Evolving Cooperation in Games with Agent-Based Models (ABM)

David Hales

dave@davidhales.com

Slides and simulation models can be found at:
www.davidhales.com/talks/szeged2016

Agent-based modelling

- Modelling what agents do in simulations is called agent-based modelling
- By experimenting with such models insights can be gained without:
  - Having a theory / analytically tractable model
  - Knowing what will happen in advance
- Hence working with such models is more like empirical science than mathematics

Robert Axelrod used ABM in his tournaments by getting people to send him different strategies as programs (agents). By playing them against each other he discovered tit-for-tat won

Ideal rationality

- If agents interact by playing games and receiving payoffs
- Agents can select a strategy using the ideal rationality of game theory, assuming agents:
  - Know the game
  - Know all other agents know the game
  - Know all other agents will act rationally
  - Know all other agents know that all other agents know all of this

Where rational means maximising expected payoff and agents are cleaver (have unlimited computation)

Evolving strategies in games

- An alternative approach is to evolve strategies
- Agents store strategies and engage in game interactions with others in a population
- The payoffs agents receive are interpreted as "fitness" such that strategies with:
  - high payoff increase in the population
  - low payoff decrease in the population
- With low probability agents randomly change their strategy (mutation)

Evolving strategies in games

- Agents can be viewed as biological entities evolving through reproduction and death
- Or cultural entities evolving their strategies through imitation (copying) and random innovation (trying something new)
- Or software entities e.g. P2P clients, where strategies are a combination of the client software + the user behaviour

Evolving strategies: a generic algorithm

- An initial population of strategies are generated (from some distribution)
- Repeat some number of generations:
  – Game interactions occur between strategies
  – A new generation of strategies are produced based on their payoffs (fitness)
  – Mutation applied to strategies by randomly changing them with low probability

In general when evolution is applied to games reproduction is asexual (i.e. no genetic crossover etc)

Evolving strategies in games

- There are a huge number of different models and research papers that look at different:
  - Games
  - Ways the interactions are structured
  - Ways of reproducing new generations
- In many cases such systems are explored through simulation or analysis or both
- This area is called "evolutionary game theory"

Four models using PD

- We are going to look at four evolutionary agent-based models using the simple one-shot, two player, Prisoner's Dilemma game:
  - Random interaction (mean-field model)
  - Interaction on fixed lattice (spatial model)
  - Evolving interaction structure (tag model)
  - Evolving network (rewire model)

The Prisoner's Dilemma

Game is a strict PD when: T > R > P > S and 2R > T + S

Player 2 \ Player 1	C	D
C	(3) R R (3)	(5) T S (0)
D	(0) S T (5)	(1) P P (1)

Why is evolving cooperation in the PD interesting?

- Many disciplines use PD and evolution to gain insights into:
  - Cooperation and conflict in animal behaviour, Evolution of Life (Biology)
  - Creation and distribution of value (Economics)
  - Societies, institutions, power, conflict, collective action, "social contracts" (Sociology, Political science)*
  - Morality, rationality (Philosophy)**
  - How to program open distributed software (P2P)***

*** Robert Axelrod (1984) Evolution of Cooperation. Basic Books**

**** Matt Ridley (1996) The Origins of Virtue. Penguin Books**

***** Bram Cohen (2003) Incentives build robustness in Bittorrent. 1st Workshop on the Economics of Peer-2-Peer Systems.**

Mean-field interaction model

- A kind of “baseline” model for evolving strategies in games

A **mean-field** interaction model

- N agents initialised randomly to C or D
- Repeat some number of generations:
  - Game interaction: each agent is paired with another randomly chosen agent, plays PD
  - Reproduction: Generate new generation of N agents where fitness = average payoff
  - Mutate newly reproduced agents by flipping strategy with some small probability

Reproduction based on fitness

- How to generate the new population of N agents based on their fitness?
- Many possible ways so long as those with higher fitness have more chance of making copies of themselves into the next generation
- We will use a very simple method called **Tournament Selection**
- Another popular method is called **Roulette Wheel Selection**

Tournament Selection

- Many variants of tournament selection but here is a very simple variant:
- Repeat until next generation is full (N)
  - Select a random pair of agents (with replacement)
  - Reproduce the one with the highest fitness
  - Or a random one if both have the same fitness

Mean-field model

- ***What do you think will happen if we run this model?***

Mean-field model

- From any initial starting condition
- Evolution will quickly lead to Defect dominating the population and stay there
- This is called an **Evolutionary Stable Strategy** (or ESS)*
- A strategy is an ESS if a population all using it can resist “invasion” by other “mutant” strategies
- All ESS are Nash Equlibria (NE) but not all NE are ESS.
- Hence a link is found between game theory and evolutionary theory which biologists discovered and applied

