An overview

Two days ago's class (March 11, 2025)

- An introduction to evolutionary game theory
  (Replicator dynamics, games in finite populations)

Yesterday's class (March 12, 2025)

- Evolution of cooperation & direct reciprocity
- Social norms & indirect reciprocity

Today's class (March 13, 2025)

- Some current research: Reciprocity in complex environments

The role of memory; the effect of changing environments; the impact of inequality

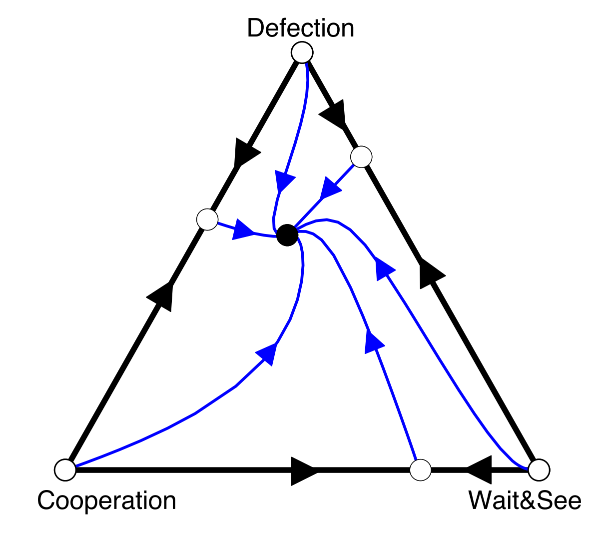

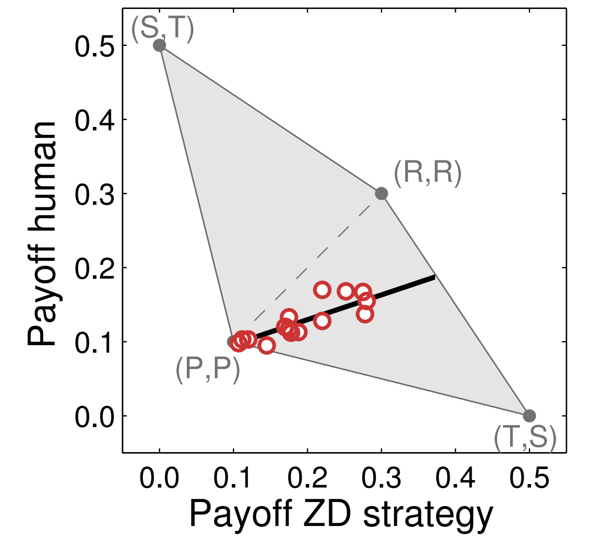

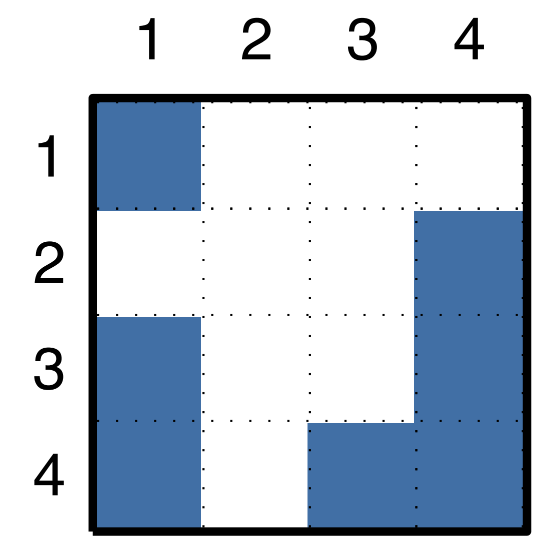

The impact of memory: Results on memory-1

Iterated Prisoner's Dilemma contains strategies that dominate any evolutionary opponent

William H. Press[a,1] and Freeman J. Dyson[b]

Remark 4.1. Robustness of results on direct reciprocity

- Yesterday, we discussed direct reciprocity, by looking at the repeated prisoner's dilemma.
- In particular, we talked about Axelrod's tournament. There I argued, that one should take TFT's success with a grain of salt. After all, the outcome of that tournament might depend a lot on which strategies you permit to participate.
- Then I argued that instead one should just consider all strategies in a simple but natural strategy space, the space of memory-1 strategies.
- But how can we be sure that those results are robust?

1. Perhaps if we did evolutionary simulations for memory-2 strategies, results would be different?
2. Perhaps strategies that are stable within the memory-1 space would cease to be stable if you allow for larger memory?

Theorem 4.2. Completeness of memory-1 strategies [P&D]

Suppose I play a memory-1 strategy **p**, and you play an arbitrary memory-k strategy **q**, and suppose the two of us get a payoff of (π_1, π_2) as a result.

Then one can find an equivalent memory-1 strategy **q*** for you, such that if you play that strategy instead, we still both get exactly the same payoffs.

When playing against a memory-1 player, anything you can do (with an arbitrary strategy), you can already do with a memory-1 strategy.

Corollary 4.3. Checking for Nash

If **p** is a memory-1 strategy, then (**p**,**p**) is a Nash equilibrium if and only if there is no profitable deviation towards another memory-1 strategy.

The impact of memory: More is different

Theorem 4.4. Checking for Nash, part II [Akin 2015]

Suppose **p** is a *nice* memory-1 strategy (it is never the first to defect, $p_0 = p_{CC} = 1$. Then(**p**,**p**) is a Nash equilibrium if and only if neither a deviation towards ALLD, nor a deviation to **q**=(0,0,1,1,1) is profitable.

Remark 4.5. Going to higher memory.

- The above results says that if I want to know whether some nice memory-1 strategy is a Nash equilibrium, I only need to check two possible deviations, instead of uncountably many possible deviations.
- Perhaps we can derive similar results for higher-memory strategies?
- More generally, it would be great to know: how does memory capacity affect the evolution of cooperation?

Remark 4.6. Why exploring larger memory is difficult.

There are at least two reasons why studying more general memory-k strategies is difficult.

- There are just too many of them:

  # of pure memory-1 strategies: 16

  # of pure memory-2 strategies: 65,536

  # of pure memory-3 strategies: $\sim 10^{19}$

  [# of pure memory-n strategies: $2^{2^{2n}}$]

- Even for two given strategies, computing payoffs becomes increasingly hard:

  Transition matrix for memory-1: 4x4
  Transition matrix for memory-2: 16x16
  Transition matrix for memory-3: 64x64
  …

The impact of memory: Reactive-n strategies

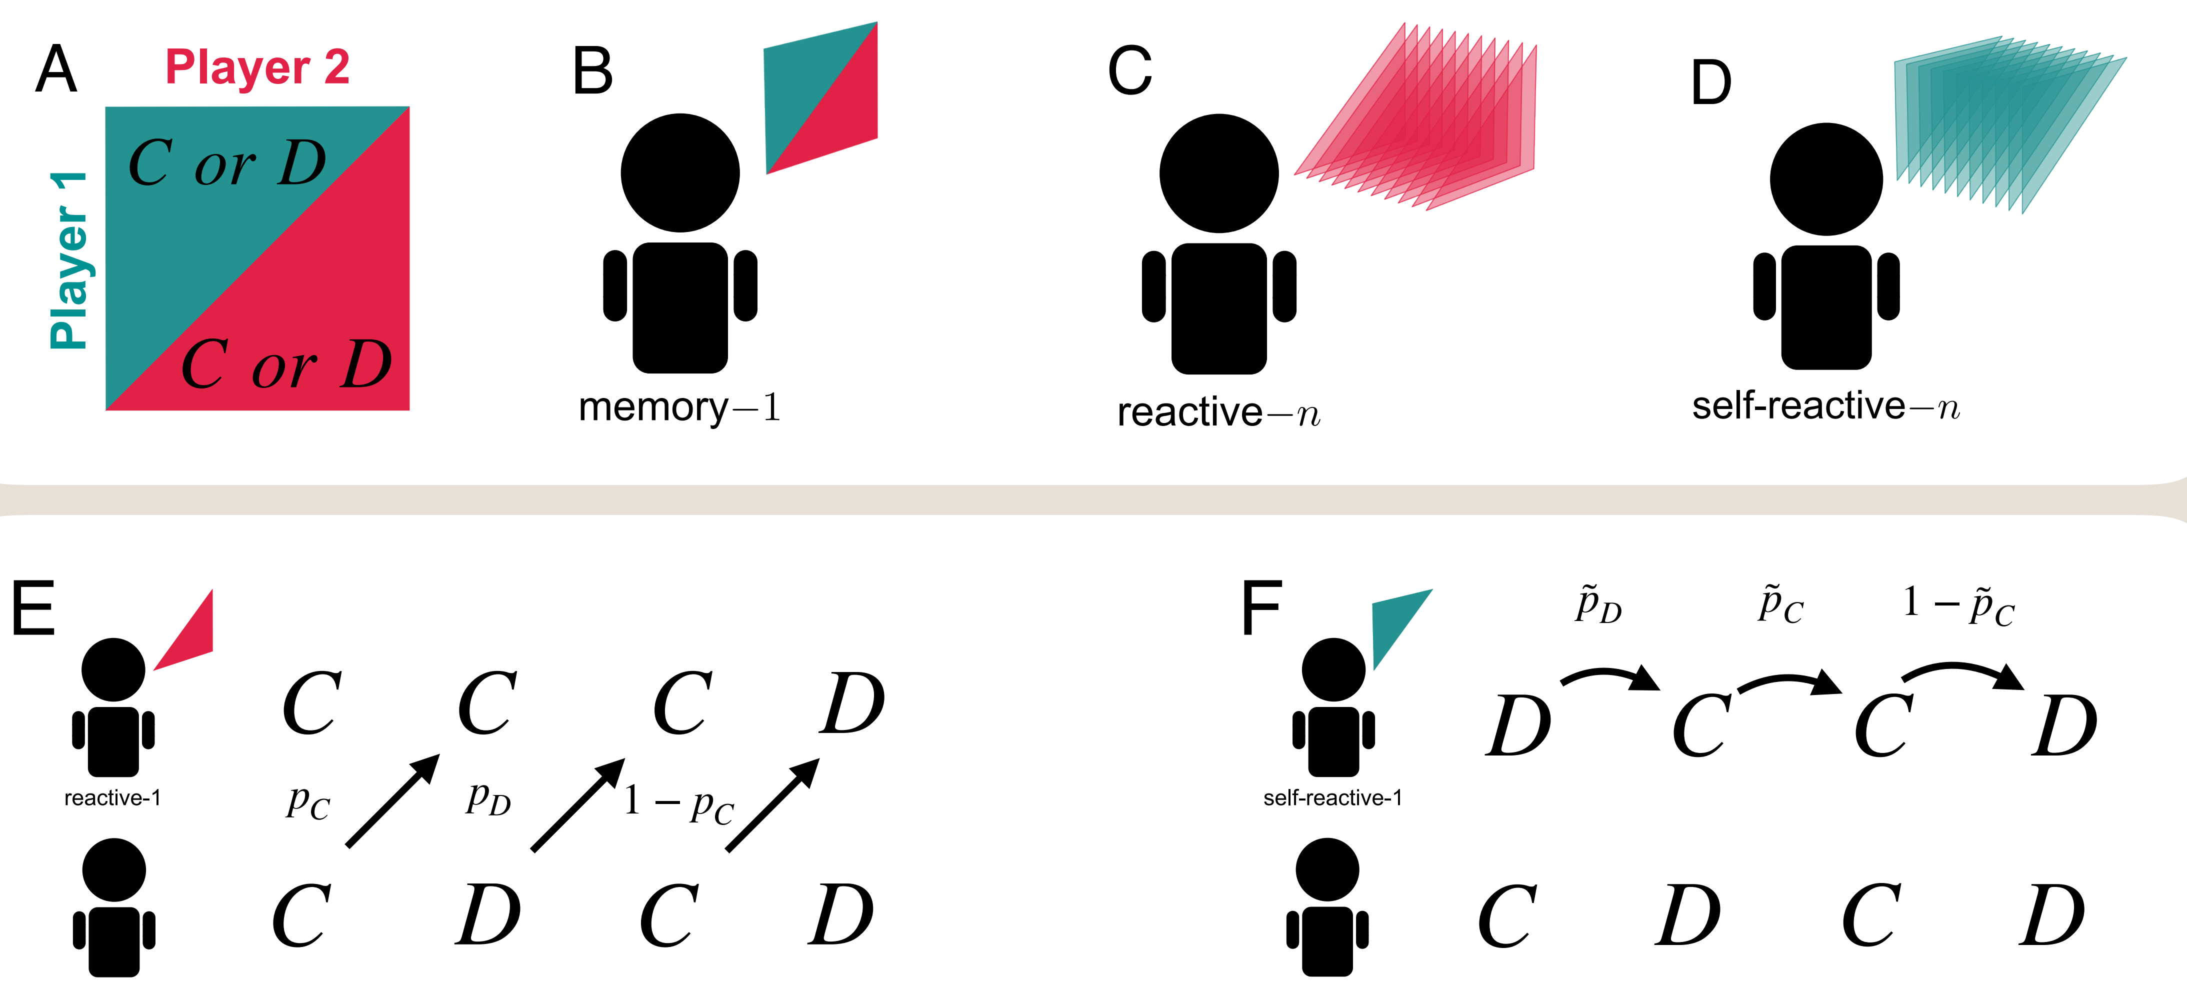

Definition 4.7. Reactive-n strategies

A memory-n strategy $\mathbf{q}$ is called reactive, if it only depends *on the co-player's* last n decisions.

