Lecture 13: Intro to decision theory & game theory

Complex Systems 530
3/24/20

Credit: Many of these slides have been borrowed directly from lectures by Michael Hayashi and Lynette Shaw!

Outline

- Super basic & brief introduction to decision & game theory! But will give us some tools for exploring further in labs
- Discussion of plans for final presentations

How smart are your agents?

Reflexive Agents
simple, static rules

Goal-Based Agents
rules adjust according to being in goal state

Utility-Based Agents
rules attempt to maximize utility function(s)

Adaptive Agents
rules update based on experience

- ***Cognitive Complexity*** **+**

Game theory

- Game theory - motivated by the realization that the study of strategically interdependent behavior can be greatly enhanced via analysis of mathematical models of conflict and cooperation between “rational” decision-makers

- First got going as a field in 1940s per publication of Theory of Games and Economics work by von Neumann and Morgenstern

Introduction

- Applied to a wide range of areas
  - Social sciences (economics, sociology, political science)
  - Biology (genetics, species)
  - Computer science and logic
- Basic idea is that if we can conceptualize the interdependencies of individuals in a system as a game, will be able to “solve” for outcomes (for individuals all the way to population levels)

Decision theory

- Sort of a one-player version of game theory, where each person decides an action based on their preferences and the expected outcome of their actions (but no considering of other individuals/ players involved)

Decision theory

- **Actions**: The set of things an individual can do. e.g. video games, nap, run simulations.
- **Outcomes**: The results of each action
  - Video games → entertainment
  - Nap → rest
  - Run simulations → work
- **Preferences**: An ordering that specifies how an individual ranks the outcomes.
  entertainment > rest > work

Preferences

- For a preference order to be **rational**, it must be *complete* and *transitive*.

- **Complete**: for every pair of outcomes, one is preferred over the other (or they can be indifferent). Formally, for every a and b, $a > b$, $b < a$, or $a = b$. (One can consider strict $>$ or weak $\succcurlyeq$ preference)

- **Transitive**: For any three outcomes, a, b, and c, if a is preferred to b, and b is preferred to c, then a must be preferred to c. Formally, $a > b$ and $b > c$ implies $a > c$.

Rationality

- Completeness and transitivity guarantee that a person will be able to identify the best alternative out of their available options

- A rational actor in the economic sense always picks the most preferred alternative

- Note that a rational choice $\neq$ good choice!

Preferences

- Preferences are often described using a **utility function** or **payoff function**, which assigns a number/value to each outcome—the ordering is then assessed based on the utility function value

- Individuals then attempt to choose their actions to maximize their utility

Game theory

- Most decisions aren’t made in isolation – It’s important to know what somebody else might do. Game theory extends decision theory to problems where other people are a factor.

Game theory

- **Game**
  - Circumstances where results depends on the actions of 2 or more individuals (players)
  - Outcomes (payoff structures) are knowable and pre-defined
- **Players**
  - Possess choices (strategies) they can play
  - Seek to maximize their own utility/payoff (self-interest) and have the information and cognitive capacity to do so (rationality)
  - Typically everybody has common knowledge

Game theory

- **Players**: the actors making decisions
- **Strategies**: sets of choices specified for each player—these may or may not be the same across players!
- **Strategy profile**: a set containing one strategy chosen by each player
- **Payoffs**: a numerical representation of the costs and benefits of each strategy profile to each player

Common assumptions for games

- **Rationality**: Each player picks the action that gives the highest payoff given what they believe the other player might do. (Players always play best responses)

- **Complete information**: Each player knows the game, all of the payoffs, and all of the actions available to every player.3

- **Common knowledge:** Each player knows that the other players are rational and have complete information.

Common assumptions for games

- In practice, at least one of these assumptions is often violated. The basic theory shown here can be extended to deal with some deviations from rationality (bounded rationality, evolutionary game theory) and incomplete information (Bayesian games, and others).

Game variants

- Games come in a wide number of varieties:
  - Non-cooperative vs. cooperative
  - Zero-sum vs. non-zero-sum
  - One shot vs. Iterated
  - Symmetric vs. Non-symmetric
  - Simultaneous vs. Sequential (Normal vs. Extensive forms)
  - Two vs. Many player

Solutions to Games

- Anticipating the outcome of a games is often oriented toward analytically solving for the “stable” configuration of choices individuals can make

- Specifically, oriented toward identifying the **Nash equilibria** of a game:

  - Given that all players know each others’ equilibrium strategies, no player can benefit from changing their own strategy while the other players’ strategies remain unchanged

Prisoner's Dilemma

		Prisoner A	
		Silent	Betray
Prisoner B	Silent	-1 -1	0 -3
	Betray	-3 0	-2 -2

Prisoner's Dilemma

Nash equilibria?