***John Maynard Smith (1982) Evolution and the theory of games. Oxford University Press**

ESS and NE

- In a 2x2 game, strategy *s* is an ESS if for all other possible strategies *m*:

  u(s,s) > u(m,s) or

  u(s,s) = u(m,s) and u(s,m) > u(m,m)

- Whereas (s, s) is an NE if:

  u(s,s) >= u(m,s)

where u(a,b) means the payoff of strategy *a* playing against strategy *b*

Aside: repeated strategies

- More complex strategies can be evolved in this way and analysed to see if they are ESS
- Axelrod noted in his book that tit-for-tat was "collectively stable" (a **weak ESS**)
- The relationship between ESS, Nash and, say, Pareto efficiency is subtle and complex even in mean-field models
- However analysis (not just simulation) can be applied to repeated games with a limited number of strategies*

*** Nowak, Sigmund, Esam (1995) Automata, repeated games and noise. J. Math. Biol. 33: 703-722**

Evolution of strategies

- Knowing a strategy is an ESS for a given system does not tell us the dynamics (trajectories) that evolution could take from any given starting point
- Evolution could stabilise into states other than an ESS (extensions such as **Evolutionary Stable States** and **Evolutionary Stable Sets** have been developed)*
- In simple systems **replicator dynamics** equations can be used to predict trajectories (assuming no mutation and infinite N)**
- In general, simulation experiments are used to see what happens when it gets complex (to check, inform, supplement or replace analysis)

*** See Gyorgy Szabo, Gabor Fath (2006) Evolutionary games on graphs. http://arxiv.org/abs/cond-mat/0607344**

**There are modified forms of these equations that address some of these issues

Fixed Lattice (spatial PD) model

- Space is introduced into game interactions

Fixed lattice model (spatial PD)

- When interactions are non-random but structured in some way
- Then dynamics often get more complex
- An example of this is a classic model that situates agents on a 2D lattice (grid)
- Constraining their interaction *and* reproduction

Nowak & May (1992) Evolutionary games and spatial chaos. Nature 359, 826-829.
Nowak & May (1993) The Spatial Dilemmas of Evolution. Int. J. of Bifurcation and Chaos, Vol. 3, No. 1. 35-78

Lattice model

- Situate agents on a 2D lattice
- Randomly Initialise strategy of agent (C or D)
- Repeat some number of generations:
  - Interaction: each agent plays a PD game with its 8 neighbours (and itself)
  - Reproduction: each agent copies the strategy of the best performing (total payoff) agent in its neighbourhood (including itself)*

* Make a random selection if several have the same best payoff

Lattice model

- Using PD payoffs: $T = b$, $R = 1$, $P = 0$, $S = 0$
- Outcomes explored for different b values
- This is a so-called “weak PD” since $P=S$
- ***What happens when we run this model?***

Lattice model observations

- Different values of b give different dynamics:
  - Often dynamical patterns in which freq. of cooperators change over time
  - Various threshold values for b leading to different dynamical regimes
- Generally, from *most* starting conditions:

  $b < 1$ cooperators take over (no longer PD)

  $1 < b < 1.8$ between 0.7 and 0.95 cooperators

  $1.8 < b < 2$ around 0.3 cooperators (chaotic)

  $b > 2$ defectors take over (no longer PD)

Lattice model observations

- Through analysis it can be shown that when:

  $b < 1.8$ only C clusters can grow

  $b > 2$ only D clusters can grow

  $1.8 < b < 2$ both C and D clusters can grow

- the latter interesting region produces complex dynamics
- hard to capture analytically because depends on interactions between C and D clusters