Example 4.8. Some reactive-n strategies

- Reactive-1 strategies: $\mathbf{q} = (q_C, q_D)$
  ALLD = (0,0), ALLC=(1,1), TFT = (1,0).
- Reactive-2 strategies: $\mathbf{q} = (q_{CC}, q_{CD}, q_{DC}, q_{DD})$
  Tit-for-Two-Tat = (1,1,1,0).

Theorem 4.9. Best responses against reactive-n players

To any memory-n strategy $\mathbf{q}$, one can find a best response $\tilde{\mathbf{q}}$ among the pure self-reactive-n strategies.

[In fact, if the game is a donation game, there is a best response $\tilde{\mathbf{q}}$ among the pure self-reactive-(n-1) strategies.]

Theorem 4.10. Stable reactive-2 strategies

A nice reactive-2 strategy $\mathbf{q} = (q_{CC}, q_{CD}, q_{DC}, q_{DD})$ forms a Nash equilibrium (in the repeated donation game) if and only if

$$q_{CC} = 1, \qquad \frac{q_{CD} + q_{DC}}{2} \leq 1 - \frac{1}{2}\frac{c}{b}, \qquad q_{DD} \leq 1 - \frac{c}{b}.$$

In particular,

- For every defection in memory, you reduce your cooperation probability proportionally.
- The exact timing of your defections does not matter.

One can derive similar conditions for reactive-3 (reactive-n).

The impact of memory: Reactive-n strategies

Remark 4.11. Evolution of Reactive-n strategies

- Consider a population of size N
- Players adopt reactive-n strategies
- They play against all other population members to get a payoff
- Strategies that yield a higher payoff are more likely to be imitated.

What is the effect of memory on evolving cooperation rates?

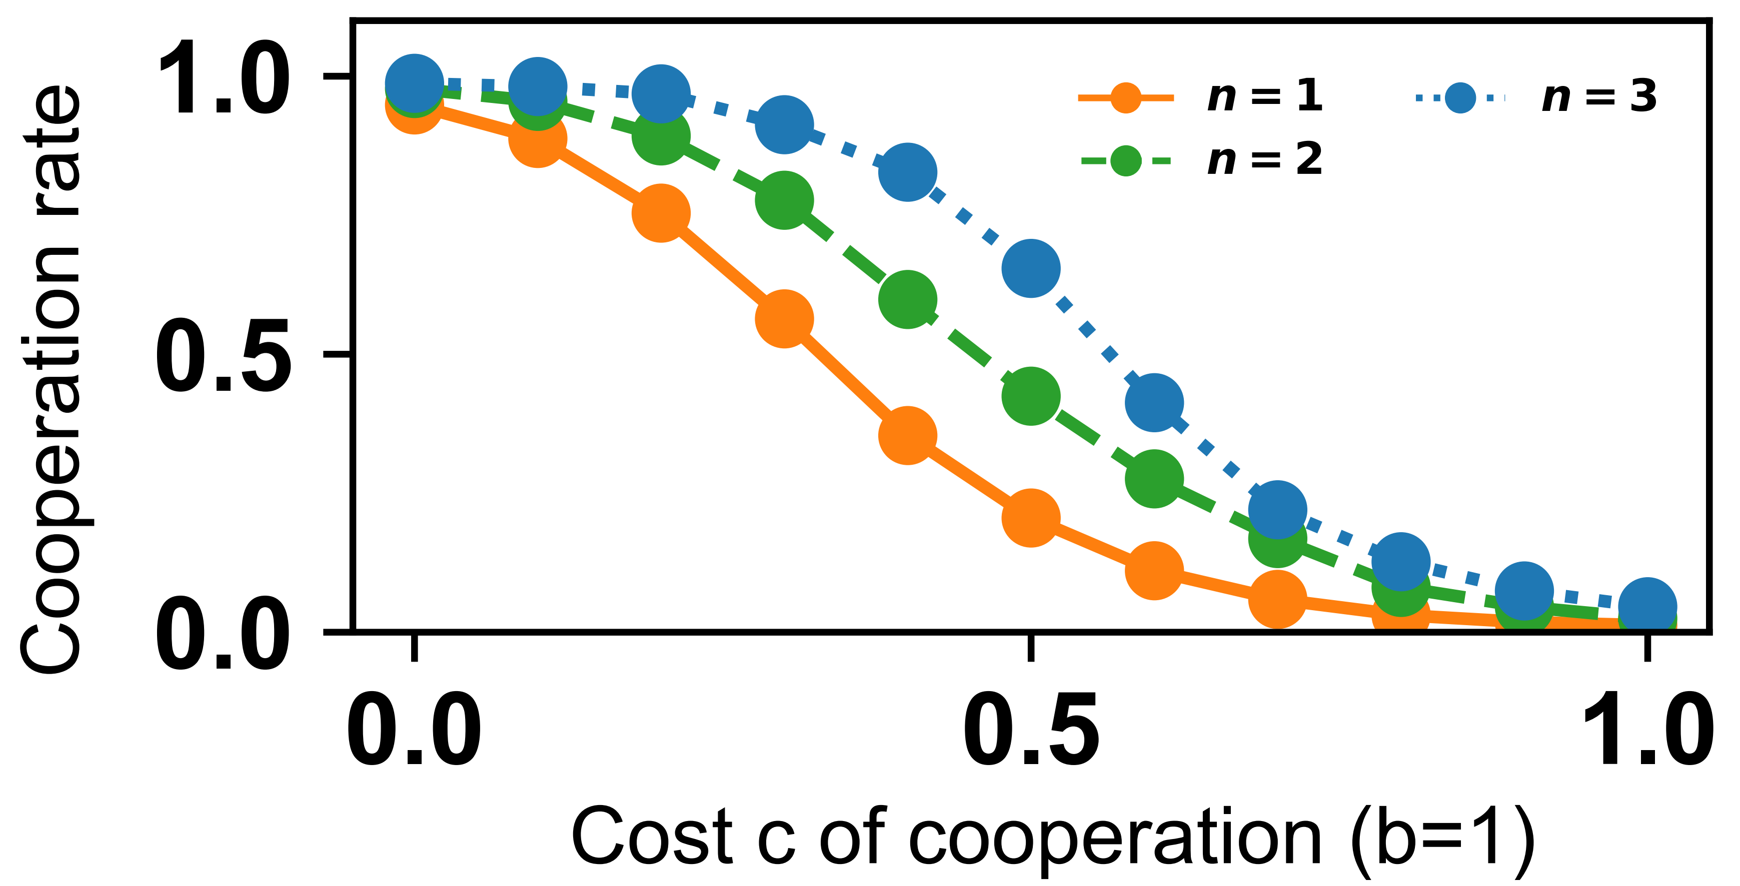

Remark 4.12. Summary

- It is interesting to explore the impact of memory capacities on direct reciprocity.
- A formal analysis can be tricky, because the size of the strategy space quickly explodes
- Still, some analytical results are feasible
- Simulations suggest that more memory helps cooperation.

An overview

Two days ago's class (March 11, 2025)

- An introduction to evolutionary game theory
  (Replicator dynamics, games in finite populations)

Yesterday's class (March 12, 2025)

- Evolution of cooperation & direct reciprocity
- Social norms & indirect reciprocity

Today's class (March 13, 2025)

- Some current research: Reciprocity in complex environments

The role of memory; the effect of changing environments; the impact of inequality

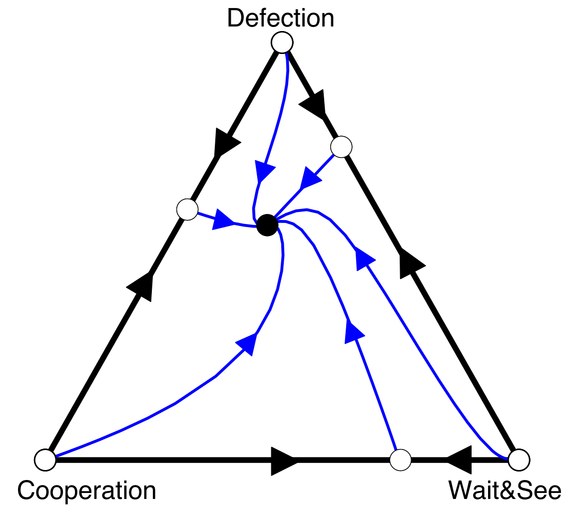

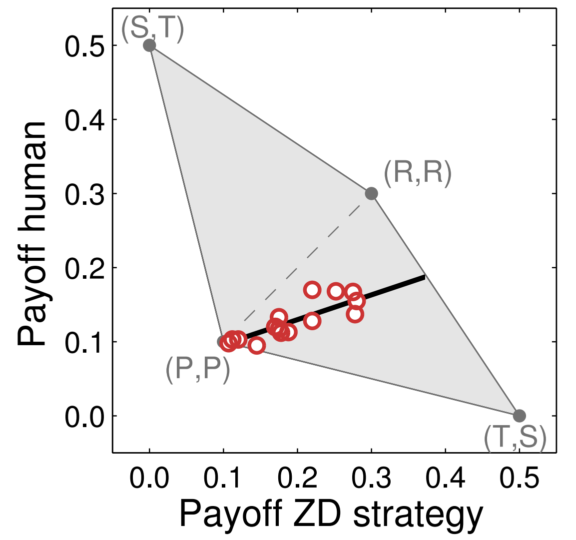

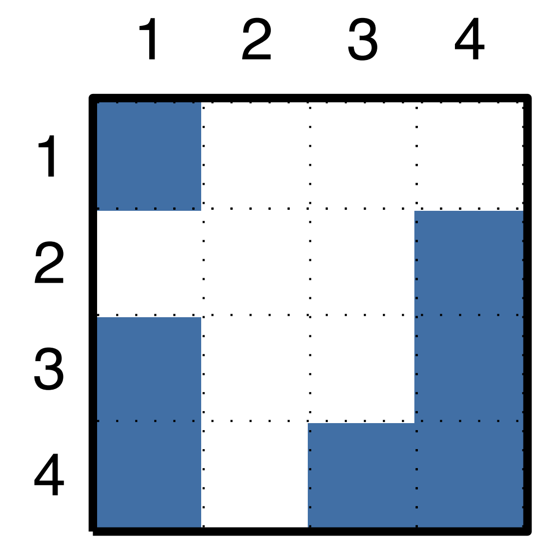

Beyond the repeated prisoner's dilemma

Remark 4.13. Is the repeated prisoner's dilemma the answer?

The repeated prisoner's dilemma is an extremely useful model to study reciprocity, but it is also quite stylised:

Symmetry: Players are identical with respect to their possible actions and payoffs.

