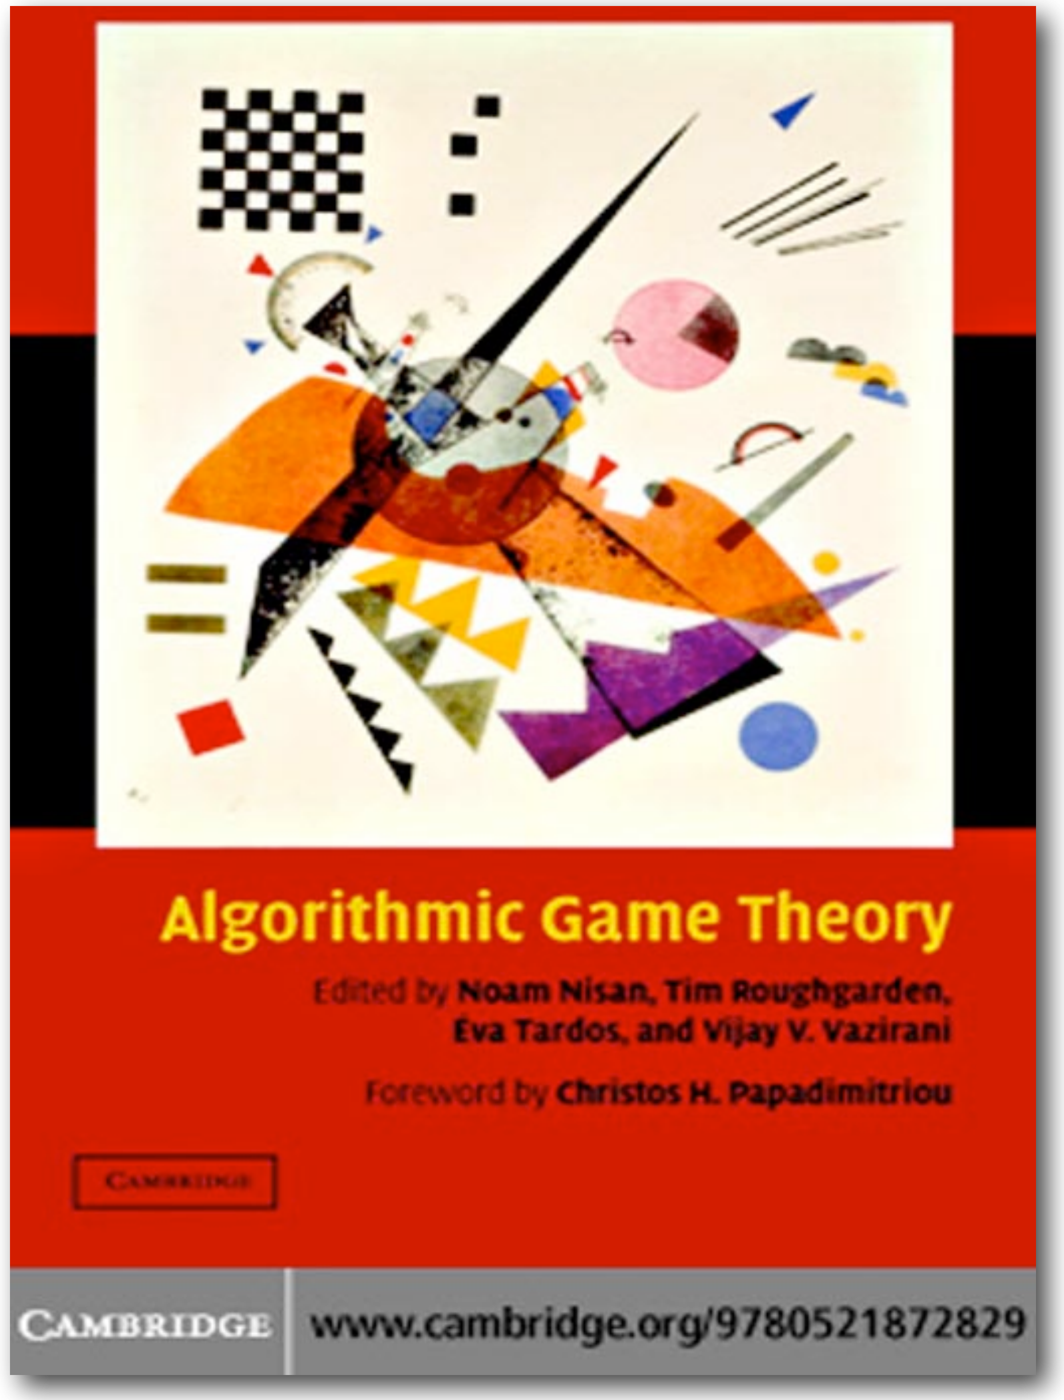

Support finding

Vertex enumeration

Lemke-Howson algorithm

Nashpy demo

The Journal of Open Source Software

Nashpy: A Python library for the computation of Nash equilibria

Vincent Knight[1] **and James Campbell**[1]

1 Cardiff University, School of Mathematics, UK

DOI: 10.21105/joss.00904

Software

- Review
- Repository
- Archive

Submitted: 31 May 2018
Published: 10 October 2018

Summary

Game theory is the study of strategic interactions where the outcomes of choice depend on the choices of all participants. A key solution concept in the field is that of Nash Equilibrium (Nash & others, 1950). This solution concept corresponds to a coordinate at which no participant has any incentive to change their choice.

As an example, consider the game of Rock Paper Scissors, which can be represented mathematically using the following matrix:

$$A = \begin{pmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{pmatrix}$$

Knight and Campbell, (2018). Nashpy: A Python library for the computation of Nash equilibria. Journal of Open Source Software, 3(30), 904, https://doi.org/10.21105/joss.00904

See also : https://nashpy.readthedocs.io/en/stable/index.html#

jupyter nashpy-demo Last Checkpoint: 13 minutes ago

File Edit View Run Kernel Settings Help

Code JupyterLa

```
[10]: #Loading the necessary libraries

[12]: import nashpy as nash
      import numpy as np

[ ]: #Define row and column matrices and initialise the game

[6]: A=np.array([[3,1],[7,0]])
     B=np.array([[3,7],[1,0]])
     rps=nash.Game(A,B)
     rps

[6]: Bi matrix game with payoff matrices:

     Row player:
     [[3 1]
      [7 0]]

     Column player:
     [[3 7]
      [1 0]]

[ ]: alculate the equilibria of the game using support enumeration (see https://nashpy.readthedocs.io/en/stable/text-

[7]: eqs=rps.support_enumeration()

[8]: list(eqs)

[8]: [(array([1., 0.]), array([0., 1.])),
      (array([0., 1.]), array([1., 0.])),
      (array([0.2, 0.8]), array([0.2, 0.8]))]

[13]: #calculate utility of the mixed Nash equilibrium

[14]: sigma_r=[1/5,4/5]
      sigma_c=[1/5,4/5]
      rps[sigma_r,sigma_c]

[14]: array([1.4, 1.4])

[ ]:
```