Growth of a Defector in an infinite sea of Cooperators

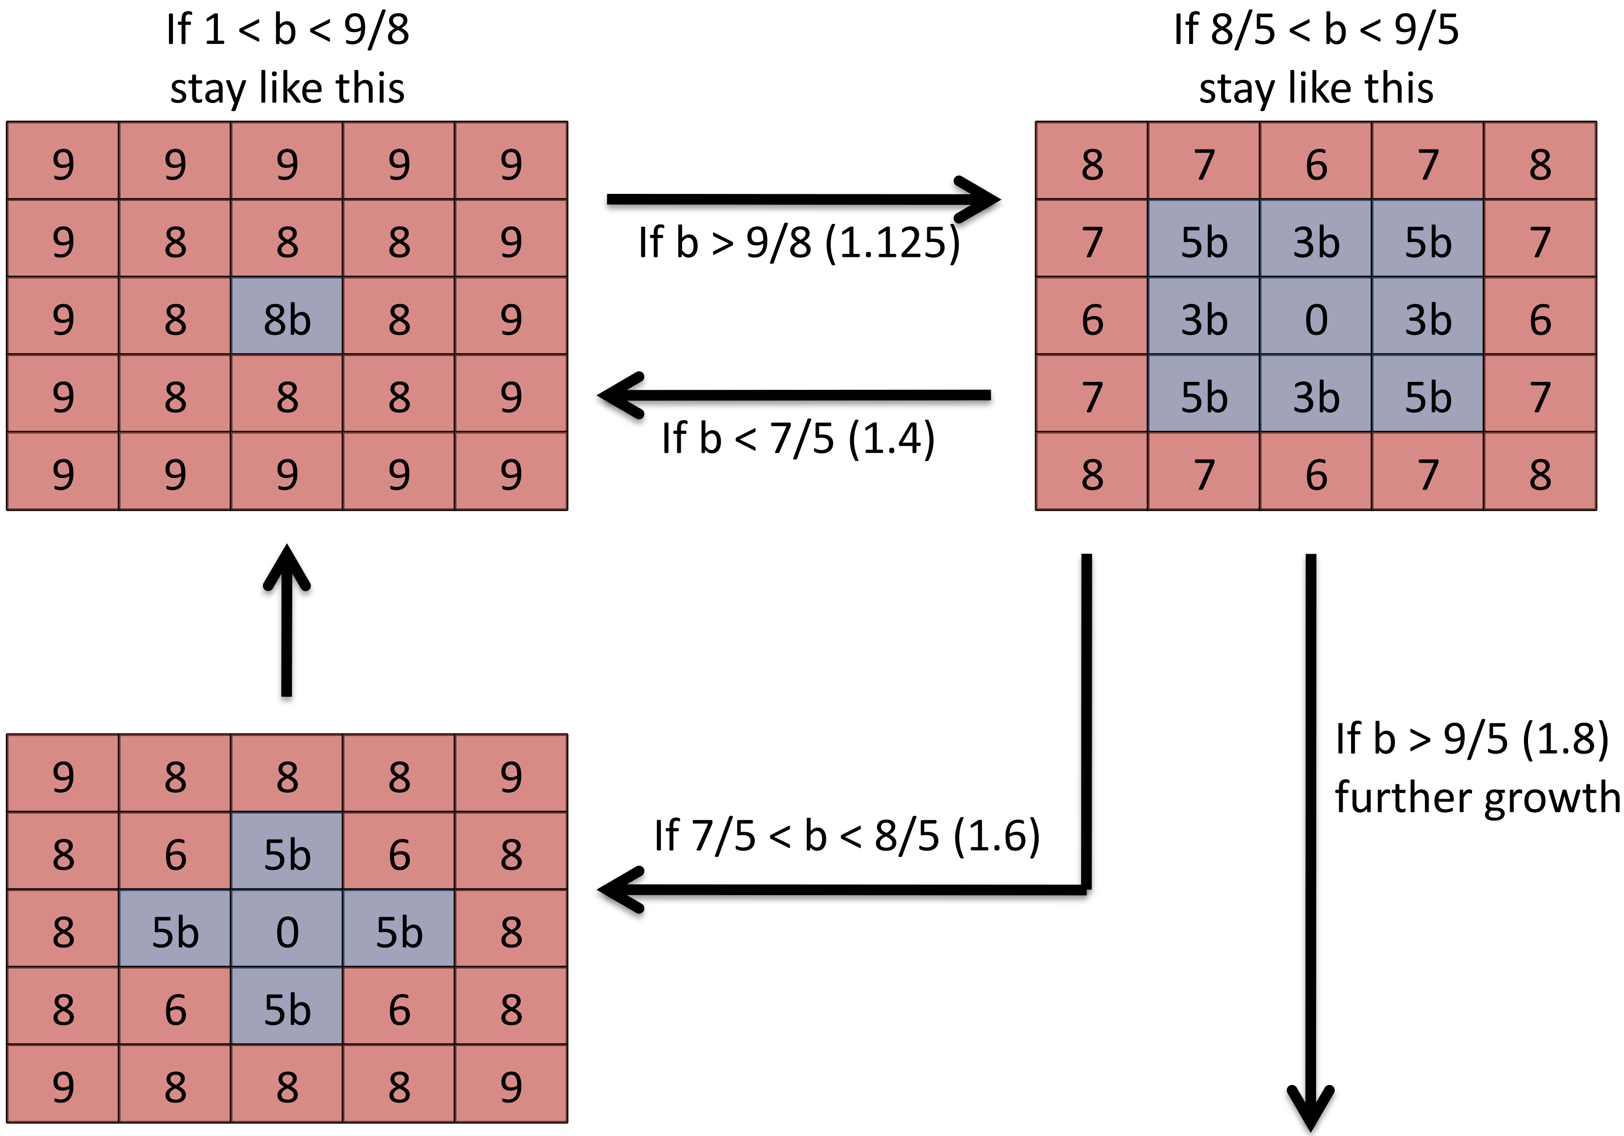

Growth of Cooperators in an infinite sea of Defectors

For $1 < b < 2$

0	0	0	0	0	0
0	b	2b	2b	b	0
0	2b	4	4	2b	0
0	2b	4	4	2b	0
0	b	2b	2b	b	0
0	0	0	0	0	0

b	2b	3b	3b	2b	b
2b	4	6	6	4	2b
3b	6	9	9	6	3b
3b	6	9	9	6	3b
2b	4	6	6	4	2b
b	2b	3b	3b	2b	b

4	6	6	6	6	4
6	9	9	6	9	6
6	9	9	9	9	6
6	9	9	9	9	6
6	9	9	9	9	6
4	6	6	6	6	4

Aside: other fixed graph topologies

- Several works examine cooperation evolving on different fixed graph topologies
- See overviews in Szabo & Fath and Allen & Nowak *
- A general idea termed “network reciprocity” proposes conditions linking topology measures to cooperation ($b/c > k$) **
- But overall general conditions are not “strong” and contingent on specifics

*** Gyorgy Szabo, Gabor Fath (2006) Evolutionary games on graphs.**
http://arxiv.org/abs/cond-mat/0607344