Stationarity: The players' environments (the games they play) do not change.

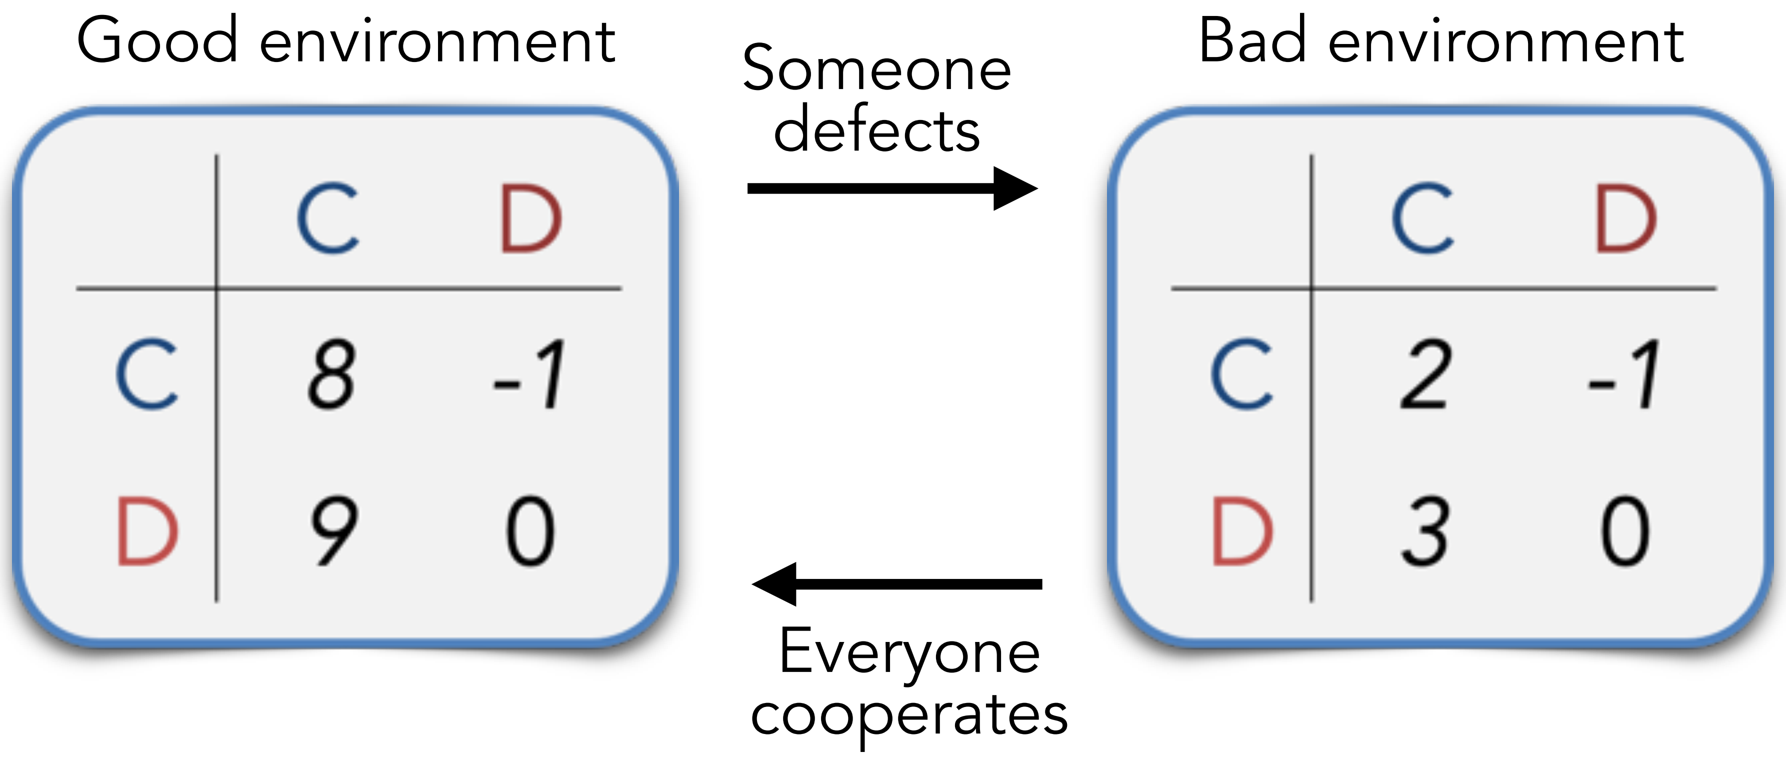

No longer clear whether strategies like Tit-for-Tat can promote cooperation.

Remark 4.14. Stochastic games

- We consider a set of n players
- These players can find themselves in m different (environmental) states
- In each state they play some (different) one-shot game, where they can either cooperate or defect.
- The players' actions determine their payoffs, but they also determine in which state they are next.
- We assume players use memory-1 strategies.
  In this case, memory-1 strategies depend on the outcome of the previous round, and on the current state. For example, if there are N=2 players and M=2 states, then $\mathbf{p} = (p^1_{CC}, p^1_{CD}, p^1_{DC}, p^1_{DD},\ p^2_{CC}, p^2_{CD}, p^2_{DC}, p^2_{DD})$
- A player's payoff in the stochastic game is the player's average payoff per round, $\pi_i = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} \pi_i(t)$

Evolution of cooperation in stochastic games

Example 4.15. Evolution of cooperation in stochastic games

- We consider a set of N players
- Players are randomly matched in groups of n, and then play the stochastic game against each other.
- Strategies that yield a high payoff are more likely to be imitated by other players.