Social dilemma space

R > T > S > P

T > R > ***S > P***

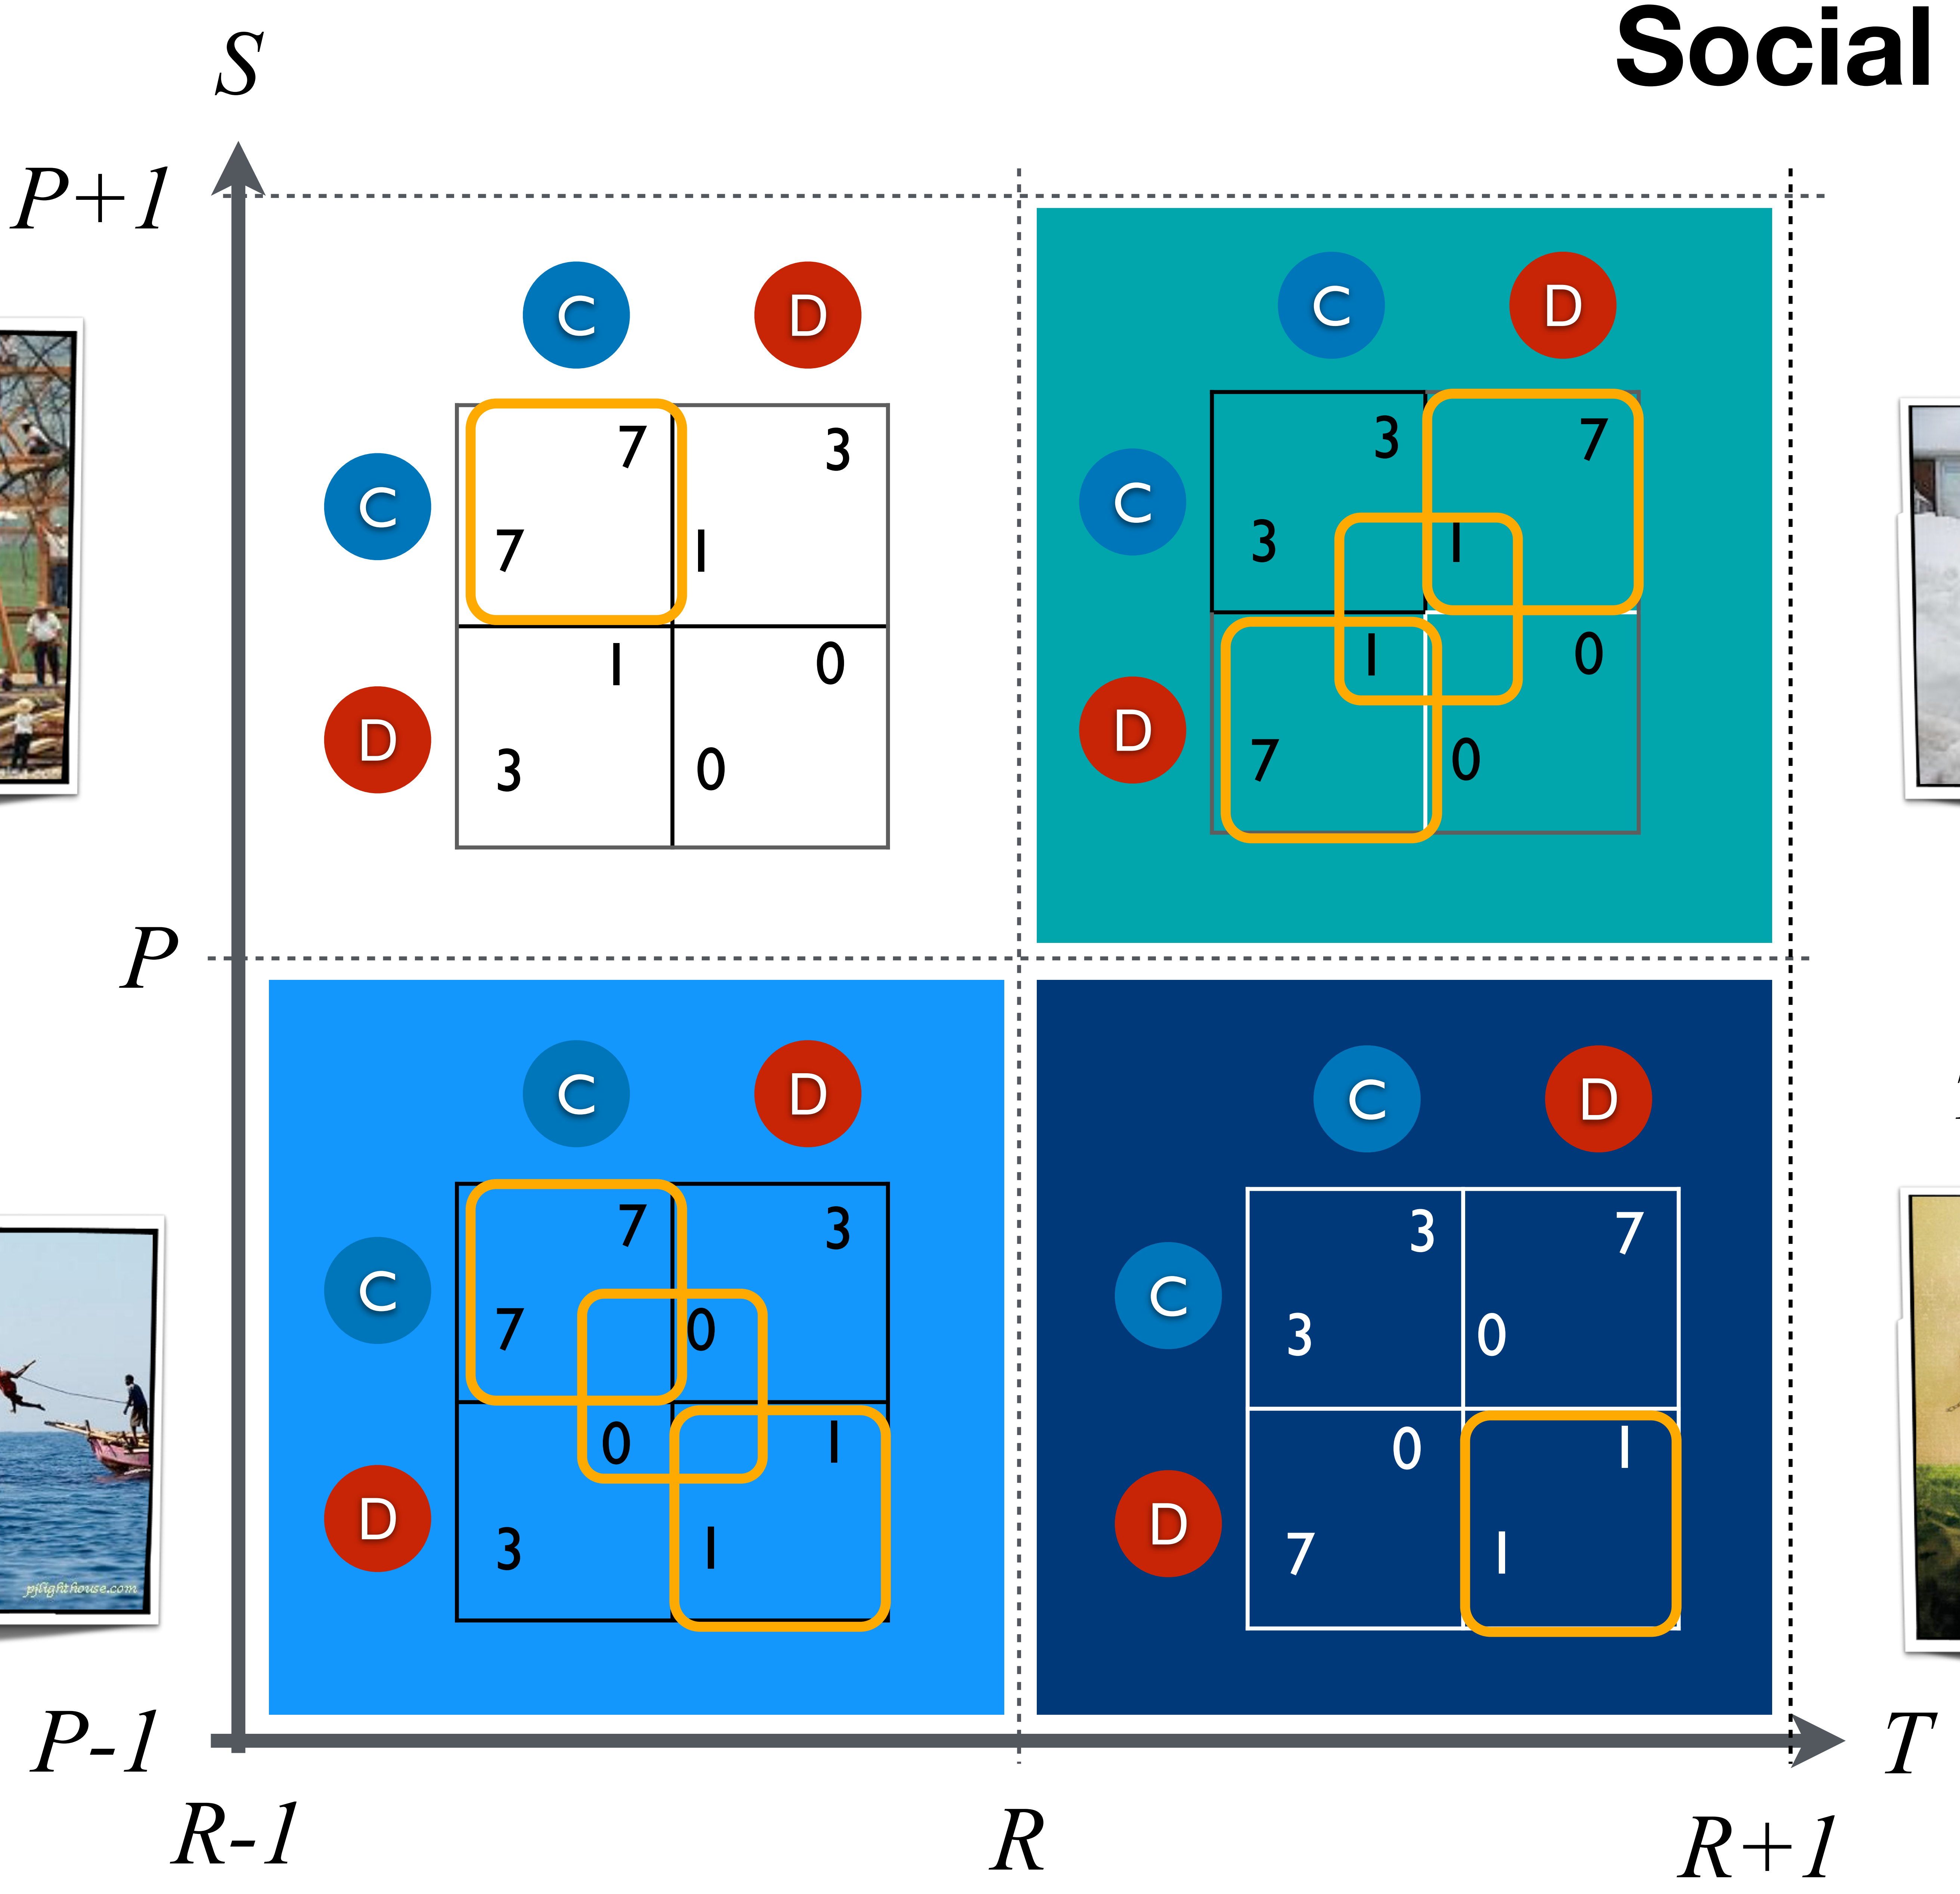

R > T *> P > S*

T > R > P > S

Game Theory and NE assumptions

People are **rational** actors that are **self-interested** and **utility** (payoff) **maximising**

A mixed NE assumes that the actions of both players are independent.

Knowing what action the row player selected does not give you any information about what the column player will do

Why should one expect Nash behaviour from rational players ?

Argument 1 ; May be obtained through introspection

Argument 2 ; If agreed upon, before the game, none of the players wants to deviate (self-enforcing)

Argument 3 ; May be the product of learning or evolution

More solution concepts

Remember …

Solution concepts *are principles according to which one can identify interesting subsets of outcomes of a game [see book Leyton-Brown and Shoham]*