*** Allen, B. & Nowak, M. (2014) Games on graphs. EMS Surv. Math. Sci. 1. 113–151**

**** Ohtsuki, H., Hauert, C., Lieberman, E., Nowak, M. A., 2006. A simple rule for the evolution of cooperation on graphs and social networks. Nature 441, 502–505.**

Tag based interaction model

- Where the interaction structure is itself evolved

Evolving interaction structure

- So far considered static interaction structures
- What if we evolve the interaction structure?
- A simple method uses “tags”
- Each agent stores a strategy AND a tag
- The tag can be observed by other agents
- It can be thought of as a genetic marker (e.g. eye colour) or a label or recognisable social cue (e.g. accent, dress)

John Holland (1993) The effect of labels (tags) on social interactions. Technical Report Working Paper 93-10-064, Santa Fe Institute.

Agents – a tag and a PD Strategy

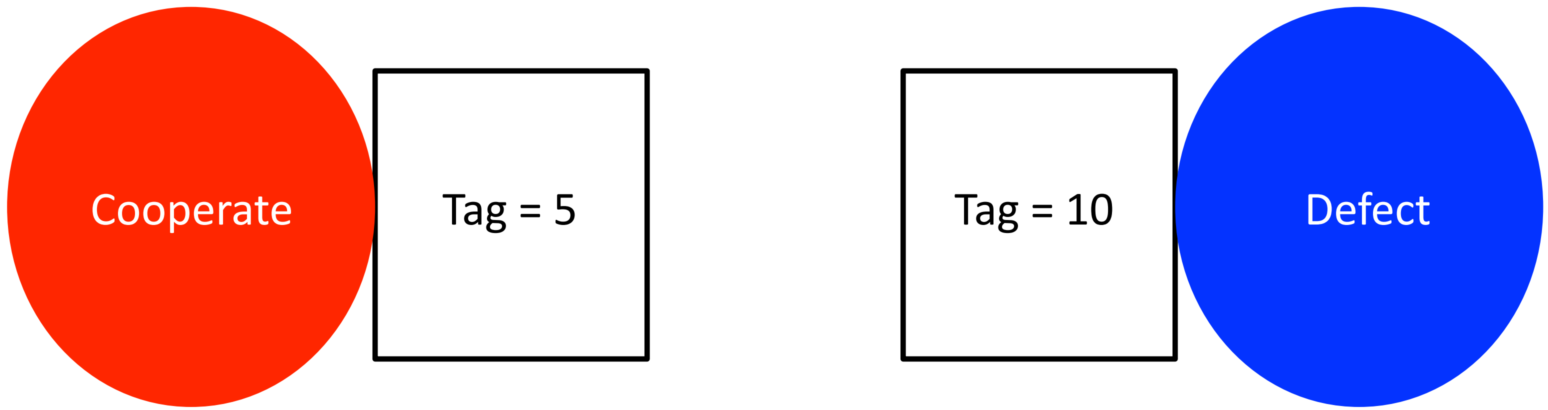

Tag can be represented by an integer, a bitstring, a colour a real number etc.

There should be many possible unique tags.