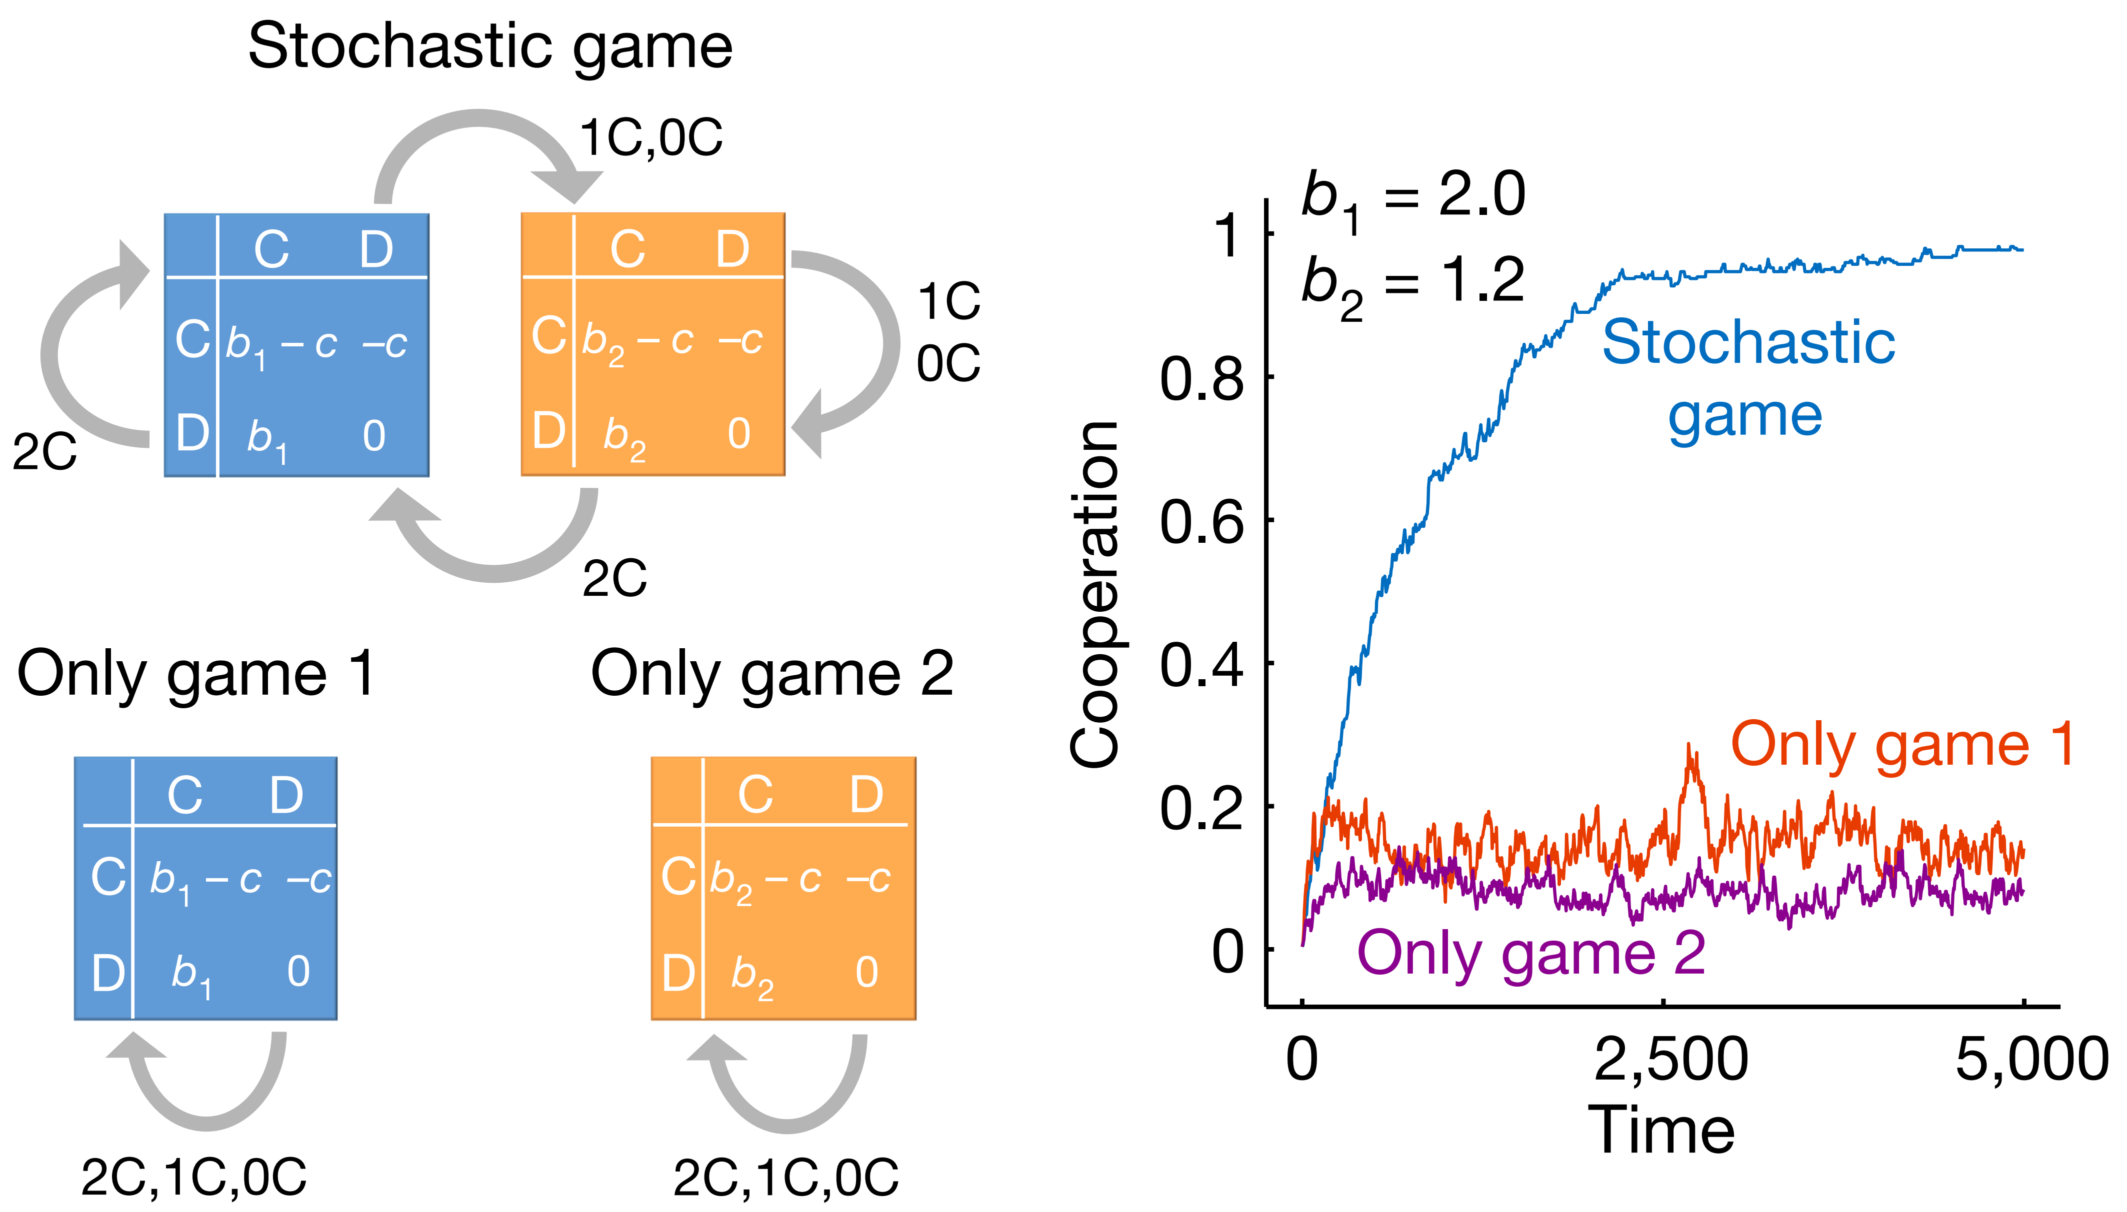

Remark 4.16. Dependence on state transitions

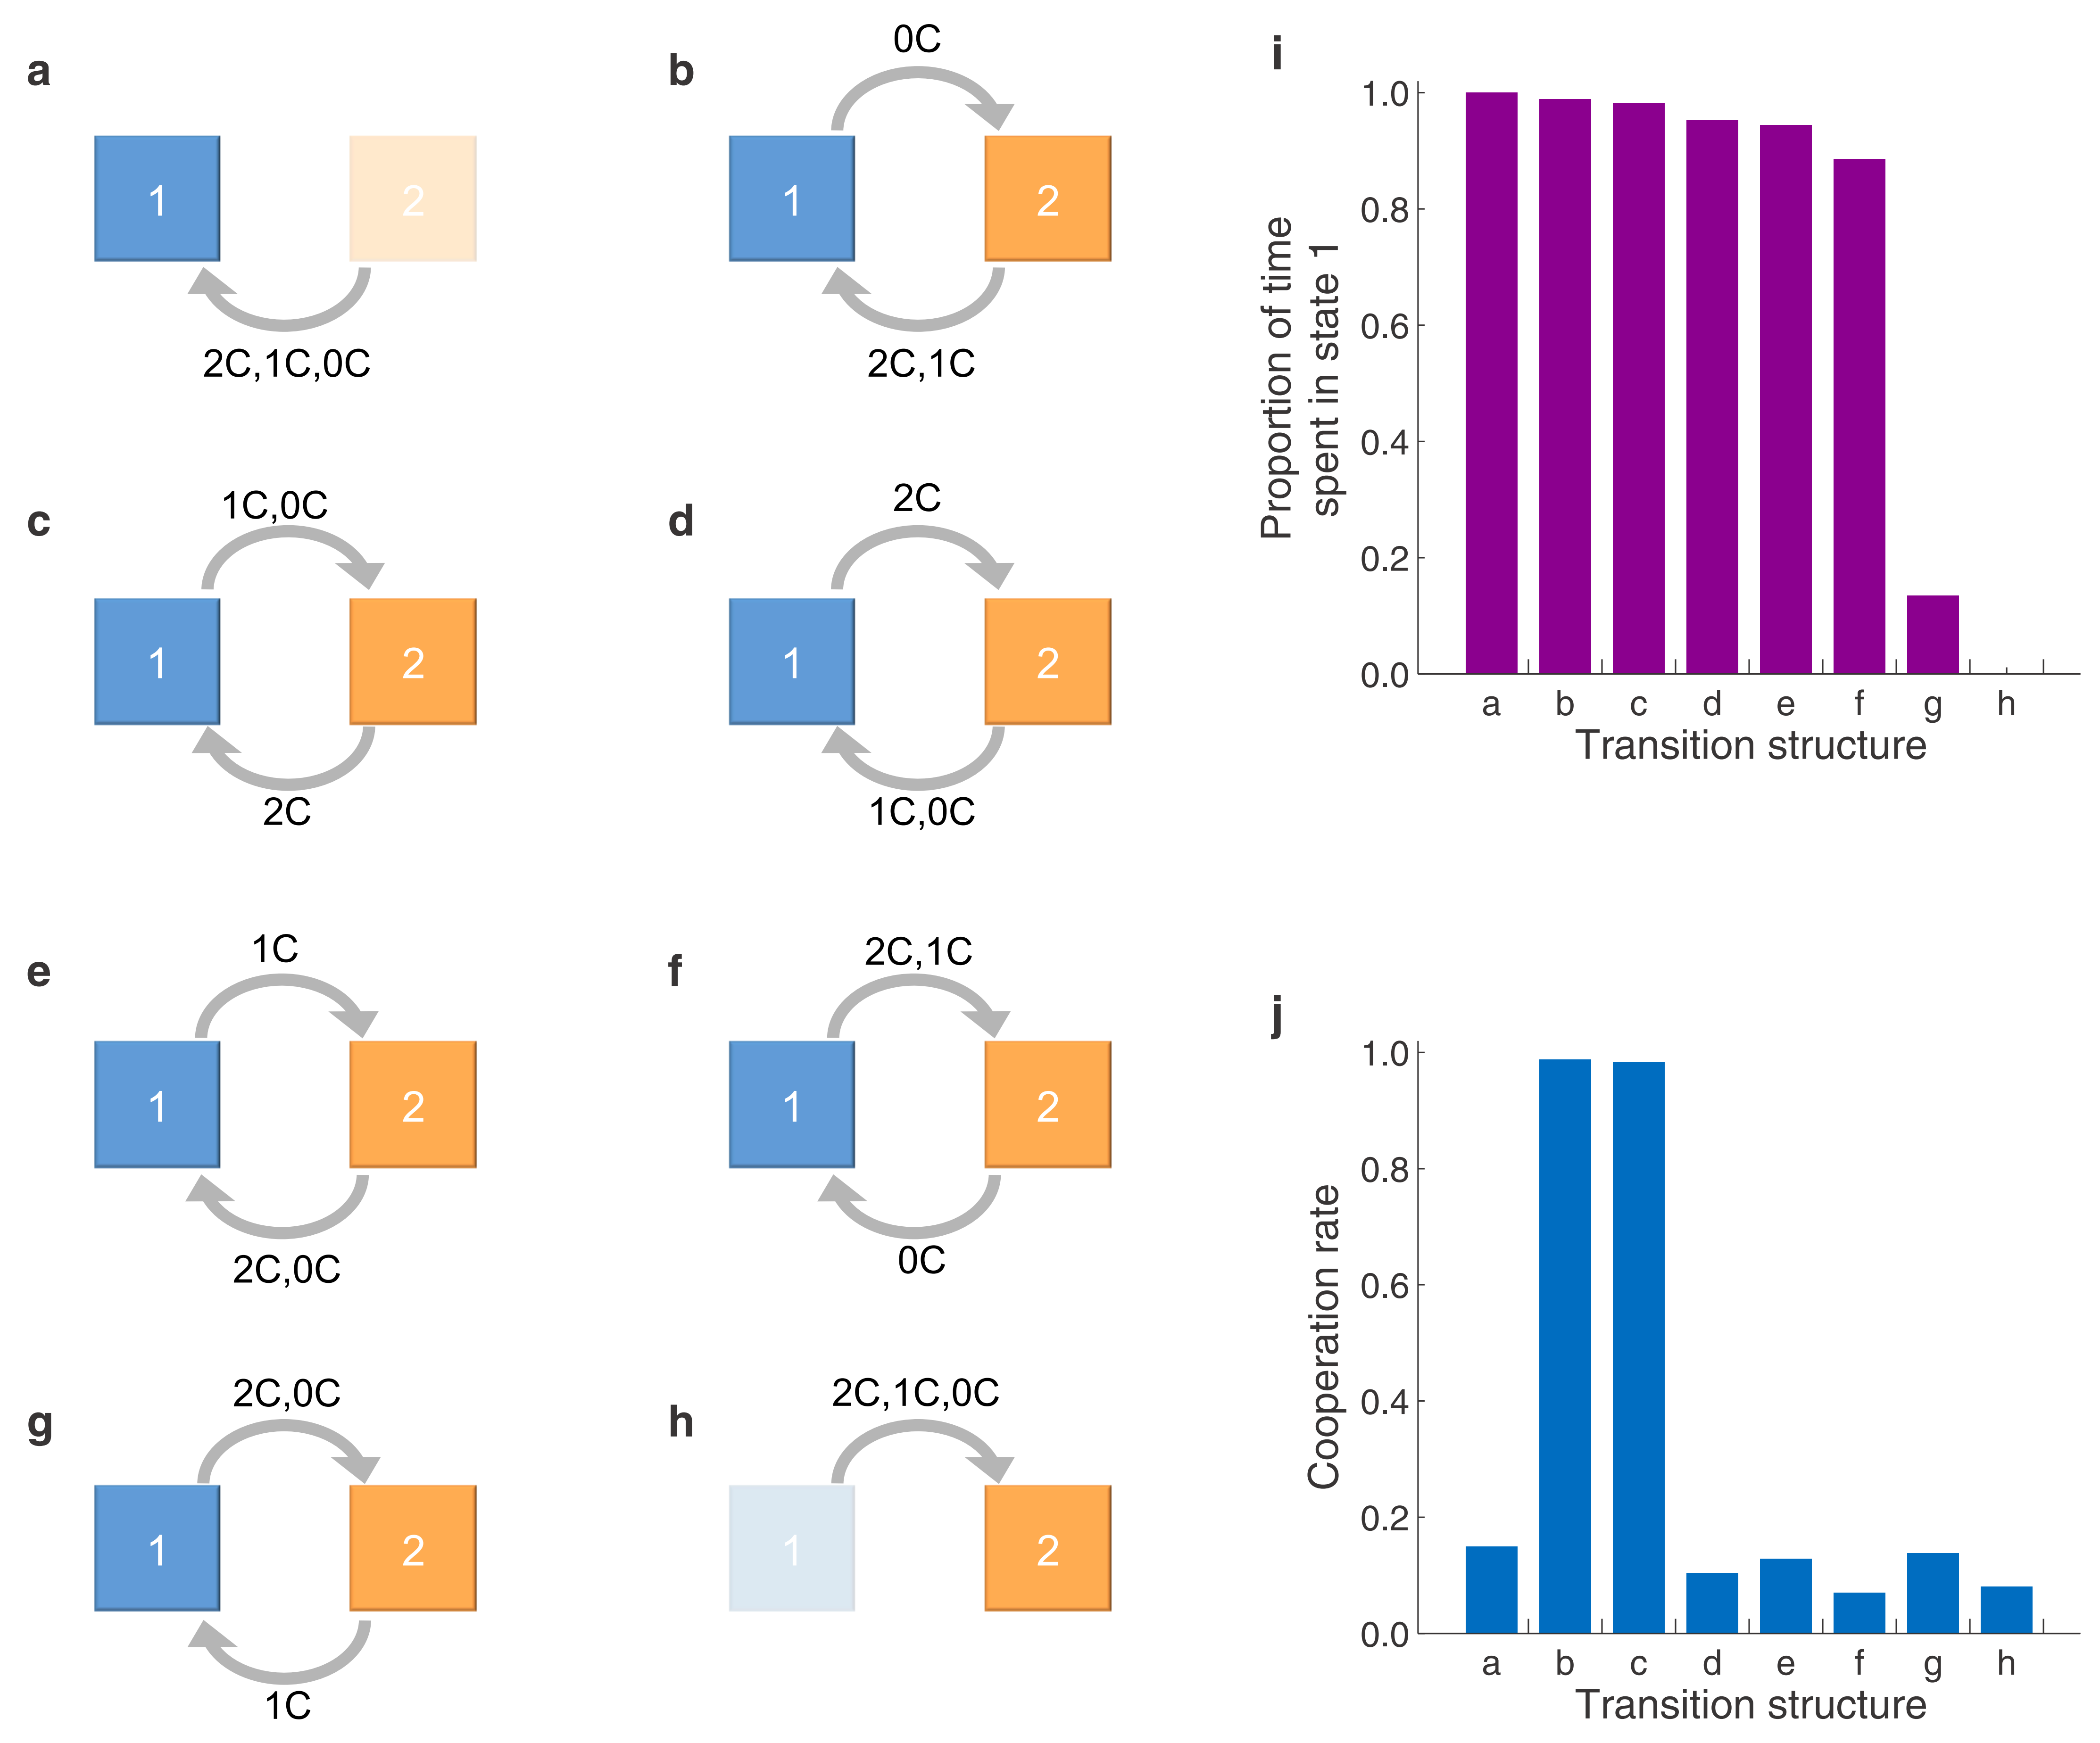

Evolution of cooperation in stochastic games

Example 4.17. Cooperation in larger groups

- Suppose now the game is played in groups of 4 players
- In each state, players play a public goods game with multiplication factor r, which might depend on the state.
- The number of cooperators determines the next state.
- We consider four different treatments, which differ in how easily the environment deteriorates (how quickly the multiplication factor decreases if few people cooperate).