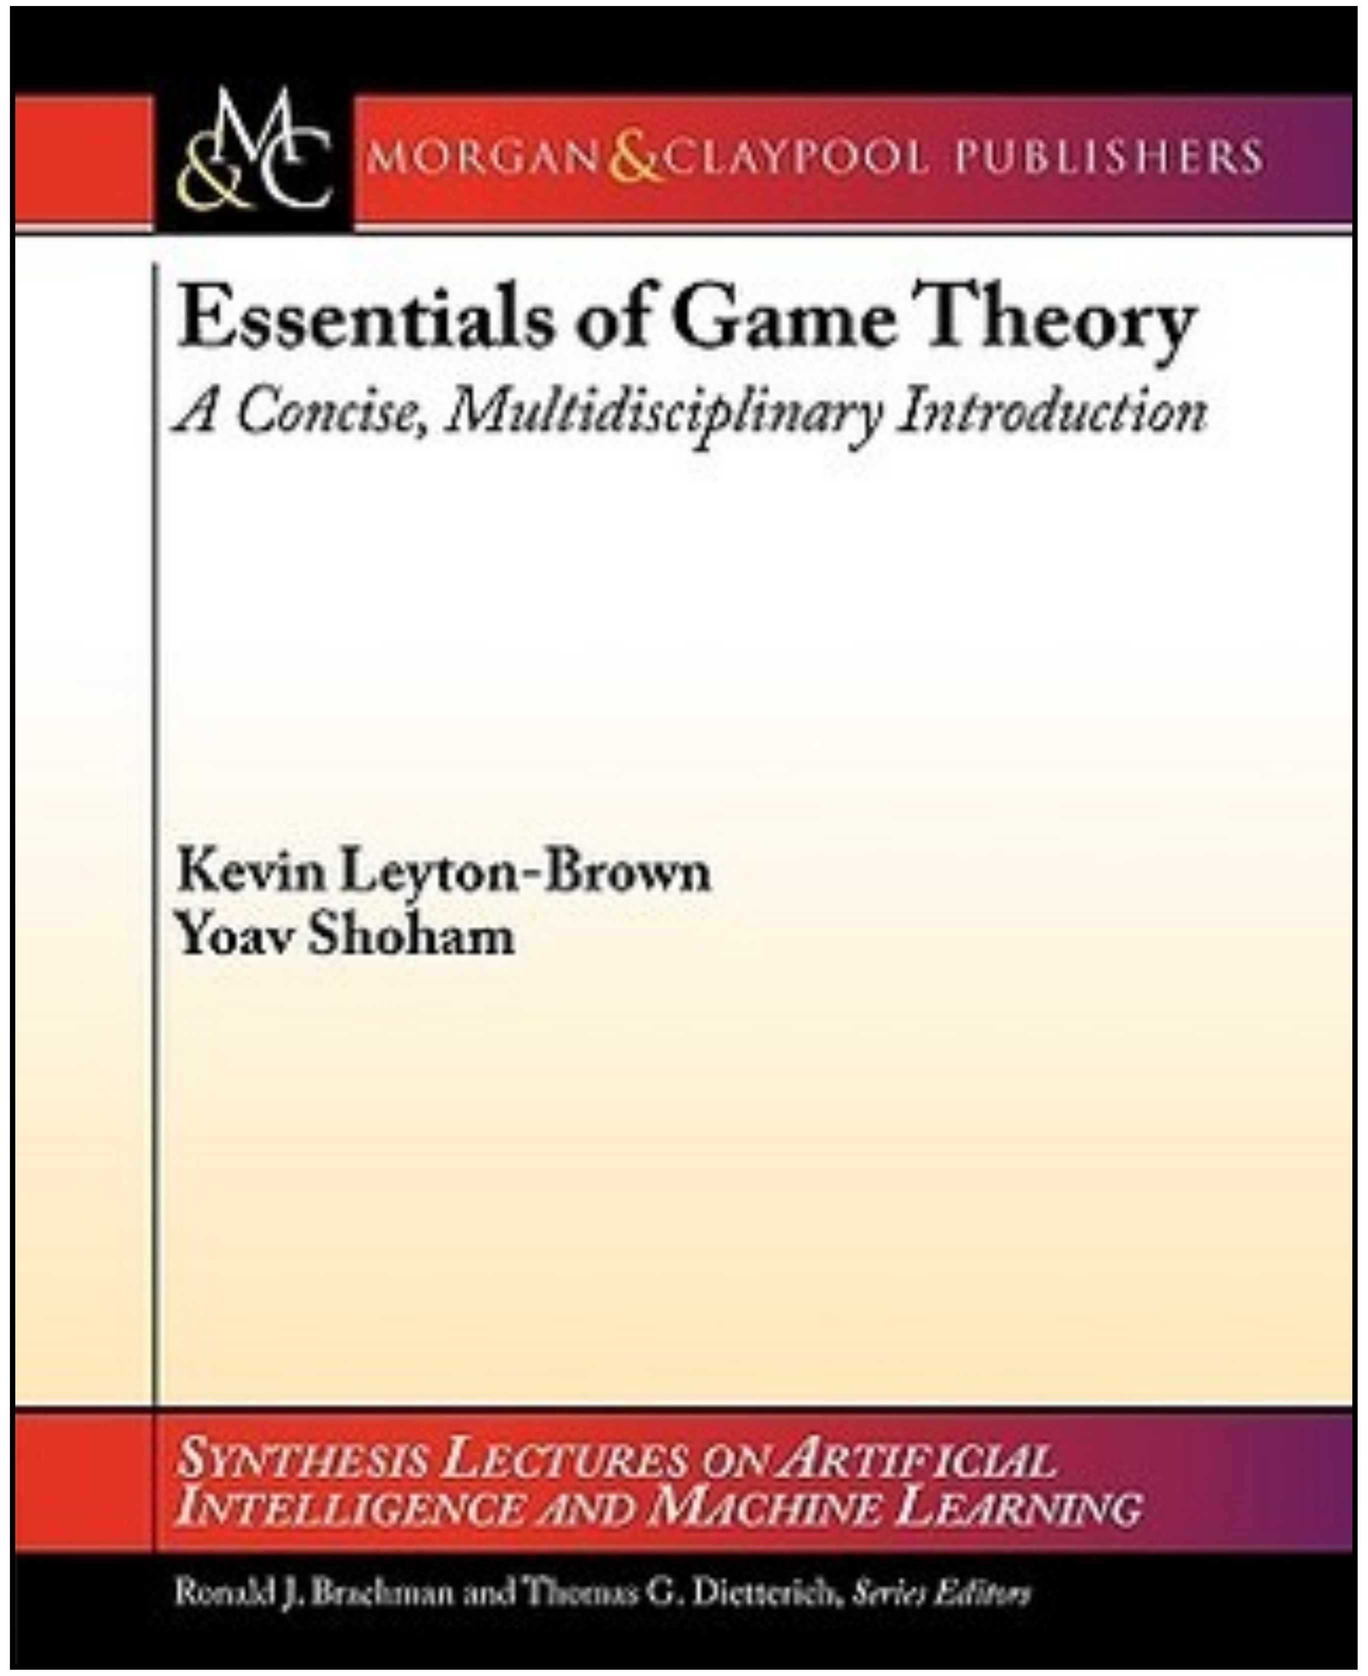

Correlated equilibria

Evolutionary Stable Strategy

https://www.gtessentials.org/toc.html

Other solution concepts; Correlated equilibria

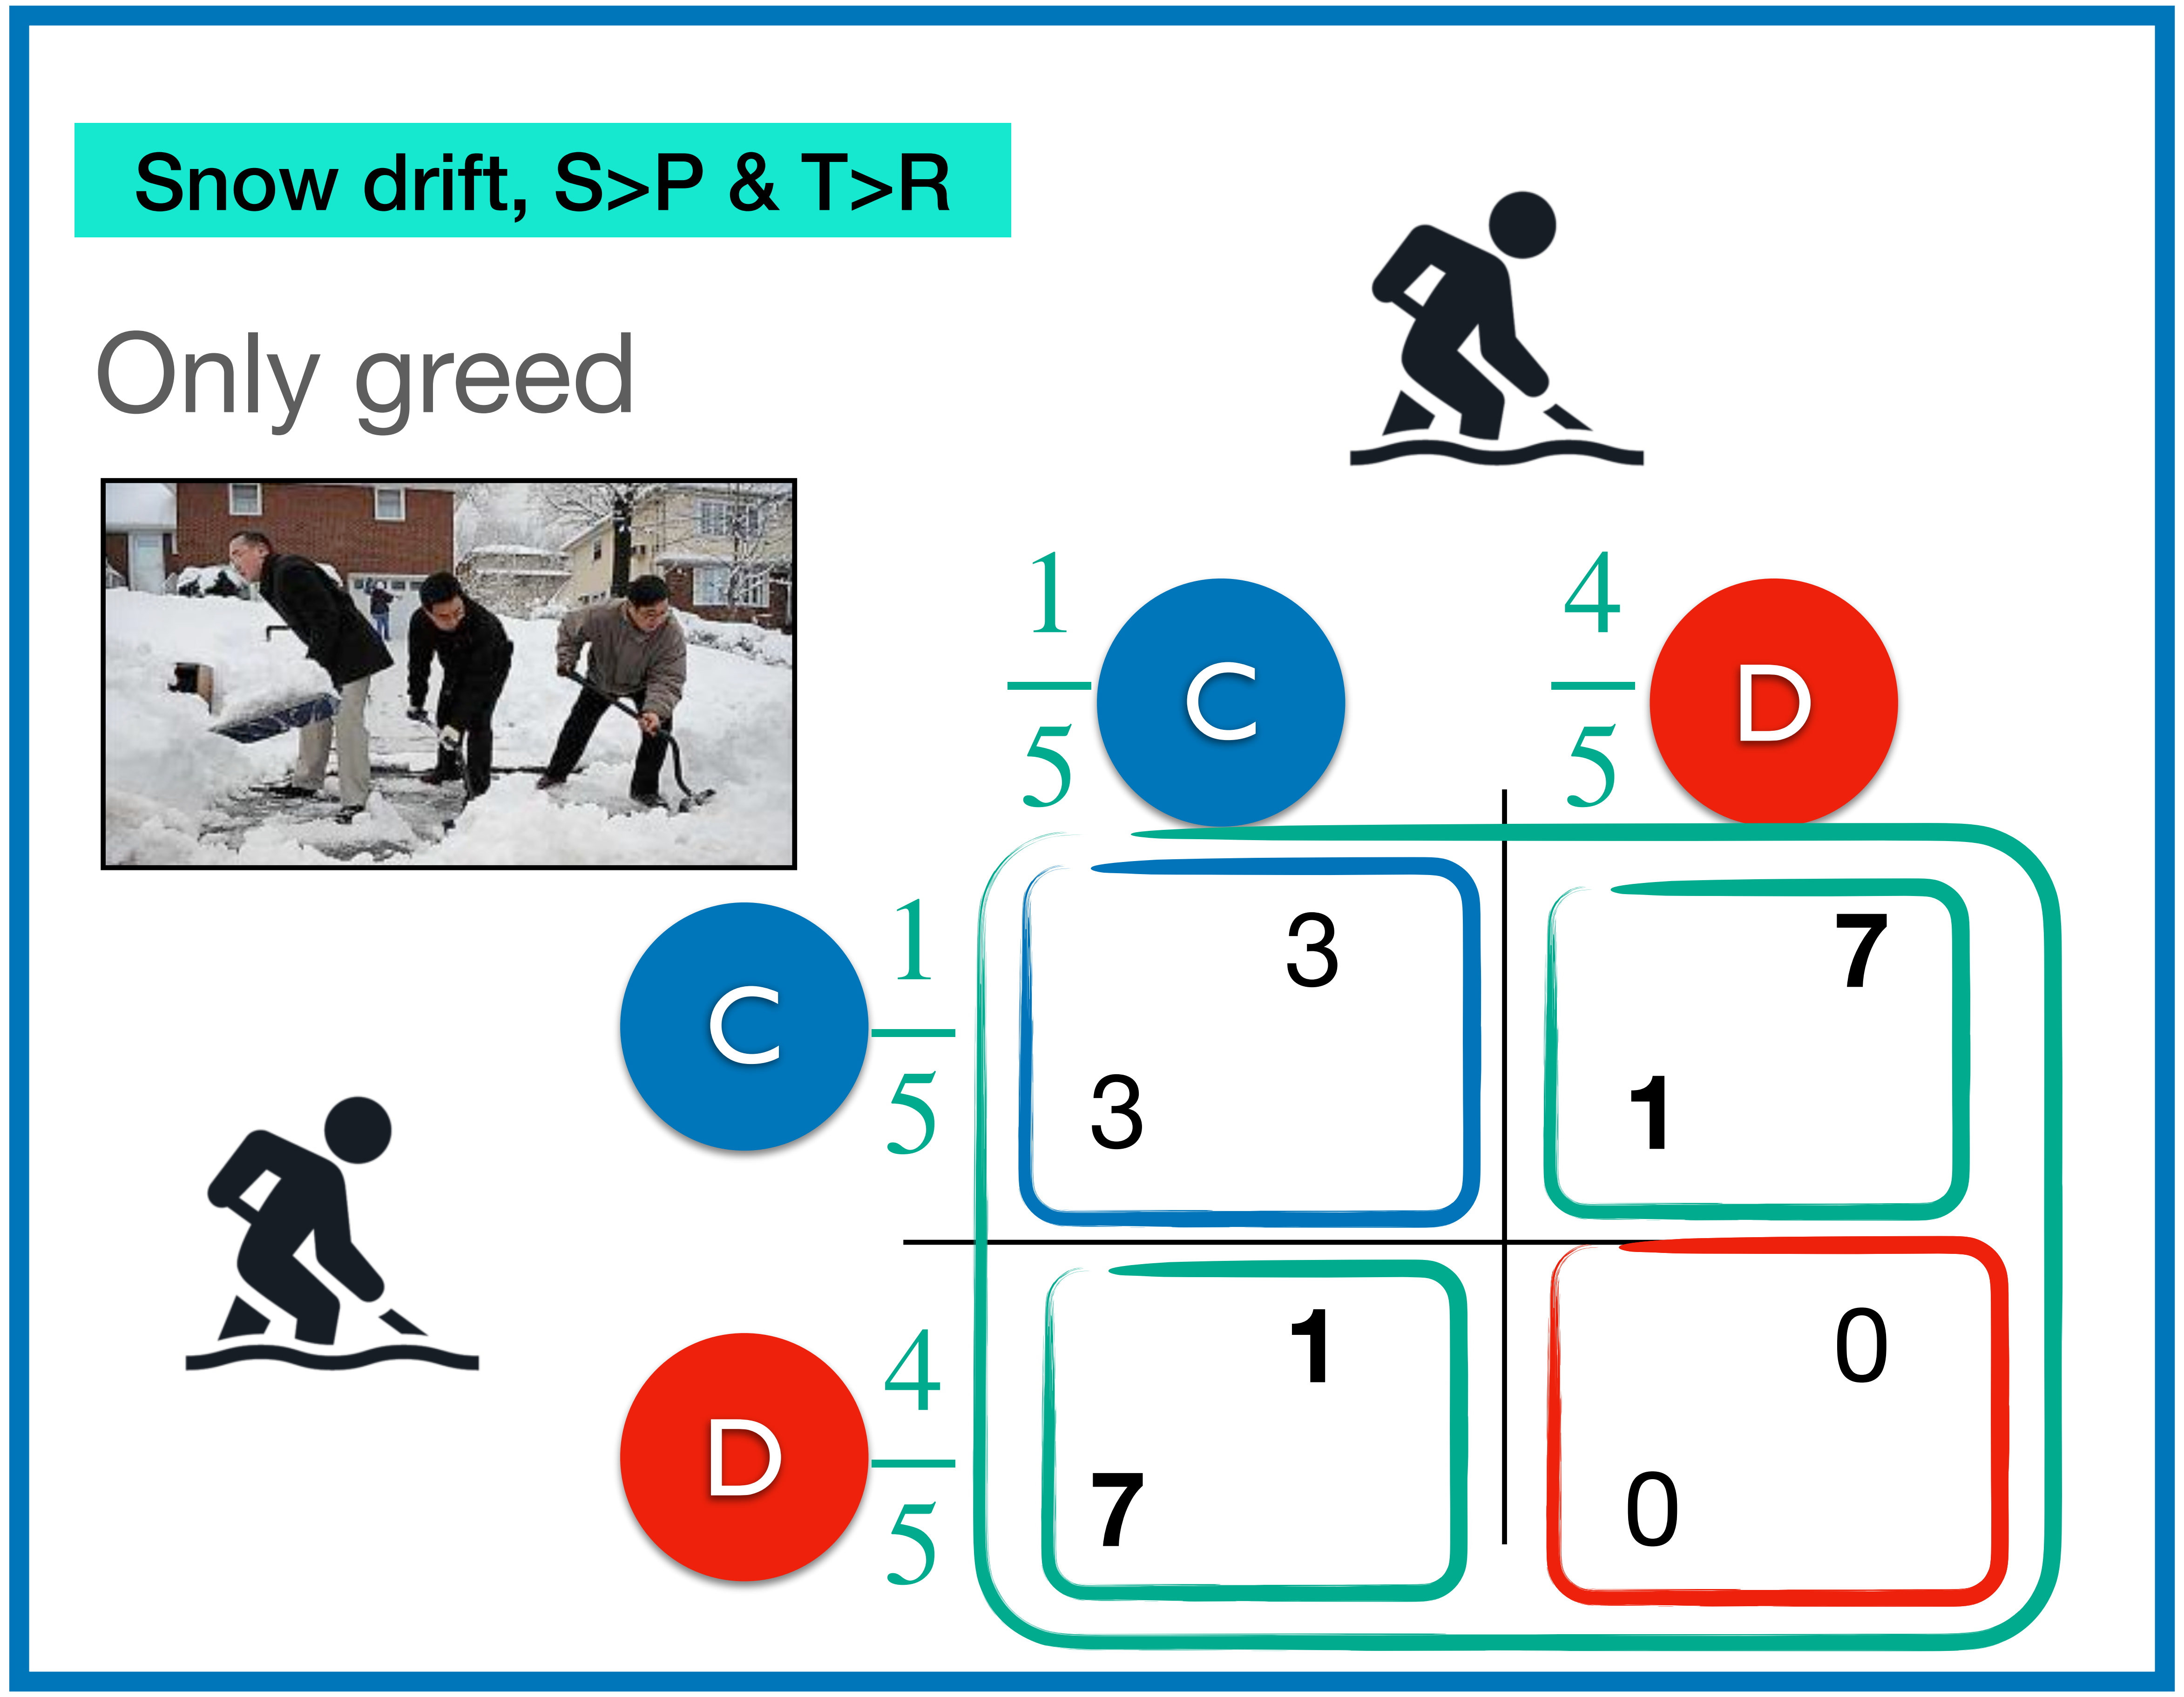

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

A mixed NE includes all possible action combinations

When $e^* = [p = \frac{1}{5}, \frac{4}{5}], [q = \frac{1}{5}, \frac{4}{5}]$ then the outcome (D, D) will occur with probability $(1-p)(1-q) = \frac{16}{25}$ reducing social welfare

What would be better is to **avoid** (D, D)**.**

Both players could **follow a coin toss (fair randomising device) to inform them about what to do**, where heads could signal (C, D) and tail could signal (D, C)

Other solution concepts; Correlated equilibria

Snow drift, S>P & T>R

Only greed

	$\frac{1}{5}$ C	$\frac{4}{5}$ D
$\frac{1}{5}$ C	3, 3	1, 7
$\frac{4}{5}$ D	7, 1	0, 0

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

Benefits of the coin toss?