A **Tag** model

- Initialise agents with random strategies and tags
- Repeat some number of generations
  - Interaction: each agent plays a game of PD with a randomly selected other *with matching tags* (or random if no match exists)
  - Reproduction: generate next generation based on average payoff
  - Apply mutation to tag and strategy with small probability

Hales, D. (2000) Cooperation without space or memory: Tags, groups and the prisoner's dilemma. Multi-Agent-Based Simulation. Springer, Berlin

Tag model

- ***What happens if we run this model?***

Tag model observations

- Over a broad range of PD payoffs
- Very high C > 0.9 levels of cooperation quickly emerge from *any* initial starting condition
- The proportion of different tags in the population continually changes
- Those sharing the same tag can be thought of as dynamic interaction groups (tag groups)
- Too few tags leads breakdown of high C

Visualising the process

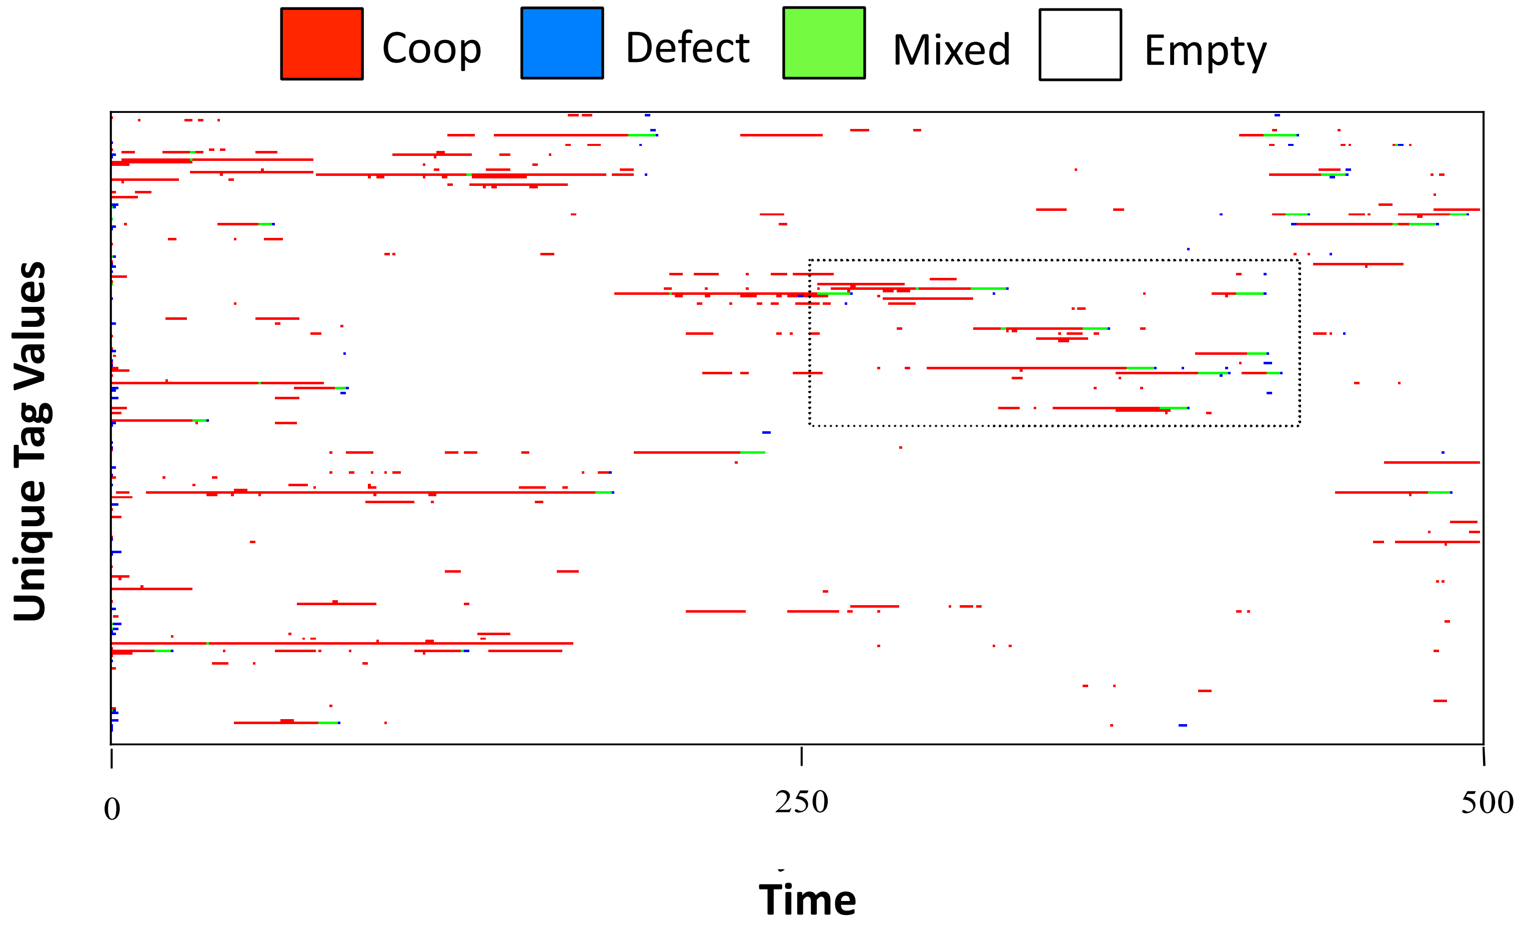

Visualising the process