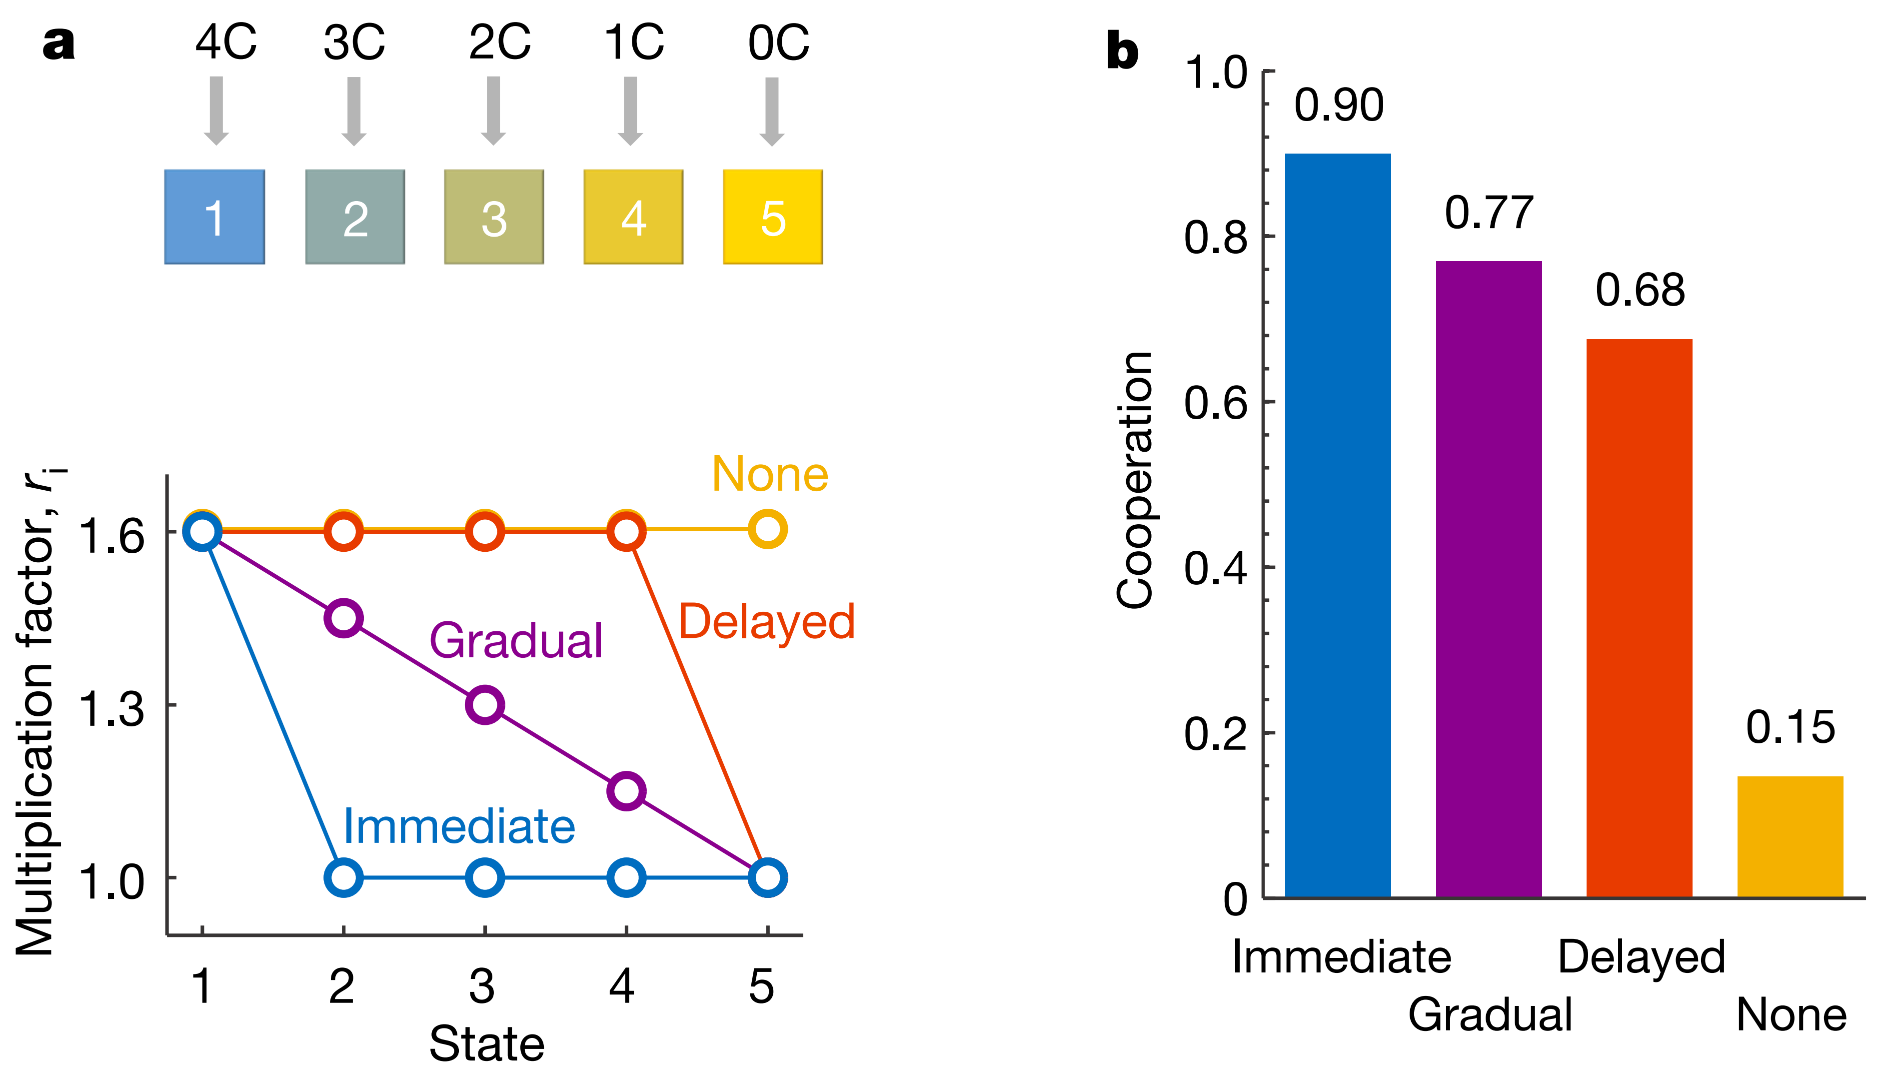

Remark 4.18. Summary

- When individual actions do not only affect payoffs, but also the players' environment, this can favour the evolution of cooperation.
- Cooperation is most favored when defection results in a quick deterioration of the environment, leading players to interact in more unprofitable games.

An overview

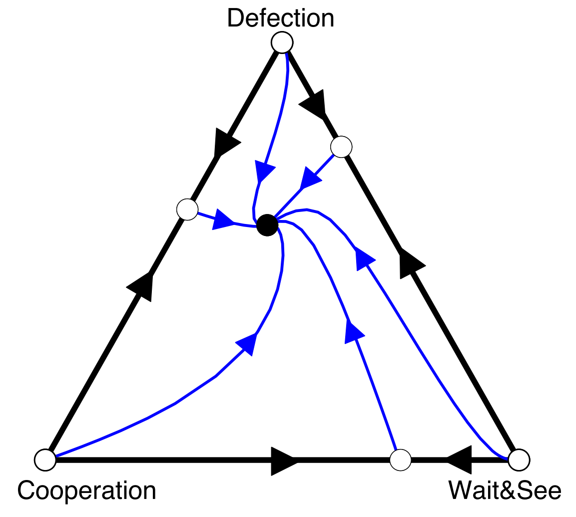

Two days ago's class (March 11, 2025)

- An introduction to evolutionary game theory
  (Replicator dynamics, games in finite populations)

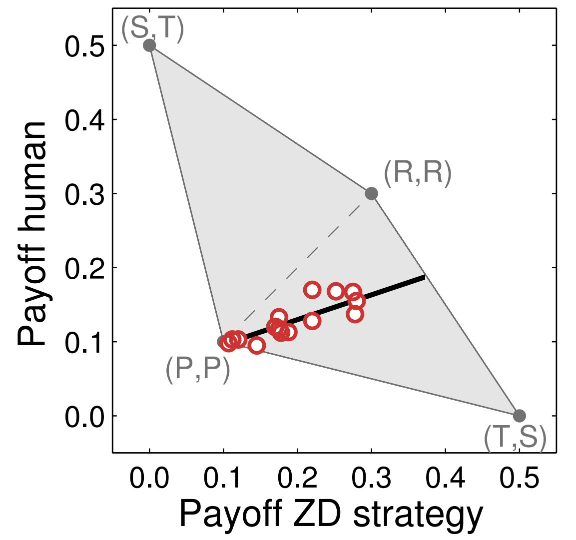

Yesterday's class (March 12, 2025)

- Evolution of cooperation & direct reciprocity
- Social norms & indirect reciprocity

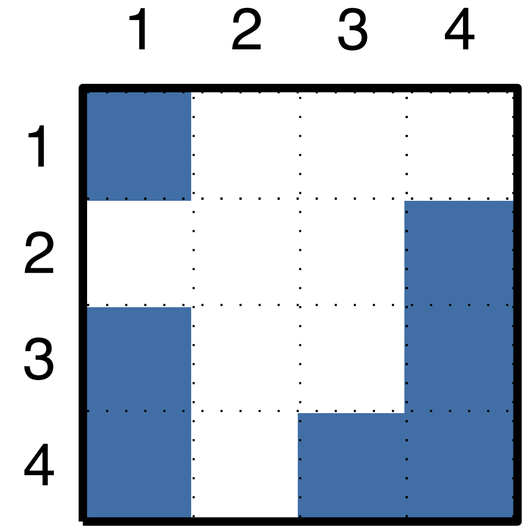

Today's class (March 13, 2025)

- Some current research: Reciprocity in complex environments

The role of memory; the effect of changing environments; the impact of inequality

Cooperation in asymmetric games

Remark 4.19. On the role of symmetric games

- In evolutionary game theory, we really like symmetric games.
- Makes the description & the math easier, simplifies interpretation of results, allows for learning by imitation
- Many social dilemmas are asymmetric

Why asymmetry is nontrivial

- No longer clear whether strategies like Tit-for-Tat are effective
- No longer clear how we should model learning in asymmetric games

Previous work

- Quite some work on cooperation in asymmetric social dilemmas
- Is endowment inequality always detrimental?