(1) (D, D) is avoided

(2) **Fairness** in shovelling is achieved (as in (C, C))

(3) **Social welfare** can exceed the mixed NE

Coin toss $(h = \frac{1}{2}, t = 1 - h)$ between $e_1 = (C, D)$ and $e_2 = (D, C)$

They would obtain $E_r(\frac{1}{2}e_1; \frac{1}{2}e_2) = E_c(\frac{1}{2}e_1; \frac{1}{2}e_2) = 3$ which is better than the mixed NE

Rewards can be made better by correlation

Other solution concepts; Correlated equilibria

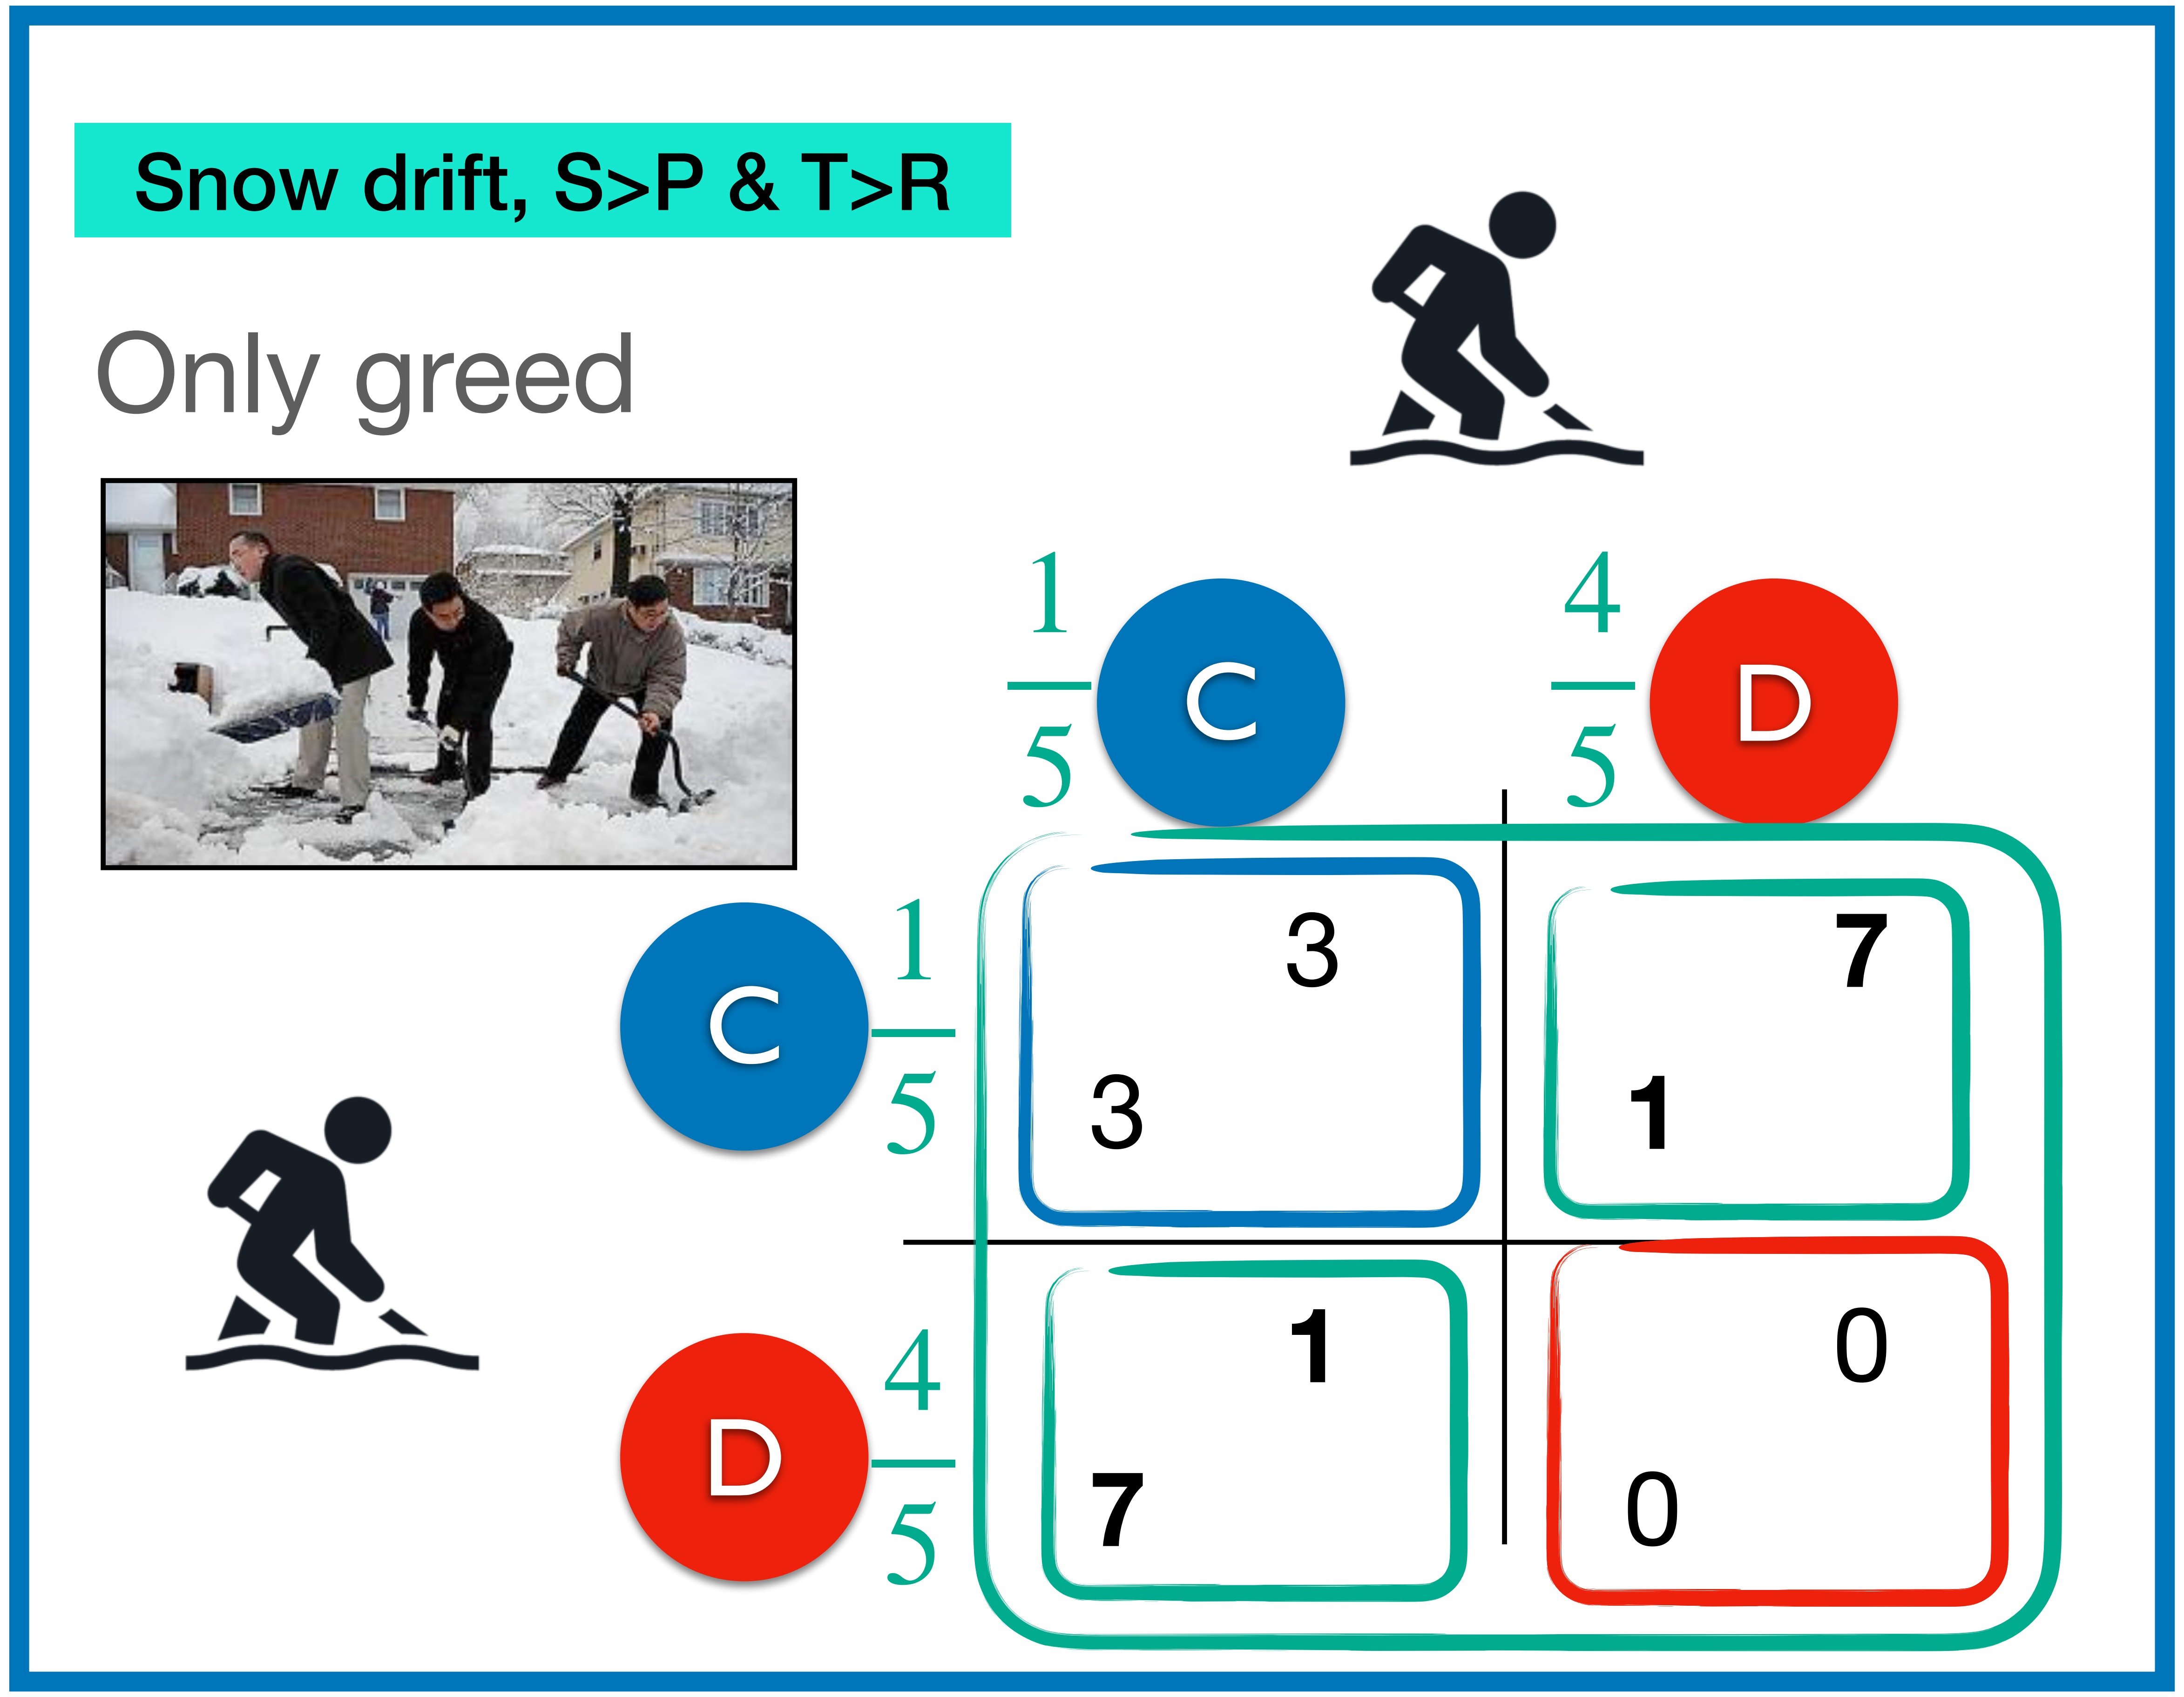

C.H. Coombs (1973) A reparameterization of the prisoner's dilemma game. Behavioral Science 18:424-428

Correlated equilibria*

A randomised assignment of (potentially correlated) action **recommendations** to the agents, such that nobody wants to deviate **

*"The idea is that each player chooses their action according to their private observation of the value of the same public signal. A strategy assigns an action to every possible observation a player can make. If no player would want to deviate from their strategy (assuming the others also don't deviate), **the distribution from which the signals are drawn is called a correlated equilibrium.**" [Wikipedia, May 2024]*

Any mixed NE is also a correlated equilibrium

*Aumann, R. J. (1987). Correlated equilibrium as an expression of Bayesian rationality. Econometrica: Journal of the Econometric Society, 1-18.