Prisoner’s Dilemma

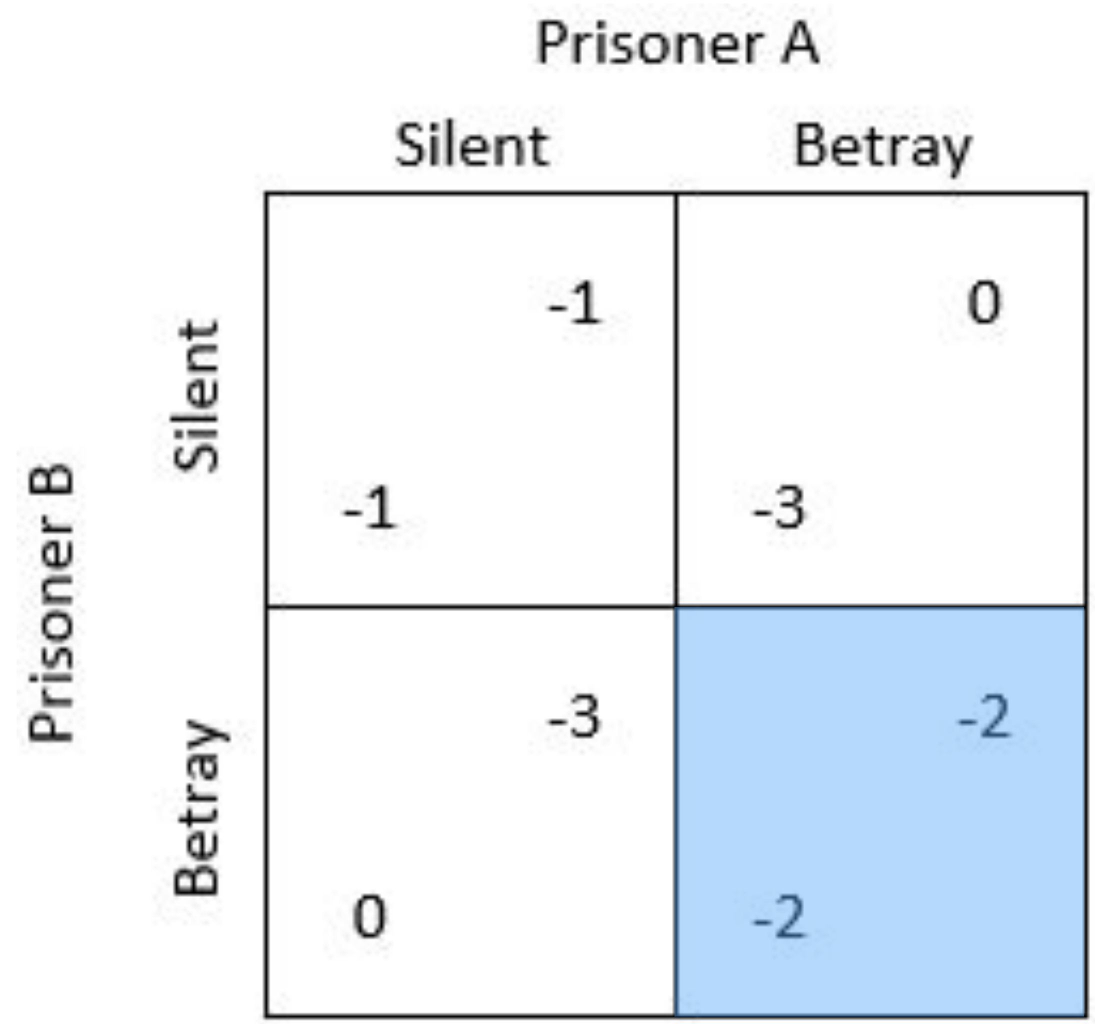

Prisoner’s Dilemma

Pure coordination game

	L	R
L	10,10	0,0
R	0,0	10,10

E.g.—suppose walking and don't want to bump into the person walking the other way—want to both swerve same direction (e.g. if both swerve to their own right, will miss each other, but if one swerves to their right and the other to their left they will bump into one another!)