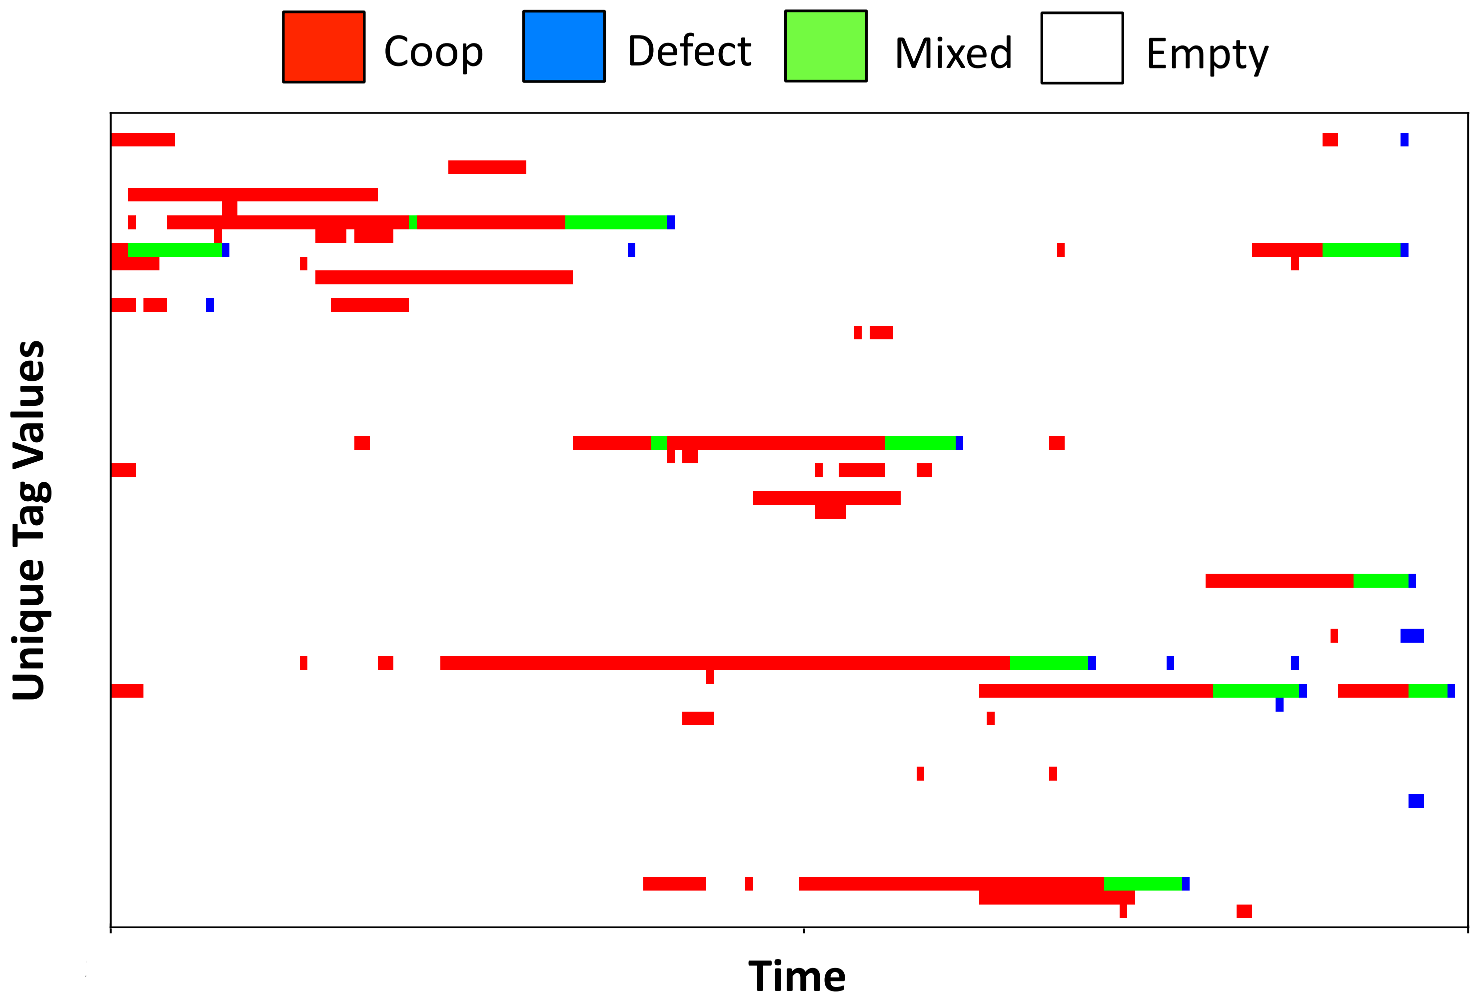

How tags work

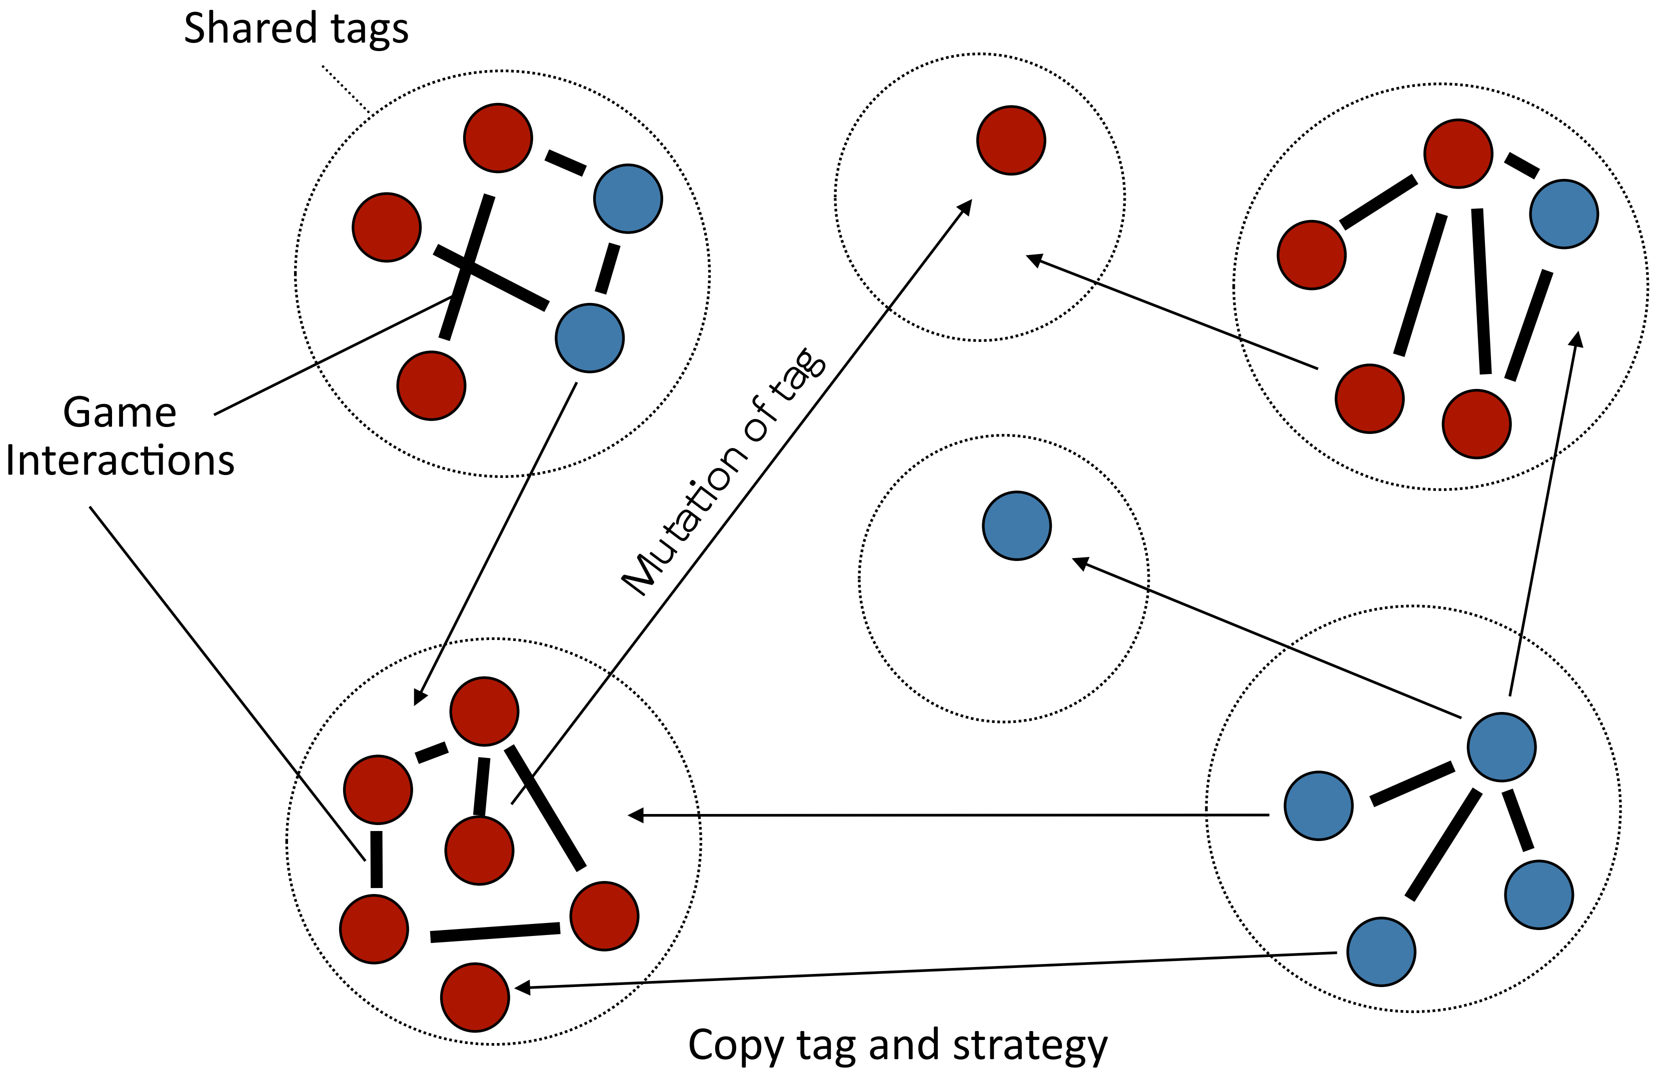

Tag models variants

- There are many variants of tag models:
  - Tags may effect the strategy played rather than the interaction structure only
  - tags may be combined with spatial interaction structures
  - Tags may be combined with a tolerance value allowing for similar tags to match

Riolo, R. L., Cohen, M. D., Axelrod, R., (2001) Evolution of cooperation without reciprocity. Nature 414, 441–443.

Jansen, V. A. A., Baalen, M., (2006) Altruism through beard chromodynamics. Nature 440, 663–666.

Traulsen, A., Nowak, M. A., (2007) Chromodynamics of cooperation in finite populations. PLoS ONE 2 (3), e270.

Network rewiring model

- Where the interaction structure is a network that evolves

Evolving interaction networks

- By evolving a network interaction structure in a similar way to tags..
- High levels of cooperation can emerge in dynamic evolving networks
- This equates to network rewiring