**https://www.youtube.com/watch?v=sQOrIpARr5E

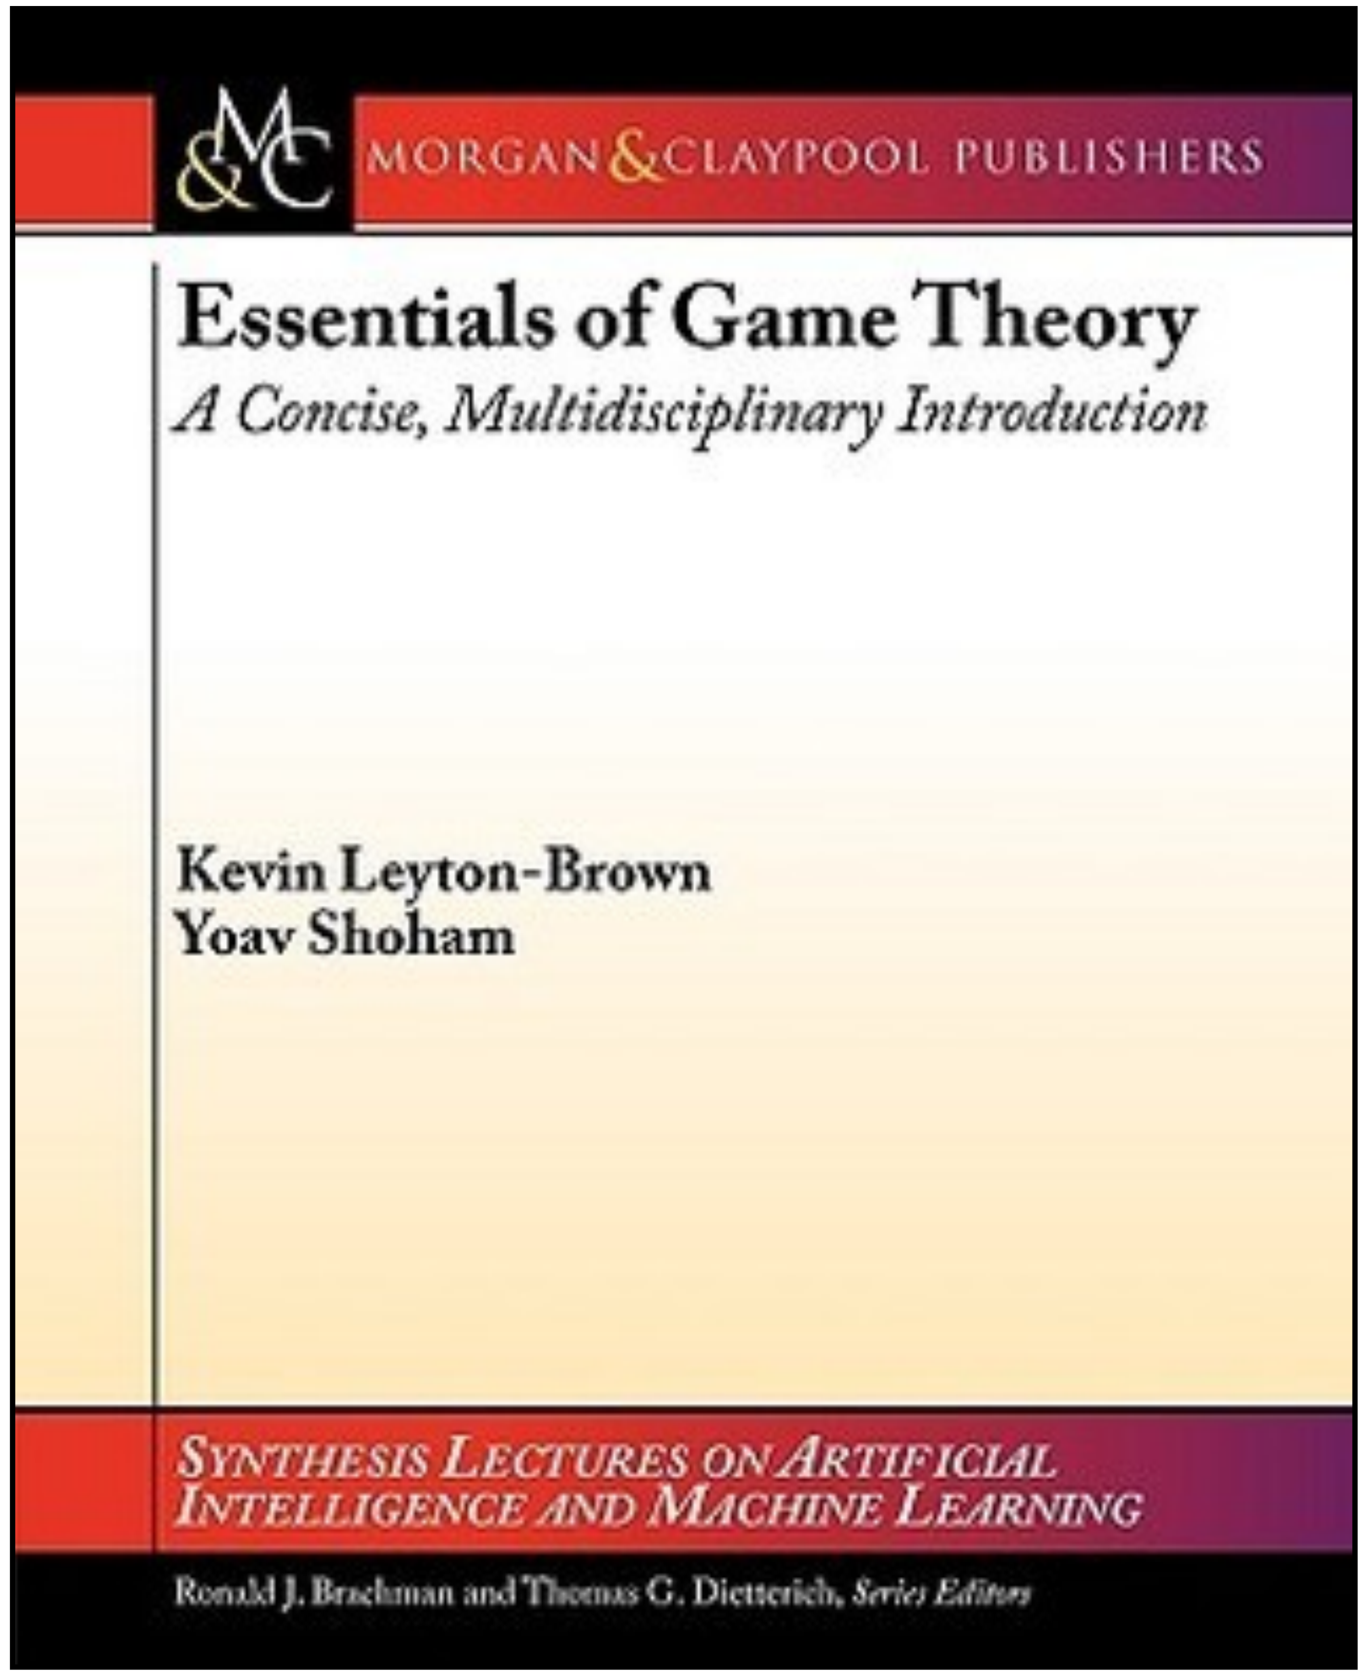

https://www.gtessentials.org/toc.html

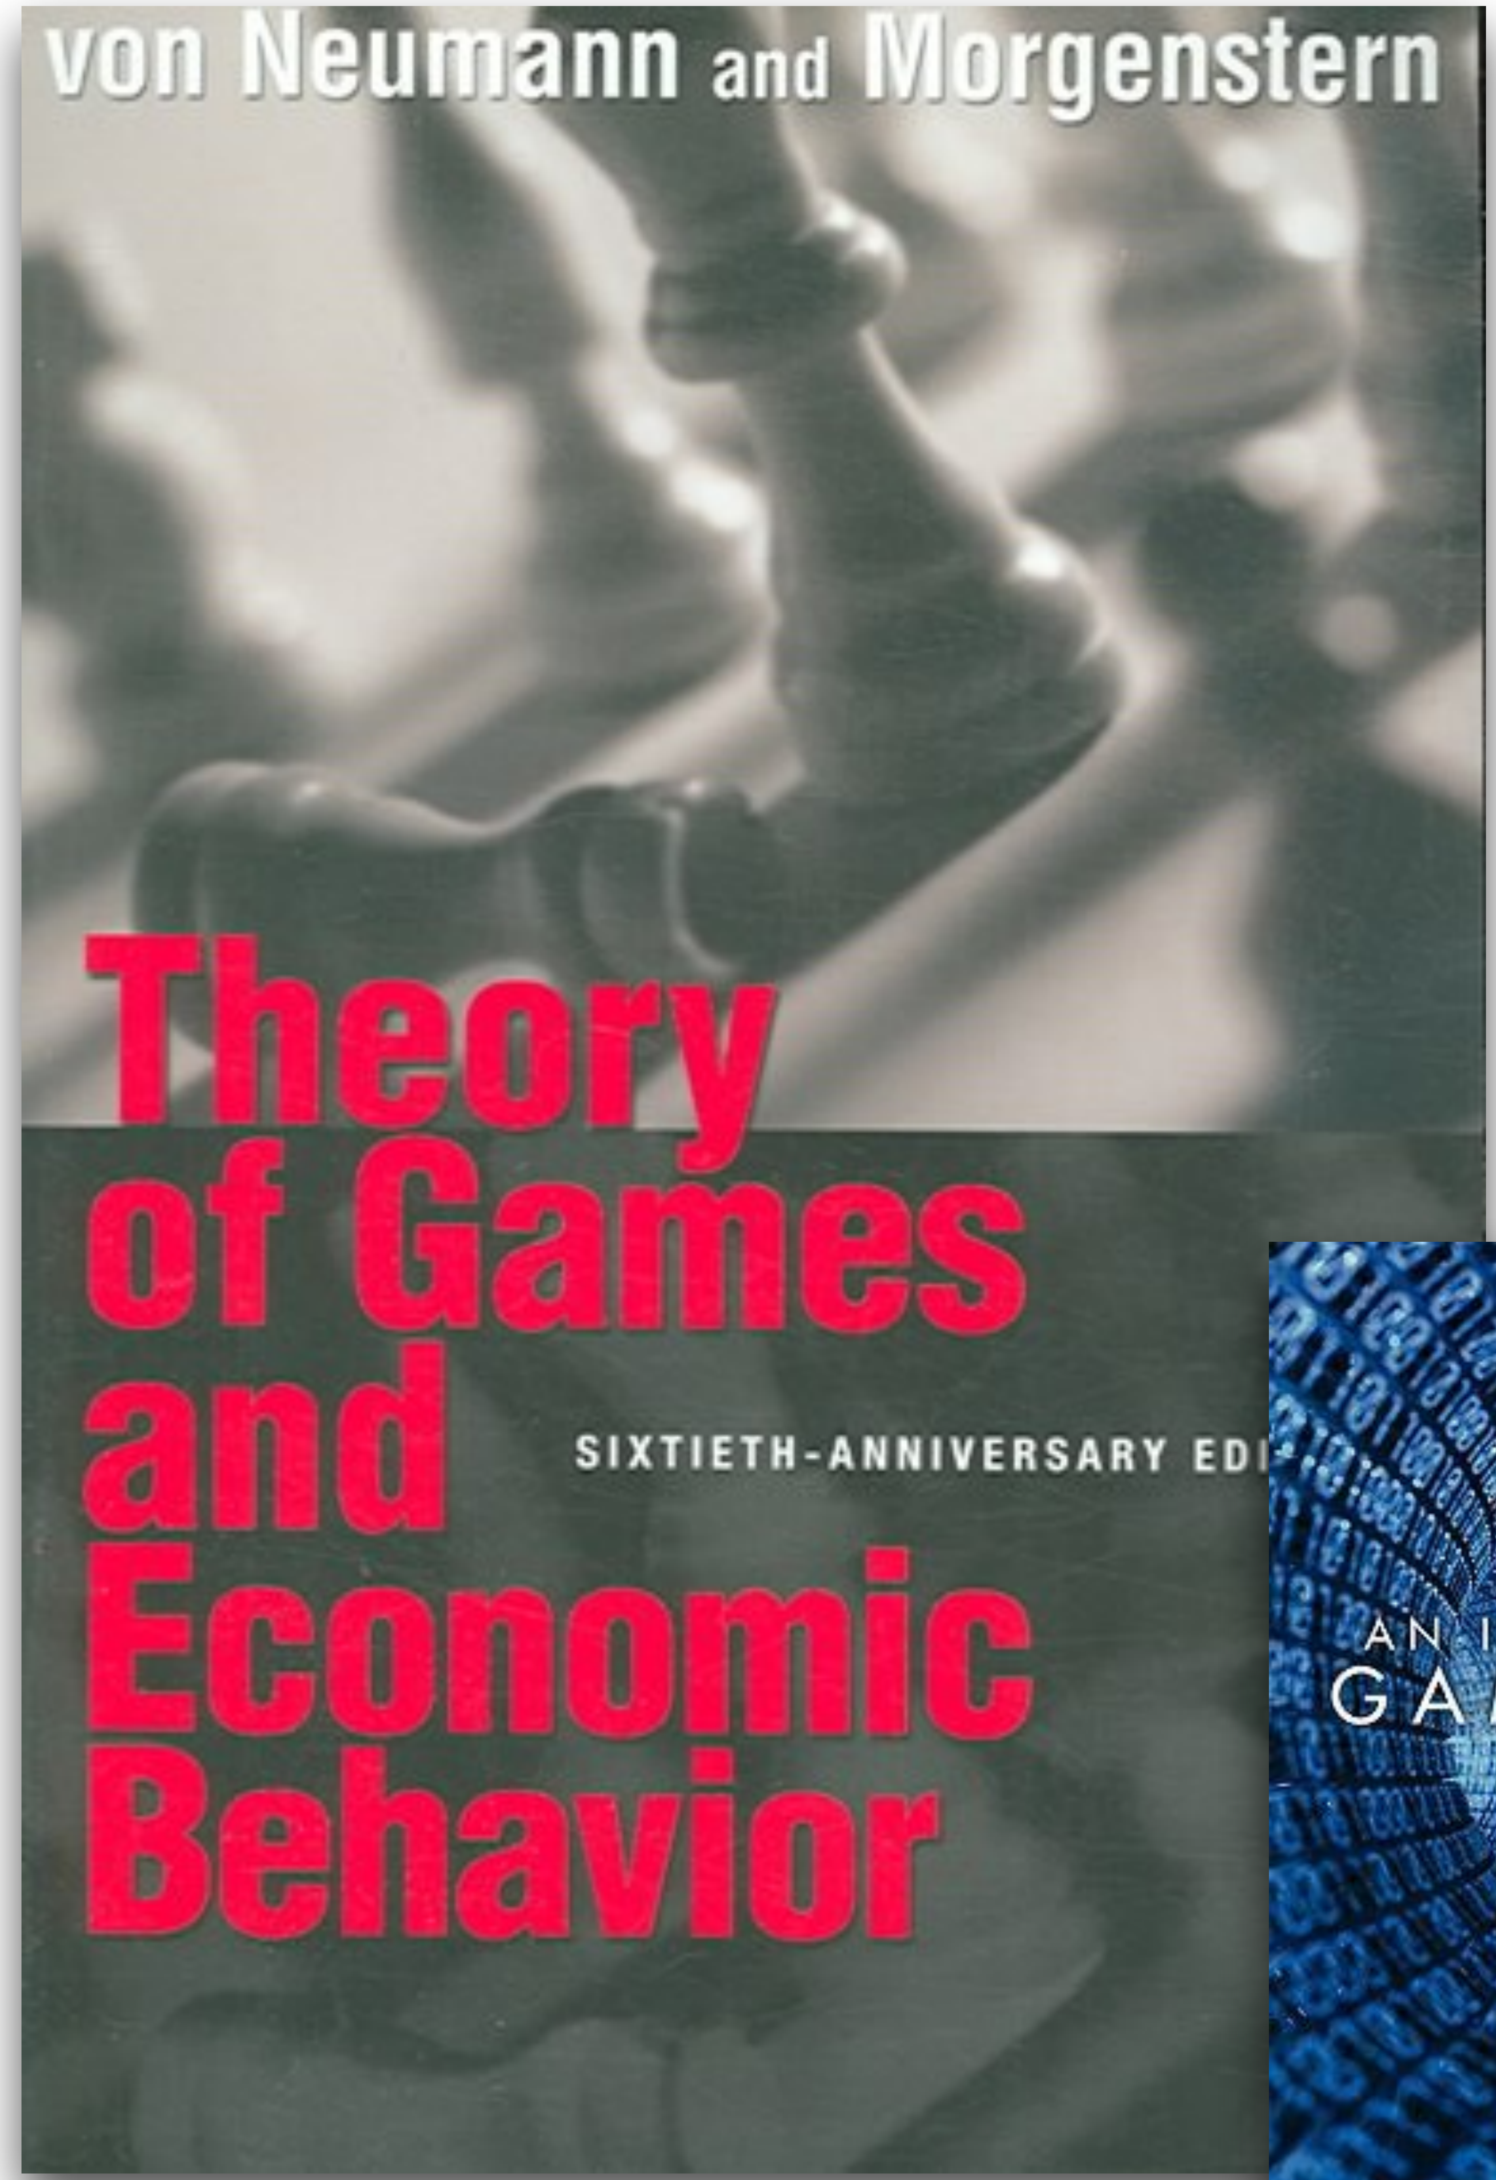

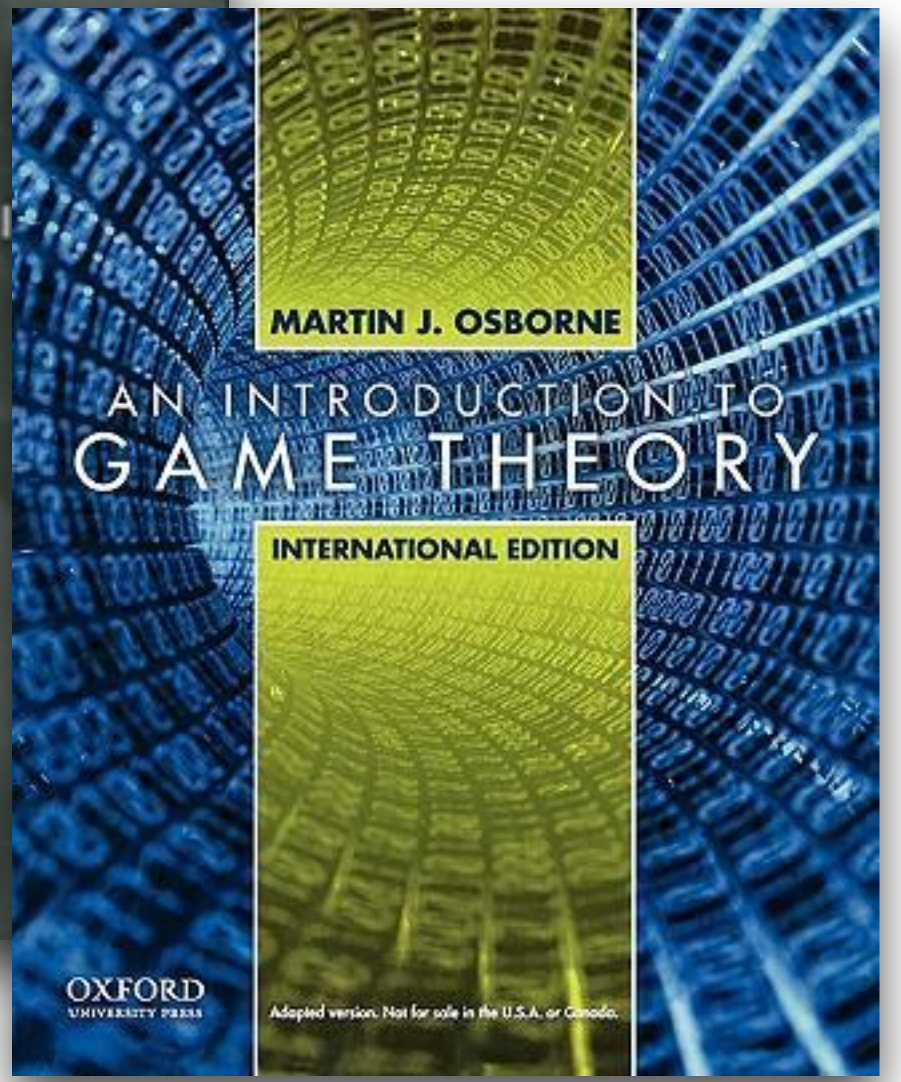

Day 2: Evolutionary Game Theory

https://blogs.bl.uk/untoldlives/2020/03/john-maynard-smith-evolutionary-biology-and-the-logic-of-animal-conflict.html

(Reprinted from Nature, Vol. 246, No. 5427, pp. 15–18, November 2, 1973)

The Logic of Animal Conflict

J. MAYNARD SMITH
School of Biological Sciences, University of Sussex, Falmer, Sussex BN1 9QG
G. R. PRICE
Galton Laboratory, University College London, 4 Stephenson Way, London NW1 2HE

Conflicts between animals of the same species usually are of "limited war" type, not causing serious injury. This is often explained as due to group or species selection for behaviour benefiting the species rather than individuals. Game theory and computer simulation analyses show, however, that a "limited war" strategy benefits individual animals as well as the species.

In a typical combat between two male animals of the same species, the winner gains mates, dominance rights, desirable territory, or other advantages that will tend toward transmitting its genes to future generations at higher frequencies than the loser's genes. Consequently, one might expect that natural selection would develo... effective weapons...

and ask what strategy will be favoured under individual selection. We first consider conflict in species possessing offensive weapons capable of inflicting serious injury on other members of the species. Then we consider conflict in species where serious injury is impossible, so that victory goes to the contestant who fights longest. For each model, we seek a strategy that will be stable under natural selection; that is, we seek an "evolutionarily stable strategy" or ESS. The concept of an ESS is fundamental to our argument; it has been derived in part from the theory of games, and in part from the work of MacArthur[13] and of Hamilton[14] on the evolution of the sex ratio. Roughly, an ESS is a strategy such that, if most of the members of a population adopt it, there is no "mutant" strategy that would give higher reproductive fitness.

A Computer Model

A main reason for using computer simulation was to test whether it is possible even ... selection ...