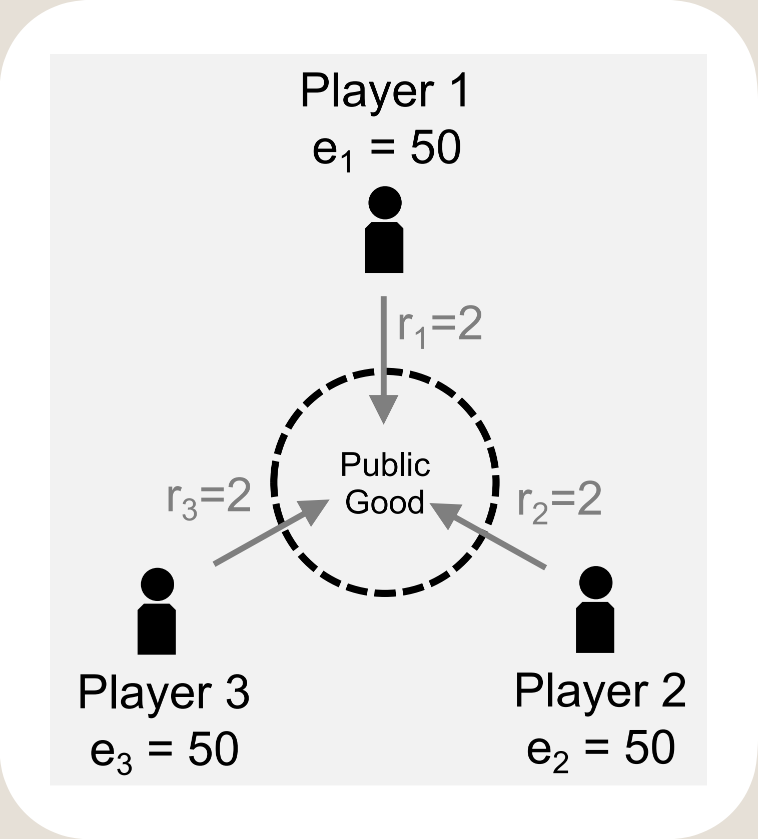

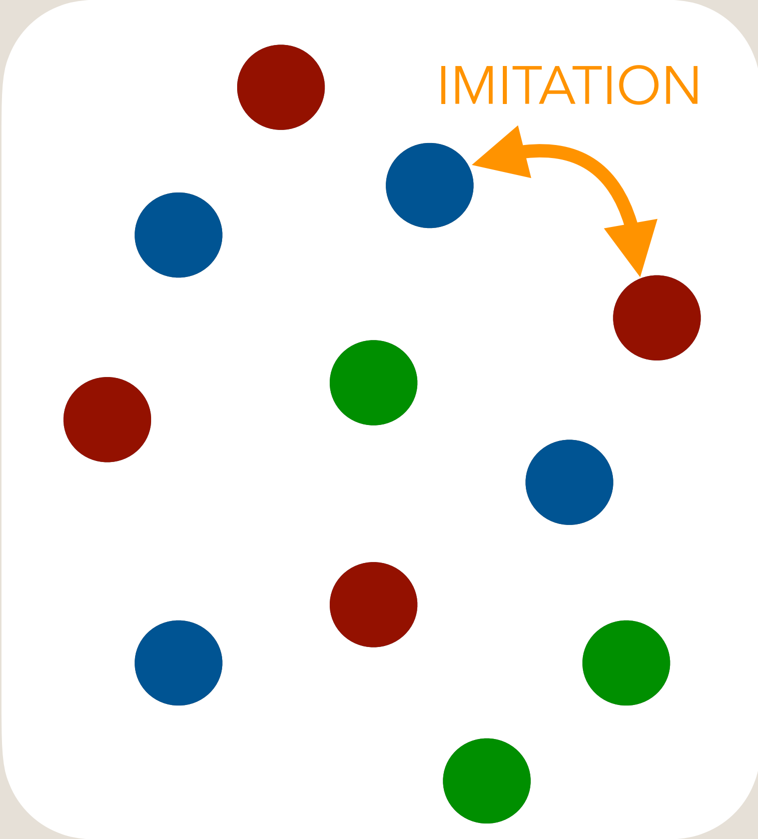

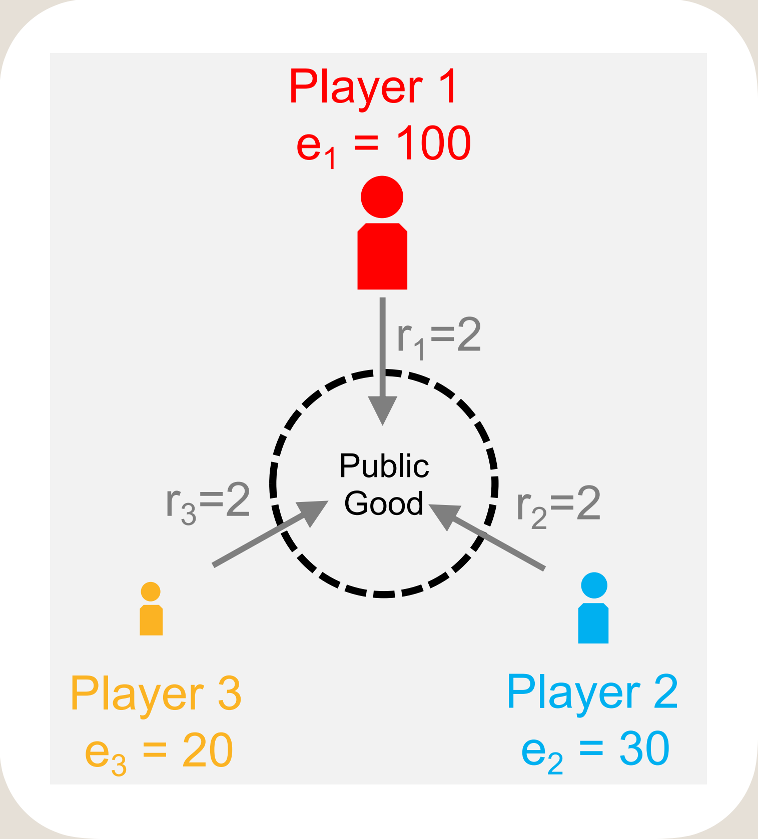

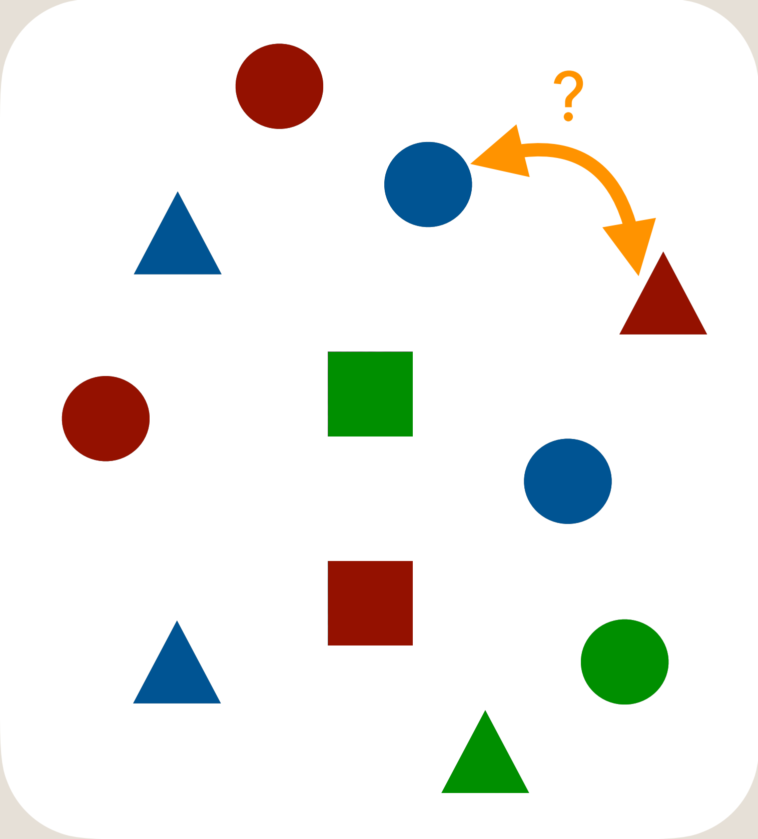

Climate policies under wealth inequality

Vítor V. Vasconcelos[a,b,c], Francisco C. Santos[a,c], Jorge M. Pacheco[a,d,e], and Simon A. Levin[f,g,h,1]

Cooperative interaction of rich and poor can be catalyzed by intermediate climate targets

A letter

Manfred Milinski · Torsten Röhl · Jochem Marotzke

Cooperation in asymmetric games

Remark 4.20. A model of games among unequals

- There is a group with n individuals who interact repeatedly
- Each round, individual i obtains an endowment e_i
- Individuals independently decide how much to contribute
- Individual i's contribution is multiplied by r_i
- Total contributions are evenly split

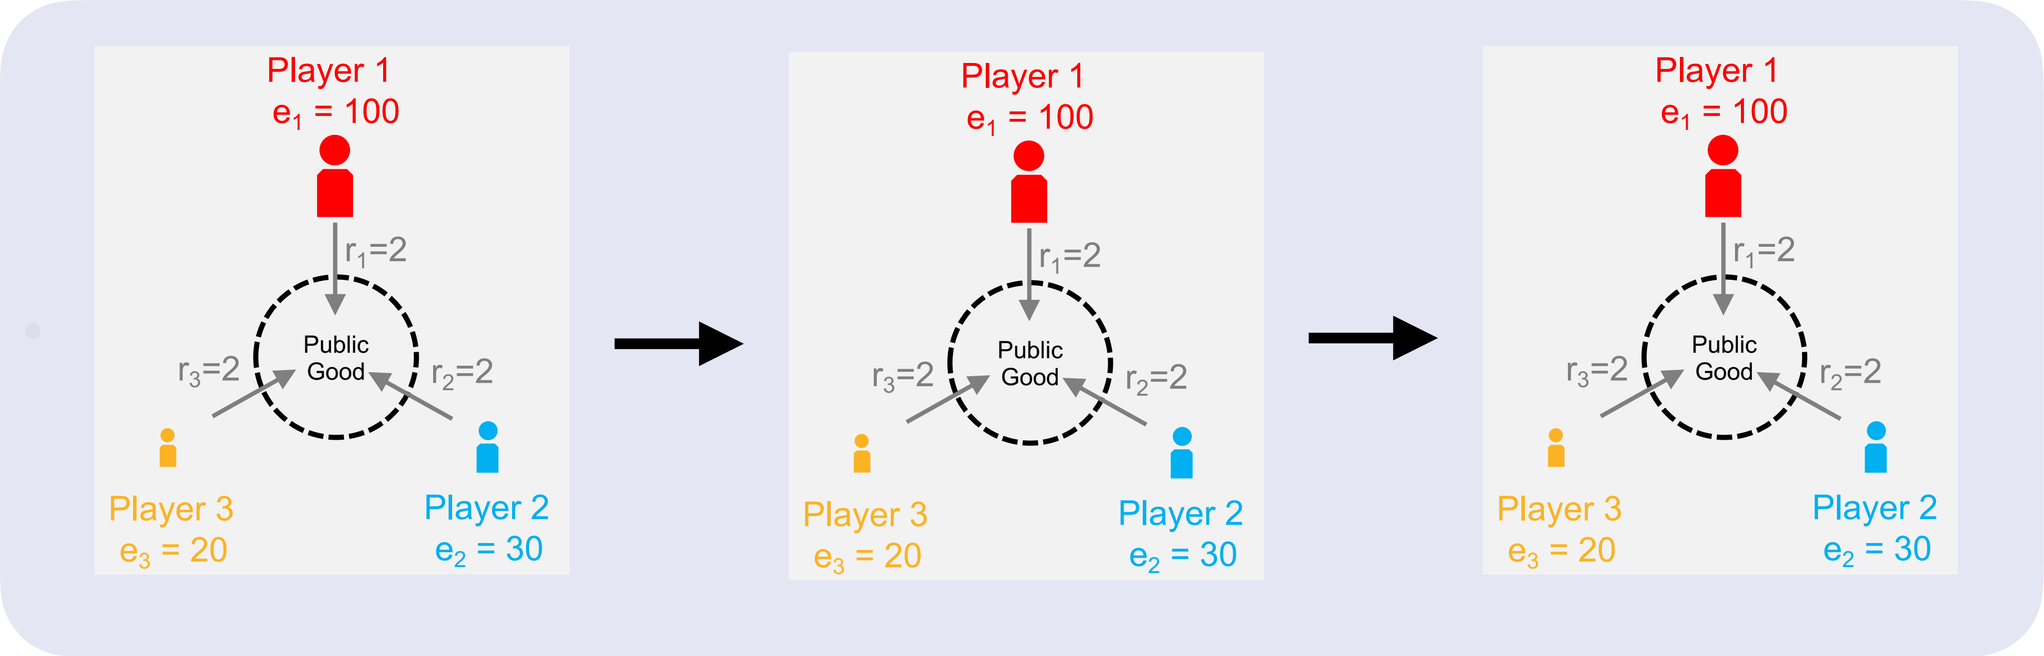

Research question

For given productivities, how should we optimally allocate endowments to maximize cooperation?