- Where nodes copy both strategies and links from each other

Hales, D. (2004) From Selfish Nodes to Cooperative Networks – Emergent Link-based Incentives in Peer-to-Peer Networks. The 4th IEEE Int. Conf. on Peer-to-Peer Computing. IEEE Computer Society Press

Santos, F. C., Pacheco, J. M., Lenaerts, T. (2006) Cooperation prevails when individuals adjust their social ties. PLoS. Comput. Biol. 2, 1284–1290

Network rewire model

Each node periodically (interaction):

plays PD with each of its network neighbours

Each node *p* periodically (reproduction):

q = select a random node

IF $\text{fitness}_q > \text{fitness}_p$ (where fitness = average payoff)

drop each link with probability d

link to node *q* and copy its strategy *and* links

mutate (with low probability) strategy and links*

*Mutate links = drop each link with probability d and connect to a randomly selected node

Network rewire model

- ***What happens if we run this model?***

rewire model observations

- For a broad range of PD payoffs:
  - High C > 0.9 emerges from *any* initial starting condition
  - Network rewires into a dynamic highly clustered topology
  - When d = 1 clusters are disconnected components
  - When d < 1 network forms a small-world topology
  - Similar process to tags

Agent-Based Modelling

- The models we looked at use very simple agents
- Many ABM use more complex agents that may learn and reason in some way
- Engage in more complex interactions than simple games where they have no simple payoff to maximise
- Such agents are often termed “boundedly rational” since they sit between ideal rationality and simple copying

Agent-based modelling

- NetLogo is a simple language / environment for experimenting with and learning ABM
- It is free and open source and comes with lots of built in ABM models and an integrated manual and tutorials
- The models I showed in this lecture were written in NetLogo and can be found at:
- www.davidhales.com/talks/szeged2016
- Module: www.davidhales.com/abm-netlogo
- Module: www.davidhales.com/msiis