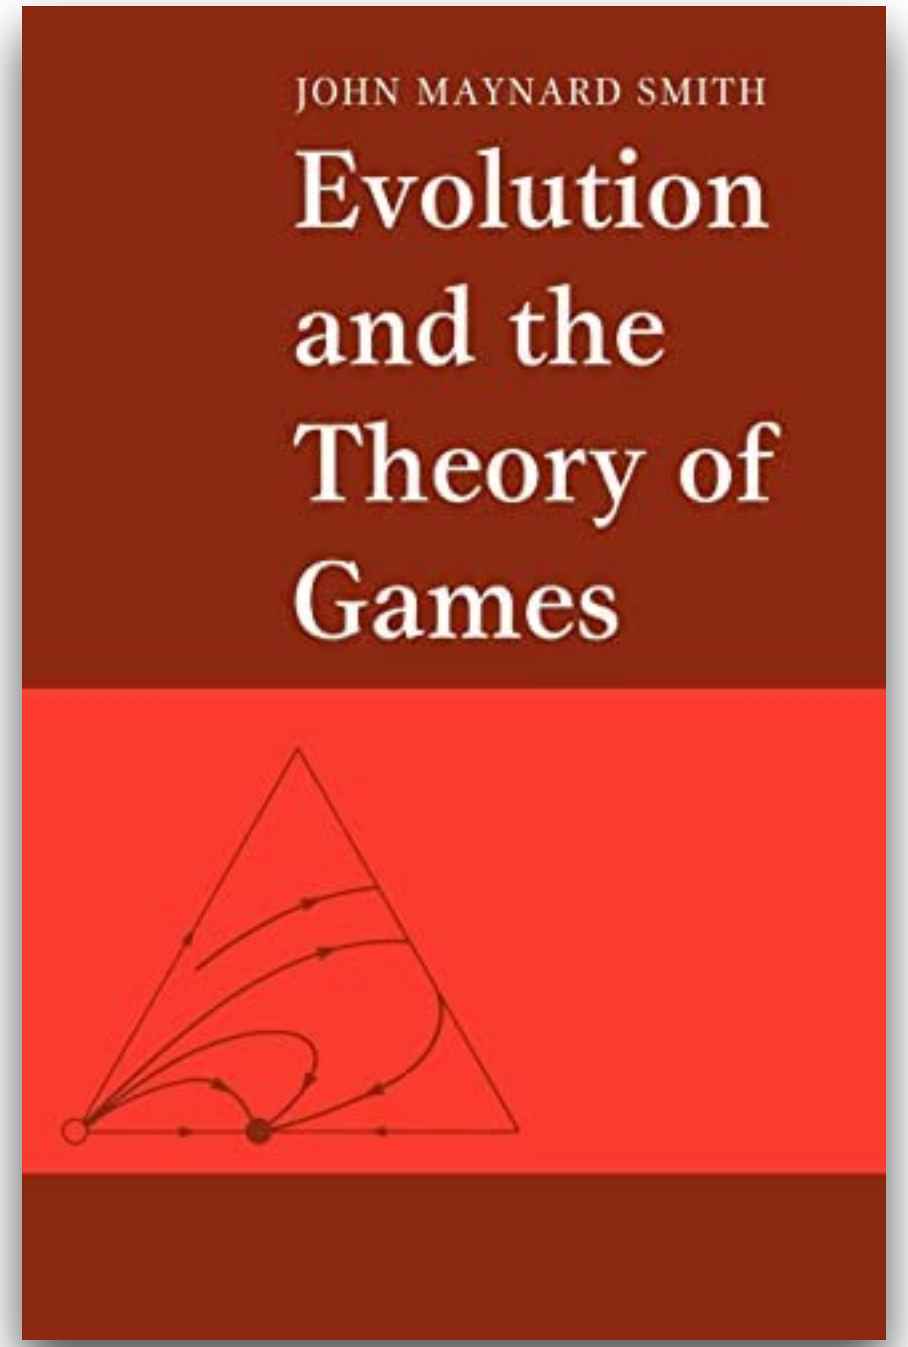

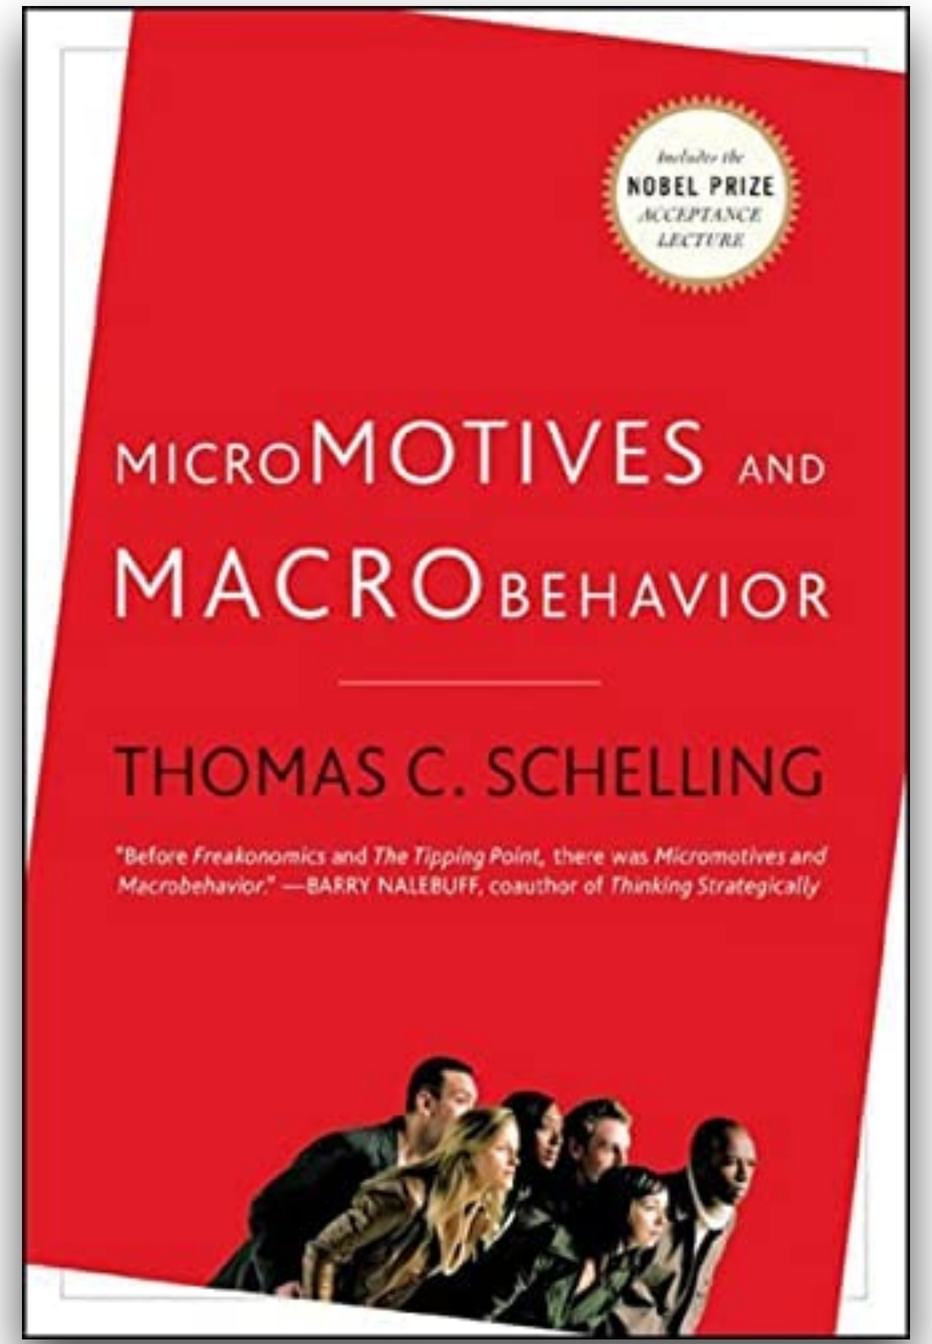

Part 3: Projects

Reproduce a paper

- Pacheco, J. M., Santos, F. C., Souza, M. O., & Skyrms, B. (2009). Evolutionary dynamics of collective action in N-person stag hunt dilemmas. *Proceedings of the Royal Society B: Biological Sciences*, *276*(1655), 315-321.
- Santos, F. C., & Pacheco, J. M. (2011). Risk of collective failure provides an escape from the tragedy of the commons. *Proceedings of the National Academy of Sciences*, *108*(26), 10421-10425.
- Vasconcelos, V. V., Santos, F. C., Pacheco, J. M., & Levin, S. A. (2014). Climate policies under wealth inequality. *Proceedings of the National Academy of Sciences*, *111*(6), 2212-2216.
- Hilbe, C., Šimsa, Š., Chatterjee, K., & Nowak, M. A. (2018). Evolution of cooperation in stochastic games. *Nature*, *559*(7713), 246-249.
- Weitz, J. S., Eksin, C., Paarporn, K., Brown, S. P., & Ratcliff, W. C. (2016). An oscillating tragedy of the commons in replicator dynamics with game-environment feedback. *Proceedings of the National Academy of Sciences*, *113*(47), E7518-E7525.
- Santos, F. C., Pacheco, J. M., & Lenaerts, T. (2006). Evolutionary dynamics of social dilemmas in structured heterogeneous populations. *Proceedings of the National Academy of Sciences*, *103*(9), 3490-3494.

Interesting but more difficult & time extensive:

- Pacheco, J. M., Santos, F. C., & Chalub, F. A. C. (2006). Stern-judging: A simple, successful norm which promotes cooperation under indirect reciprocity. *PLoS computational biology*, *2*(12), e178.
- van den Berg, P., & Wenseleers, T. (2018). Uncertainty about social interactions leads to the evolution of social heuristics. *Nature Communications*, *9*(1), 2151.