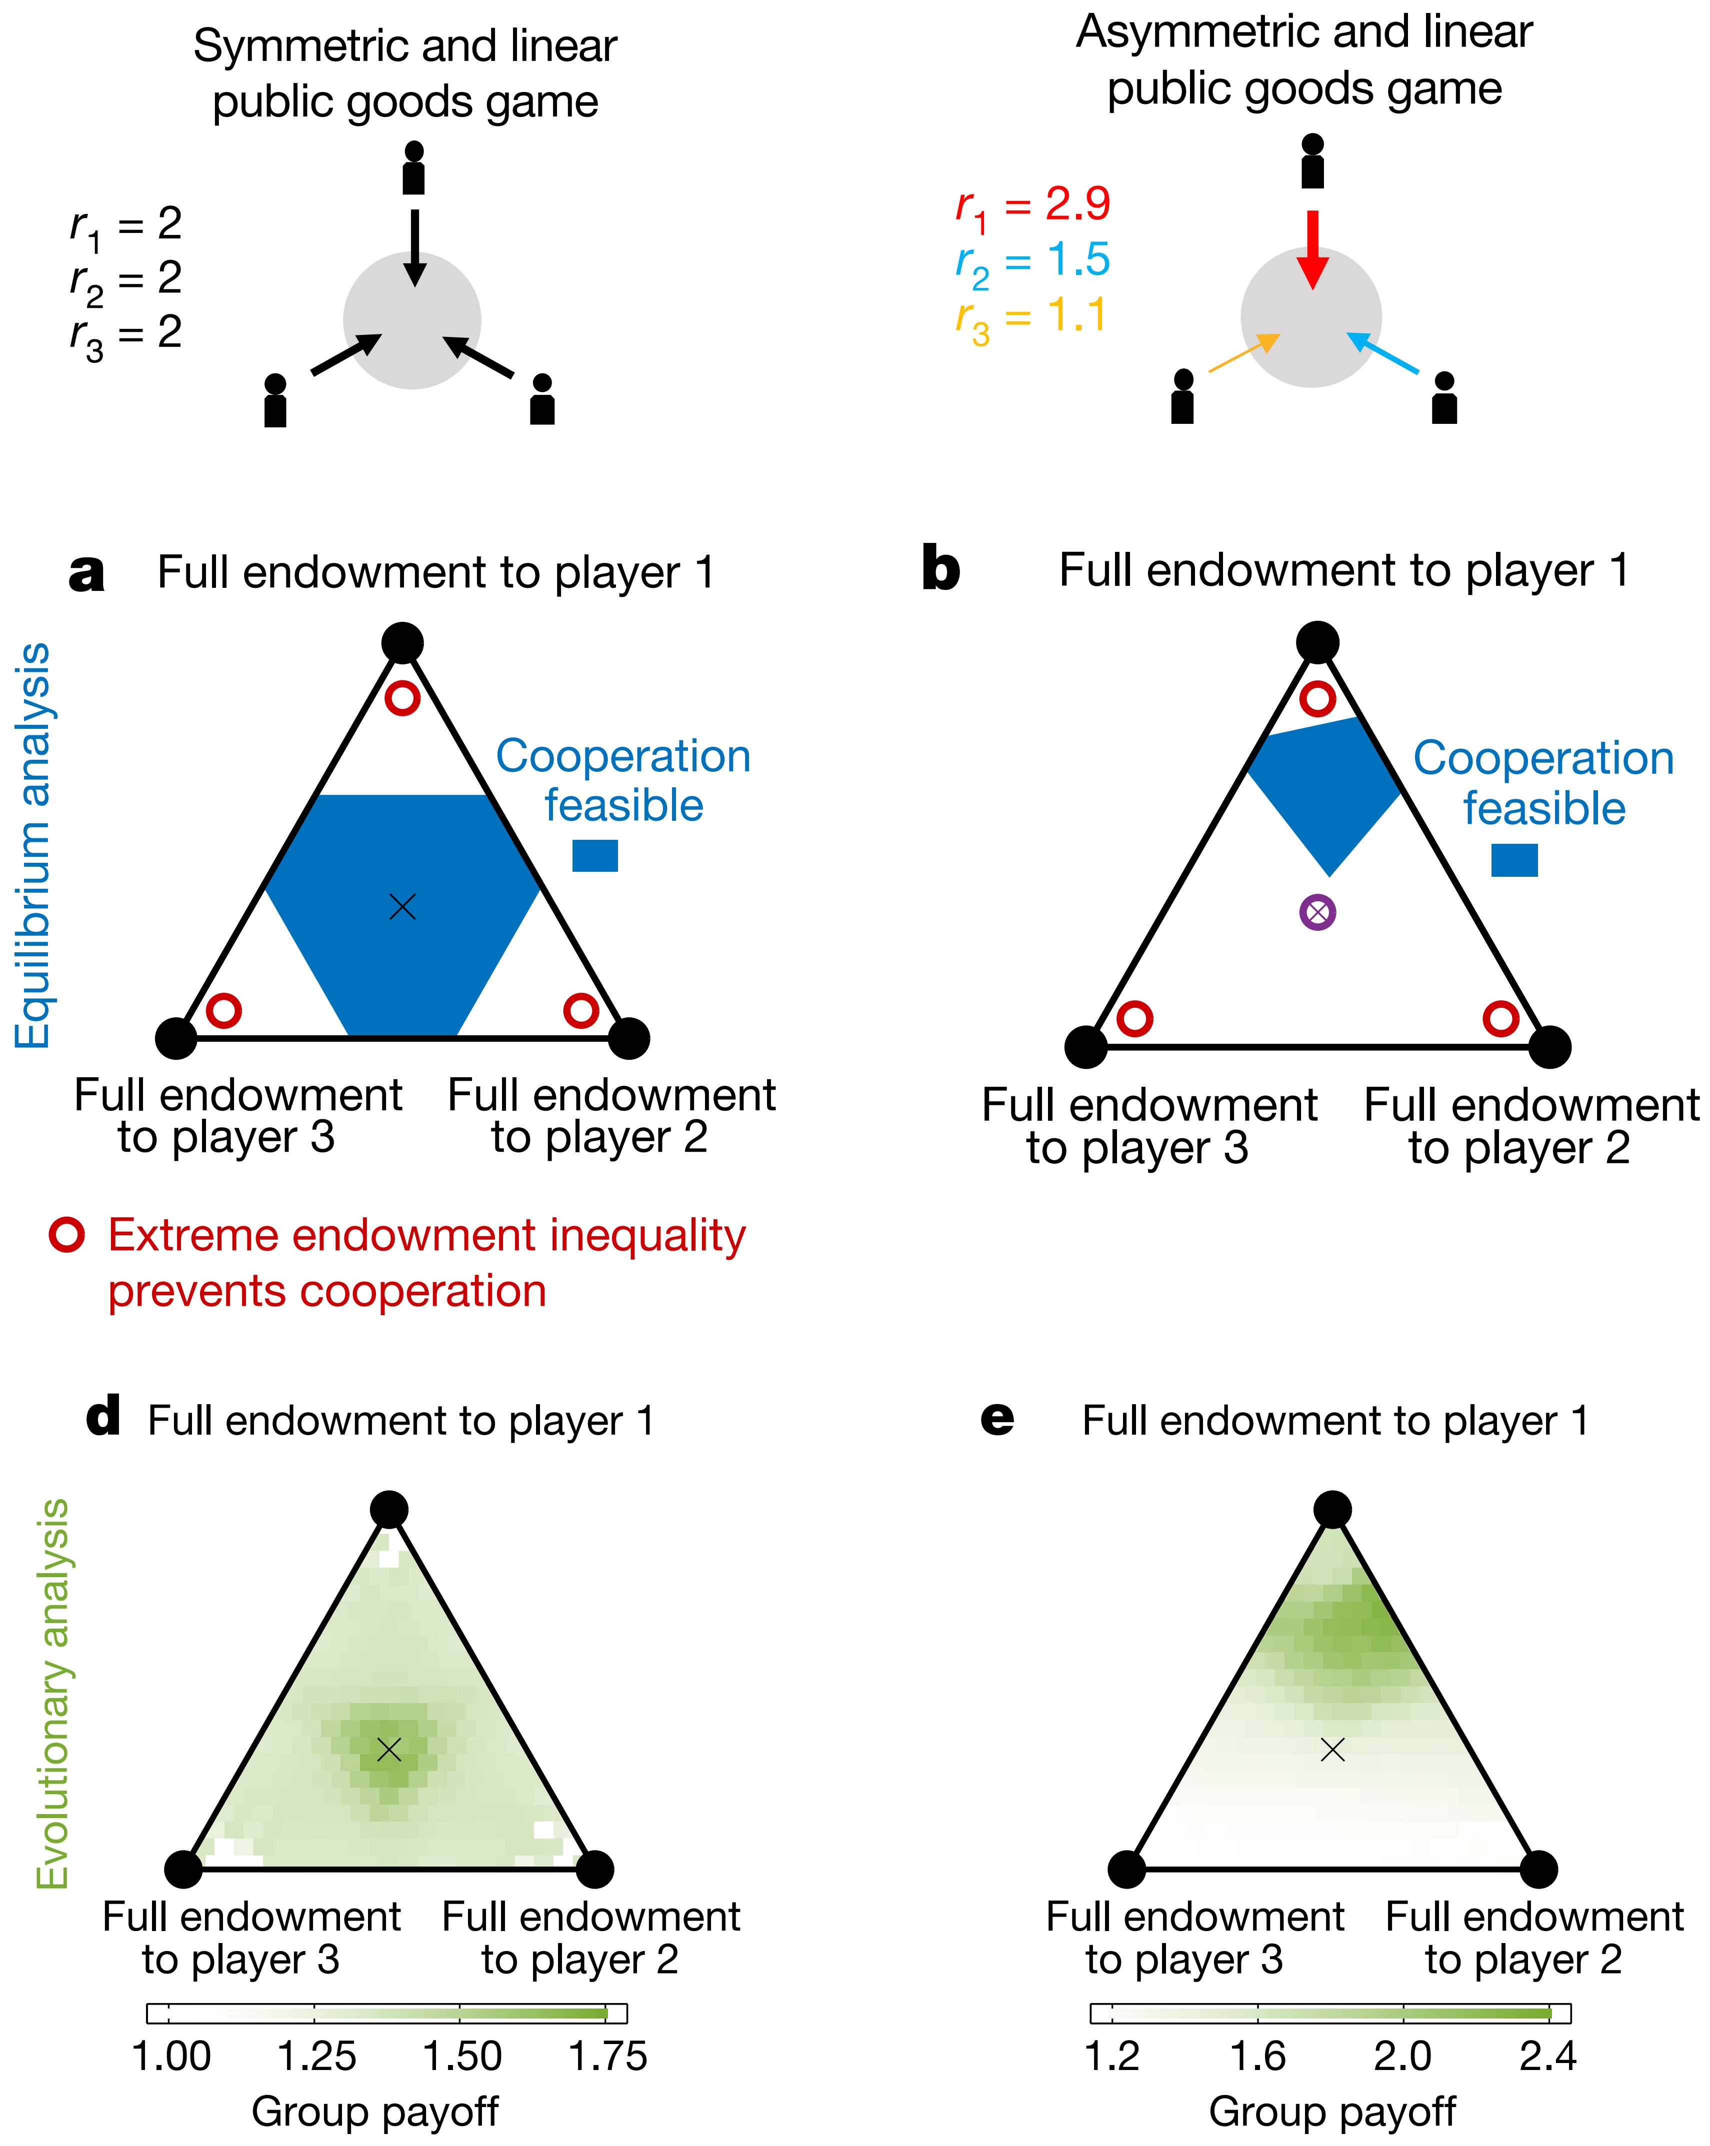

Cooperation in asymmetric games

Remark 4.21. An experiment

- To test these qualitative predictions, we did an experiment
- ~400 participants recruited through Amazon Turk
- They play a repeated public good game in groups of two
- All groups interact for at least 20 rounds
- Endowments can either be equal or unequal; productivities can either be equal or unequal.