Propose your own project

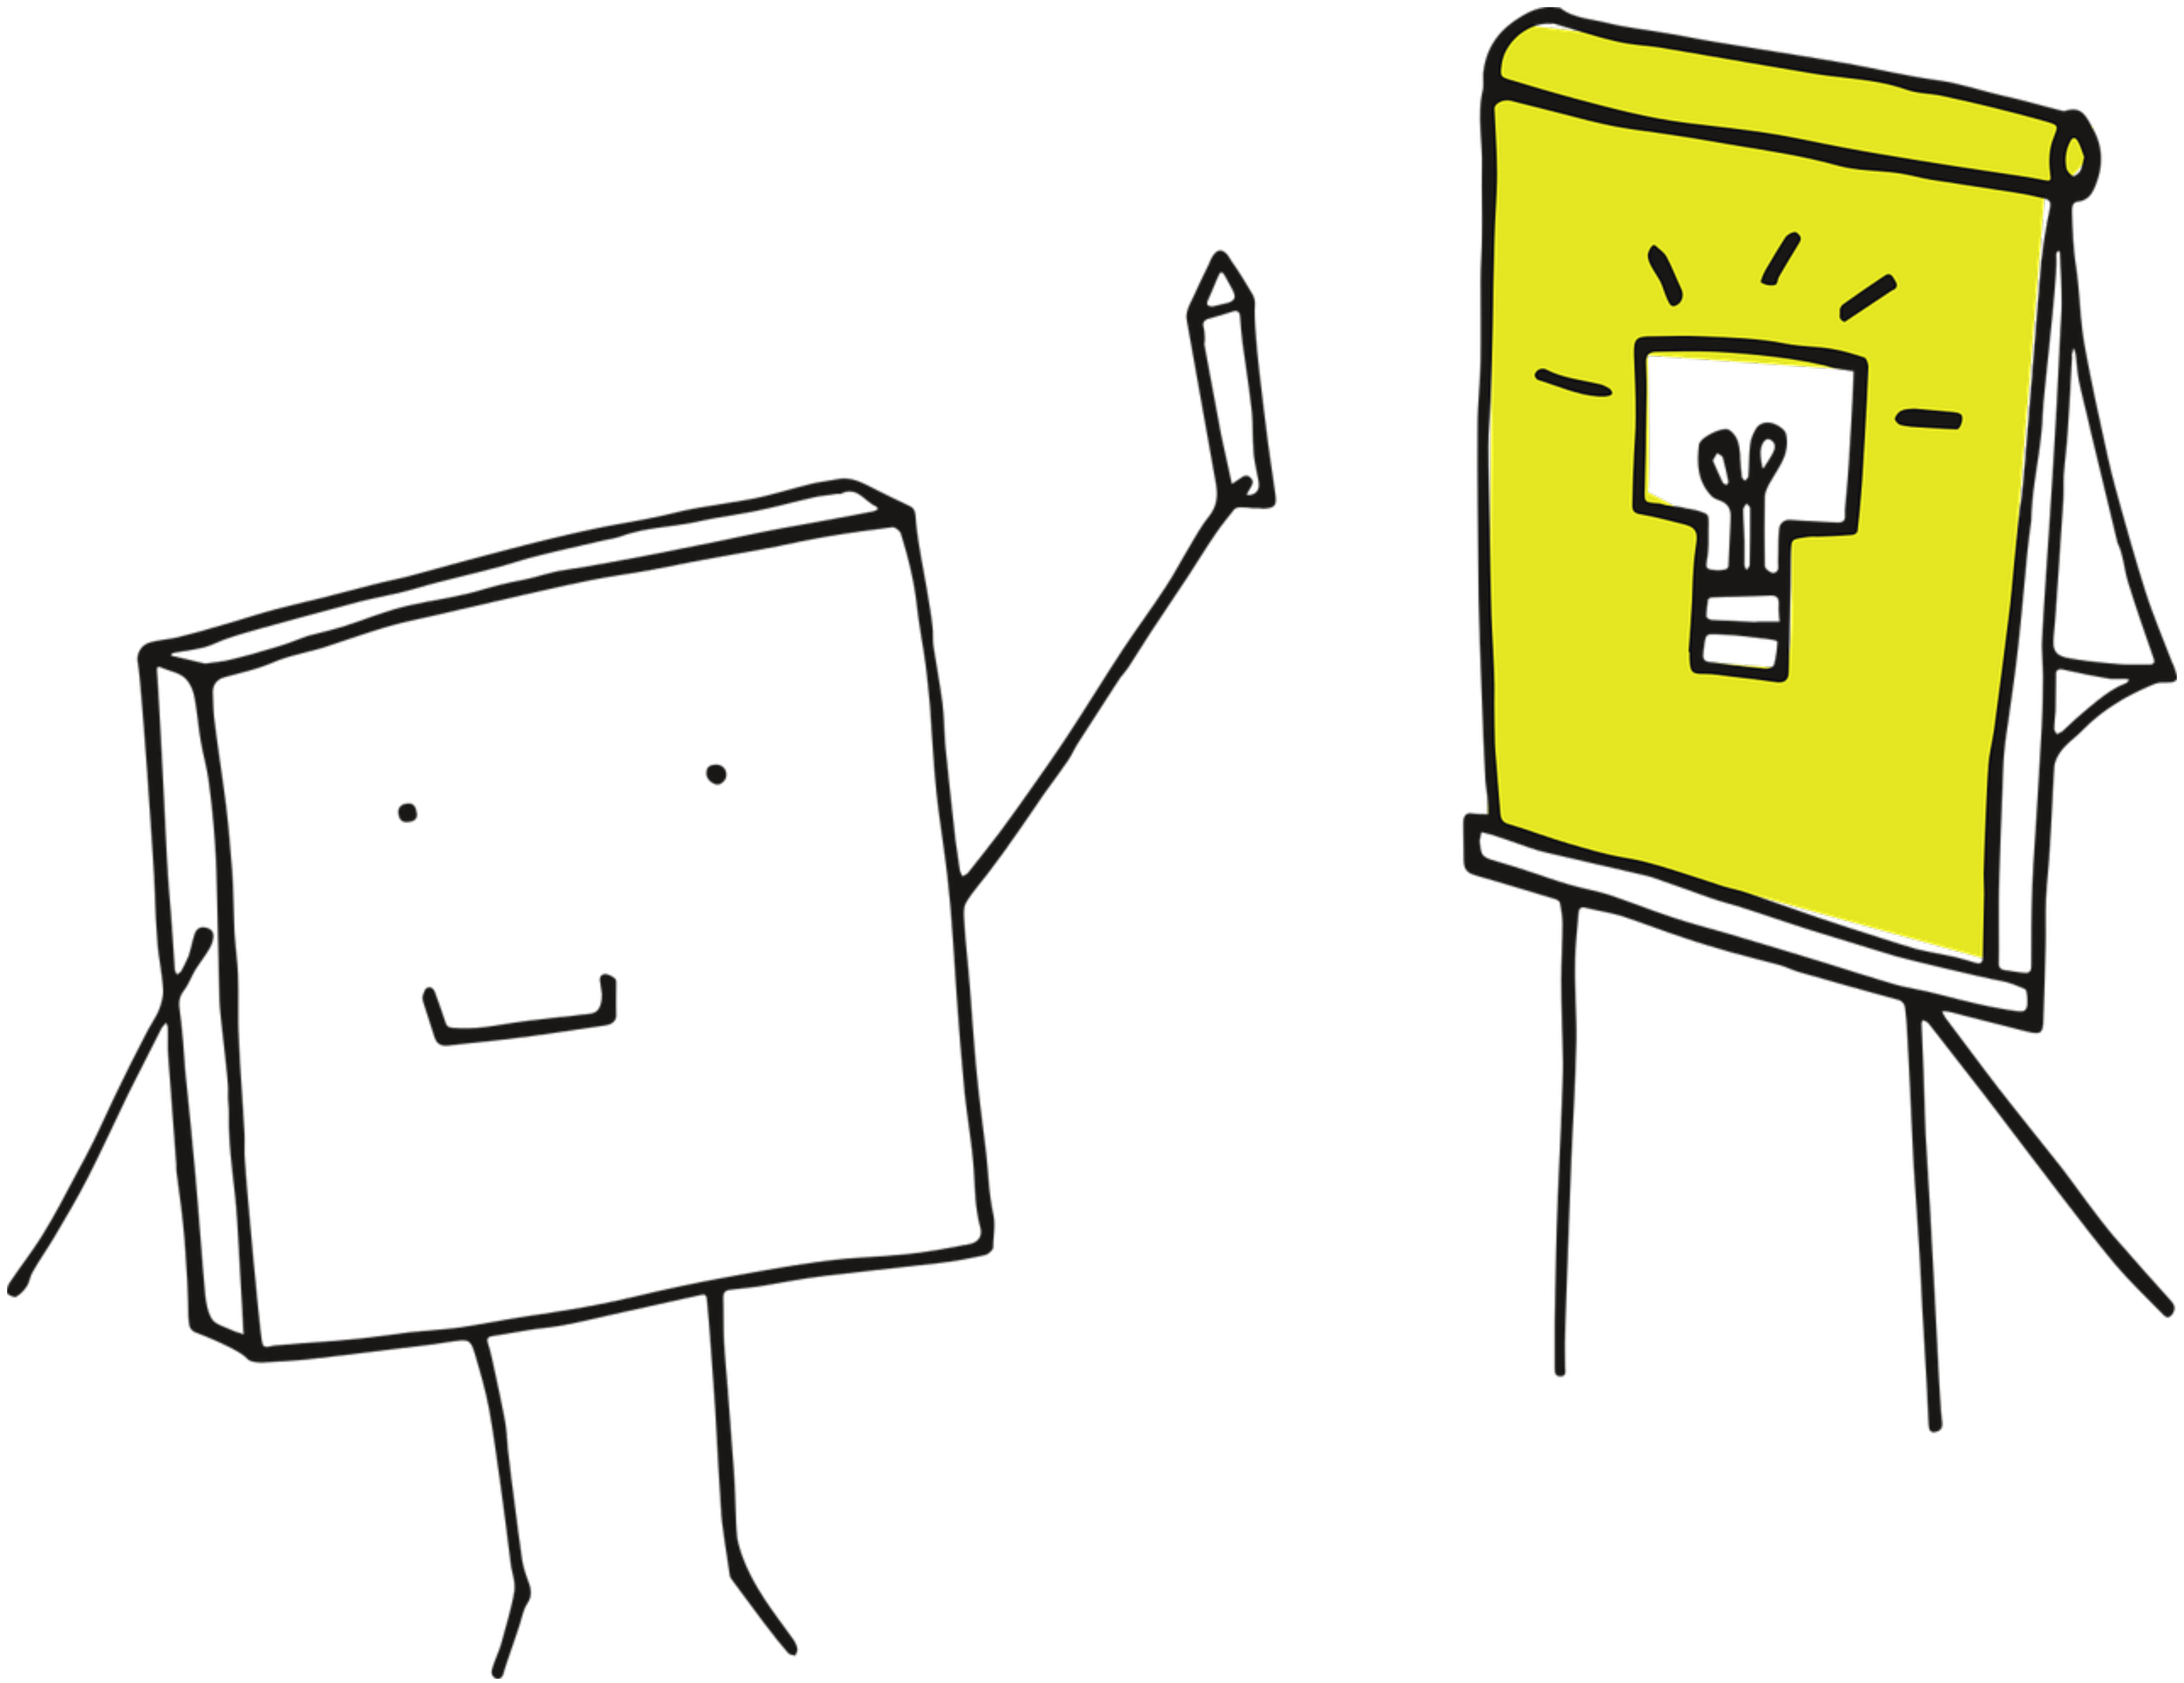

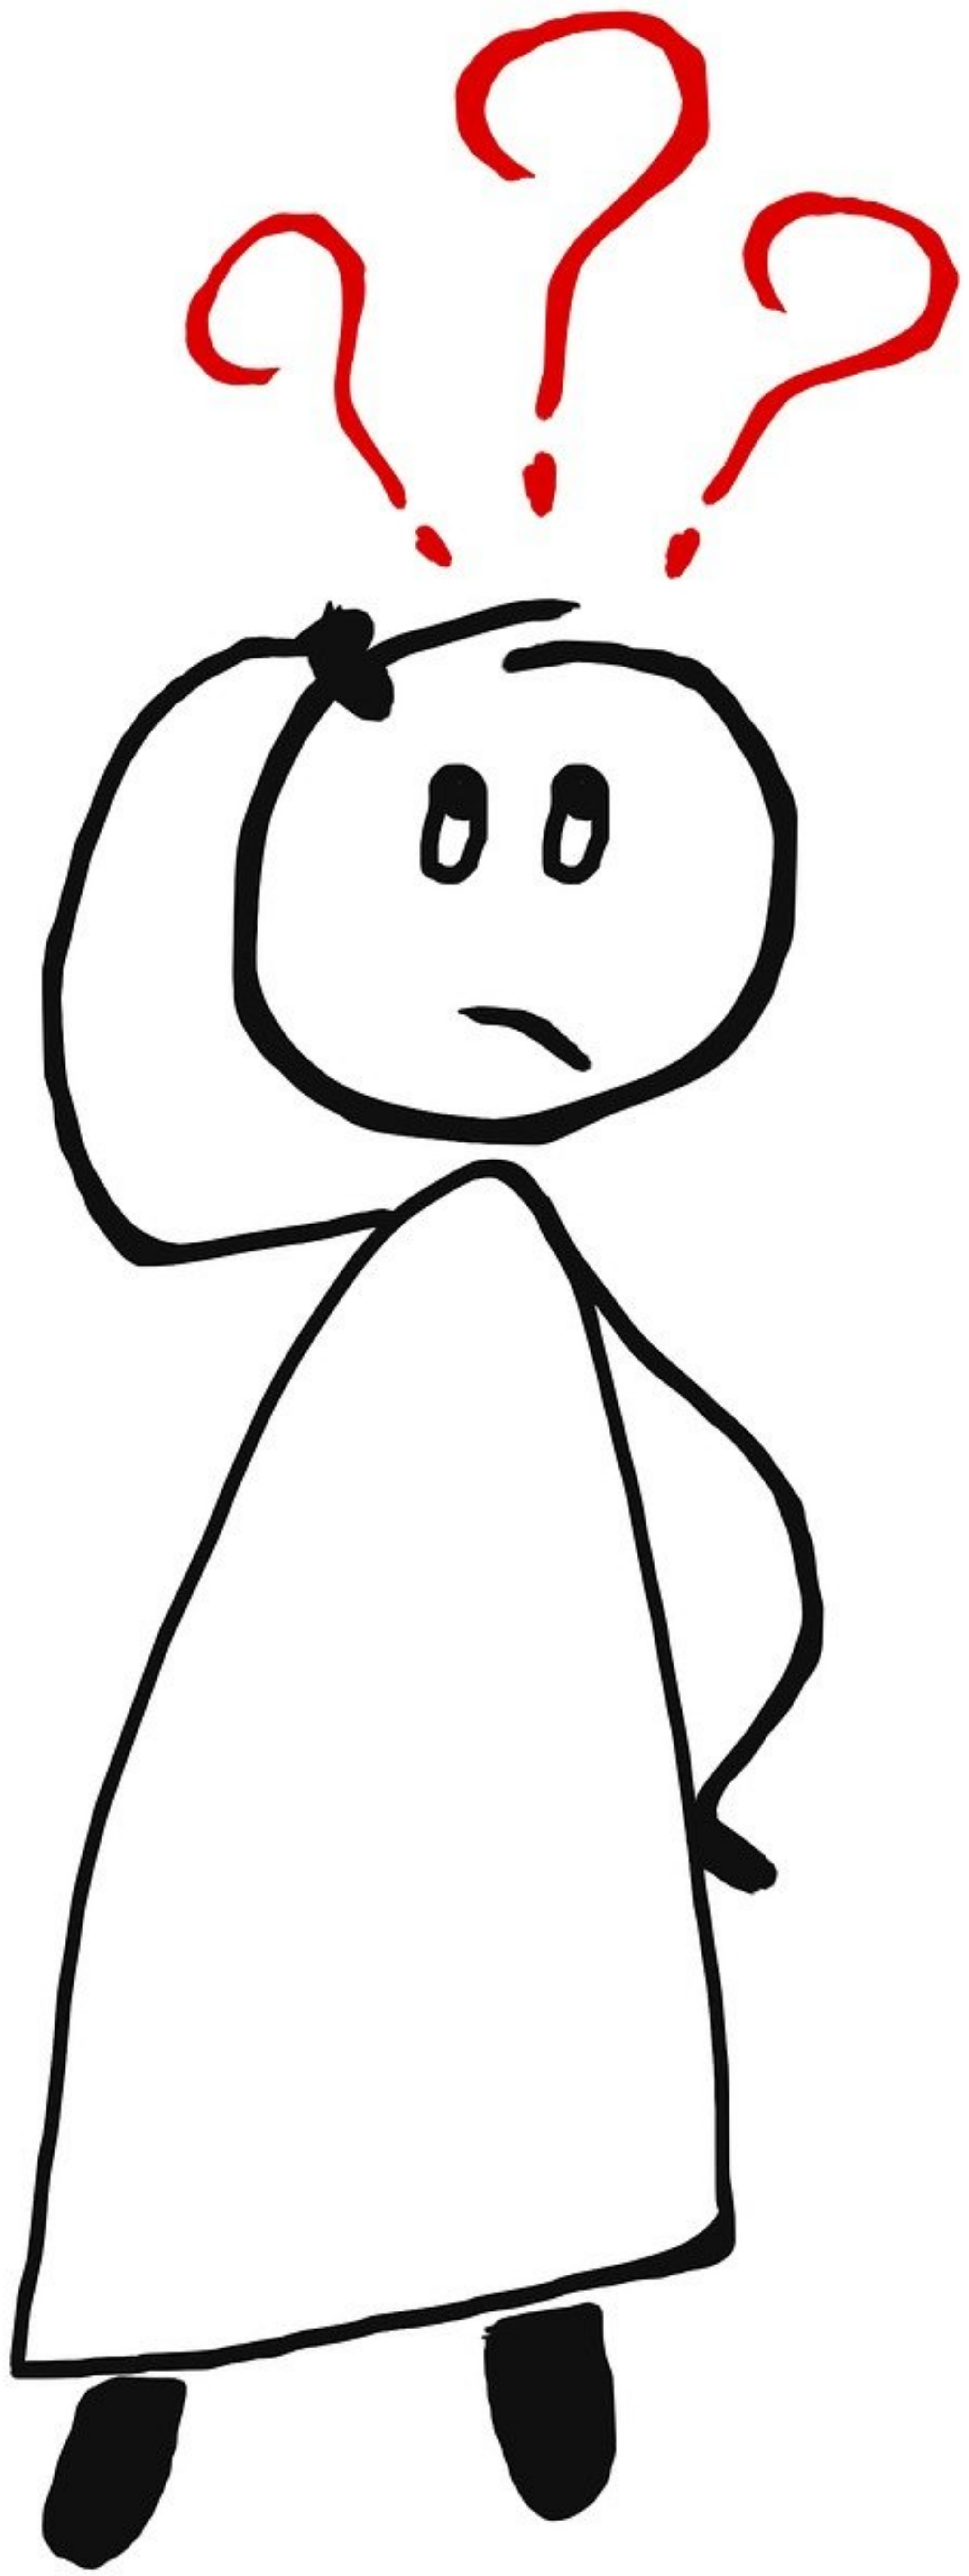

ULB UNIVERSITÉ LIBRE DE BRUXELLES

VUB VRIJE UNIVERSITEIT BRUSSEL

fnrs LA LIBERTÉ DE CHERCHER

Questions ?

elias.Fernández.Domingos@ulb.be
@esocrats
https://github.com/Socrats