a

Equal productivities — Unequal productivities

Equal endowments

1 Full equality

	Player 1	Player 2
Endowment	50	50
Productivity	×1.6	×1.6

3 Productivity inequality

	Player 1	Player 2
Endowment	50	50
Productivity	×1.9	×1.3

Unequal endowments

2 Endowment inequality

	Player 1	Player 2
Endowment	75	25
Productivity	×1.6	×1.6

4 Aligned inequality

	1	2
Endowment	75	25
Productivity	×1.9	×1.3

5 Misaligned inequality

	1	2
Endowment	25	75
Productivity	×1.9	×1.3

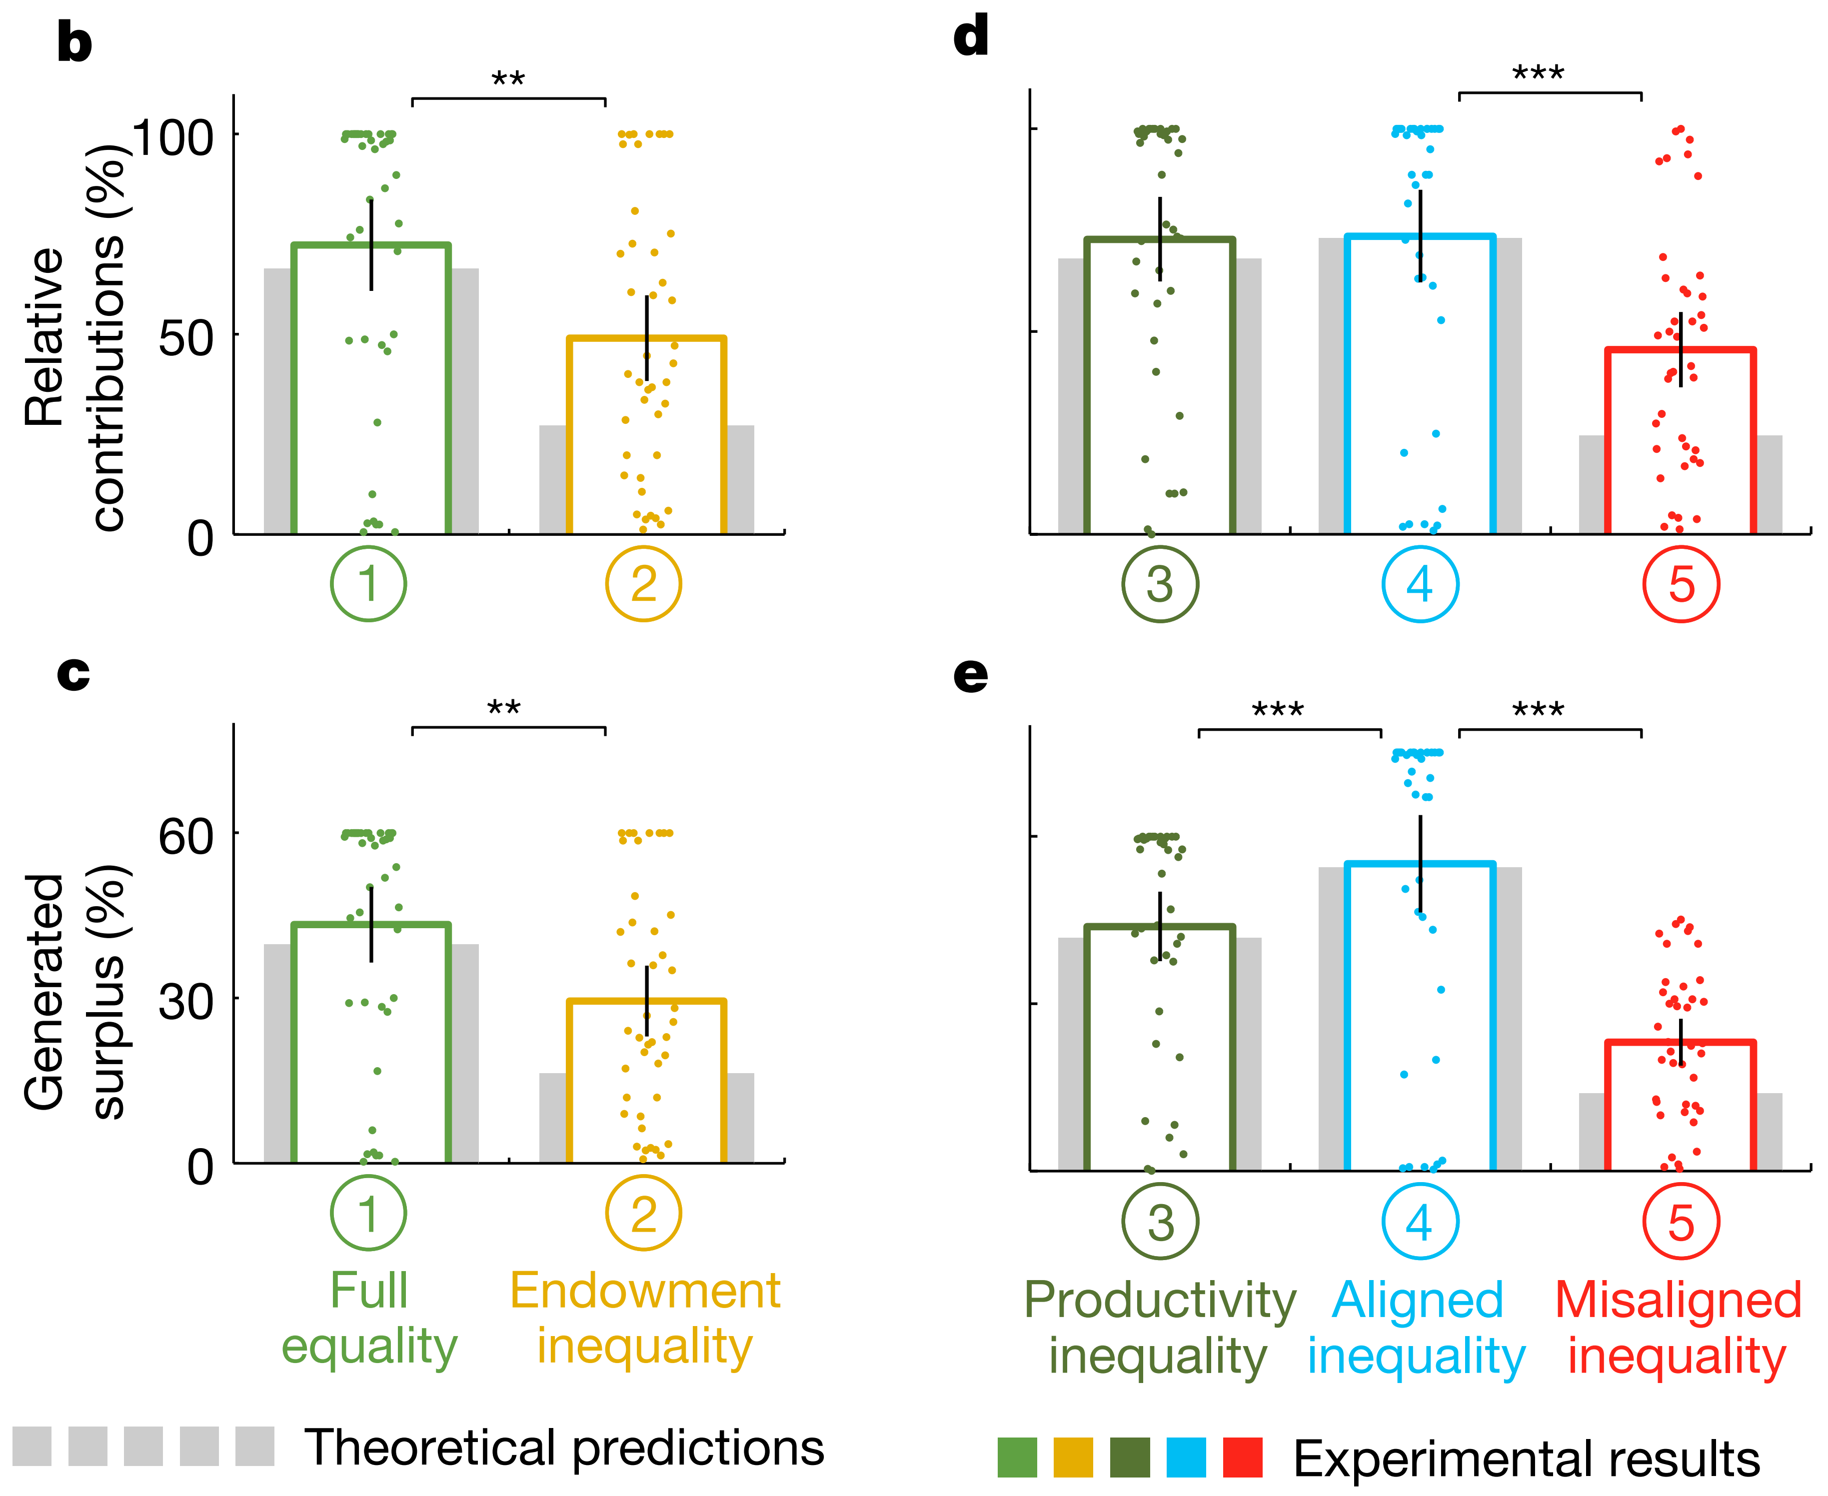

Summary: How to allocate endowments to maximise cooperation?

- When players are equally productive, they should get equal contributions.
- Otherwise, more productive players should get higher endowments (*'aligned inequality'*)

Research Group Dynamics of Social Behavior

Summary

1. In my lectures, I first provided some introduction to evolutionary game theory (replicator dynamics, Moran process).
2. Then we used to these techniques to further explore the evolution of cooperation (in particular: direct and indirect reciprocity).
3. Thanks to the organisers, and thanks for being such an engaging audience!

Where to find more information:
http://web.evolbio.mpg.de/social-behaviour/

Collaborators and funding

Martin A. Nowak, Harvard University
Krishnendu Chatterjee, IST Austria
Stepan Simsa, Rossum.ai
Maria Kleshnina, Queensland University Of Technology
Oliver Hauser, University of Exeter
Nikoleta Glynatsi, RIKEN
Franziska Lesigang, TU Vienna
Ethan Akin, City University New York
Karl Sigmund, University of Vienna
Yohsuke Murase, RIKEN
Moshe Hoffman, Harvard University
Josef Tkadlec, Charles University Prague
Arne Traulsen, MPI for Evolutionary Biology
Manfred Milinski, MPI for Evolutionary Biology
Laura Schmid, Nature Communications
Other Dynos, all over the world

IT:U interdisciplinary transformation university austria

erc