Pure coordination game

	L	R
L	10,10	0,0
R	0,0	10,10

Multiple Nash equilibria

Other coordination games

Assurance game

	Party	Home
Party	10,10	0,0
Home	0,0	5,5

BoS

	Party	Home
Party	10,5	0,0
Home	0,0	5,10

Stag hunt

	Stag	Hare
Stag	10,10	0,5
Hare	5,0	5,5

(Or sometimes, 6,0 for the non-matching cases)

Mixed strategies: best response correspondences

Stag hunt

	Stag	Hare
Stag	10,10	0,5
Hare	5,0	5,5

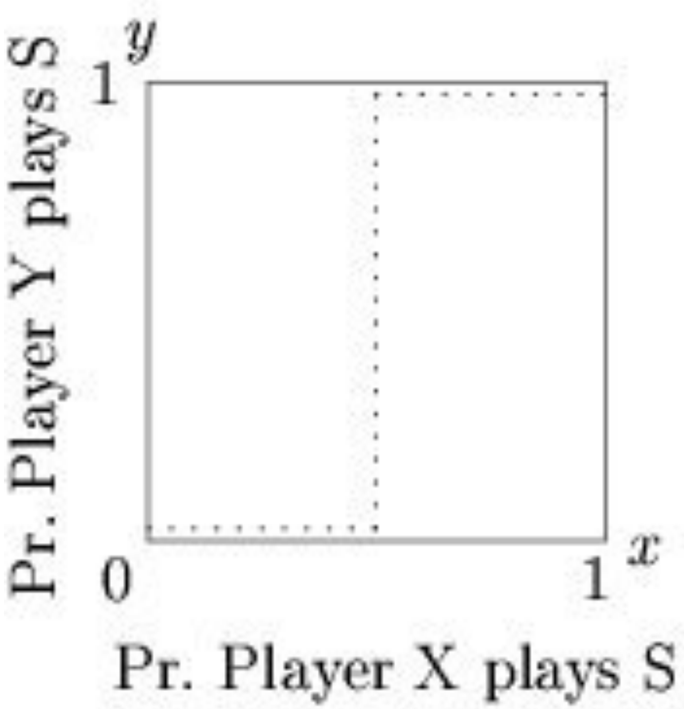

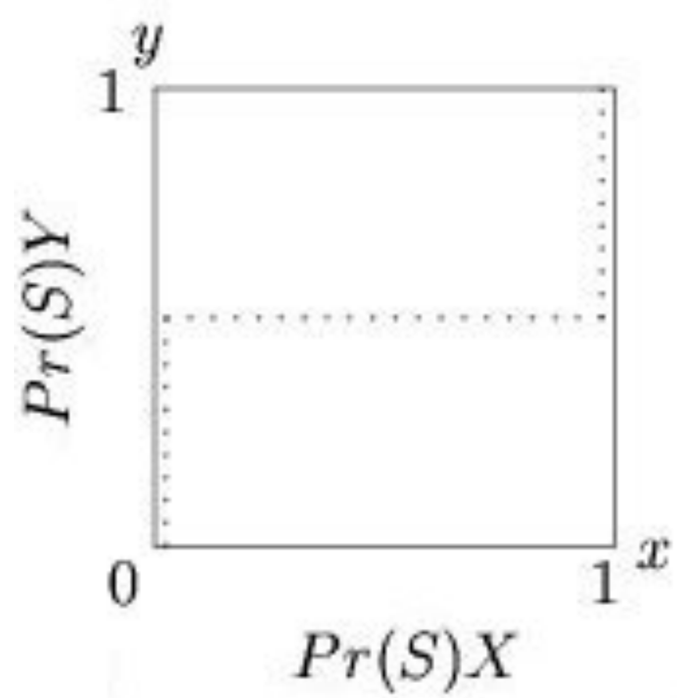

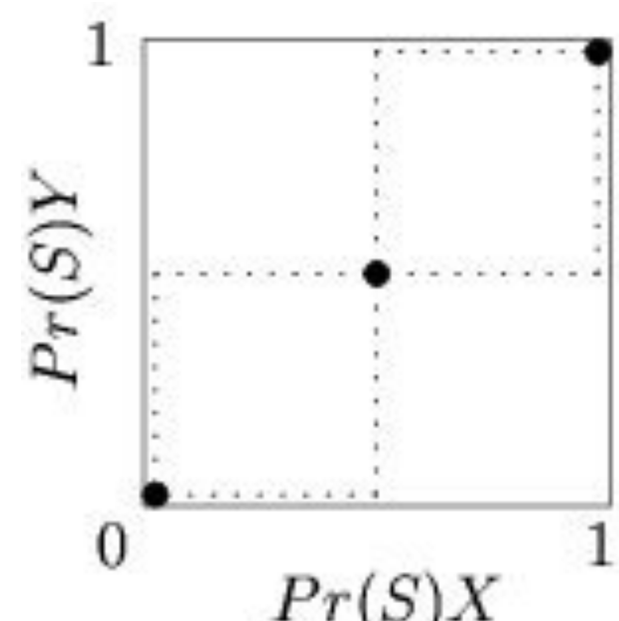

Matching pennies game

Even \ **Odd**	Heads	Tails
Heads	+1,-1	-1,+1
Tails	-1,+1	+1,-1

Zero-sum game
Nash equilibria?

Matching pennies game

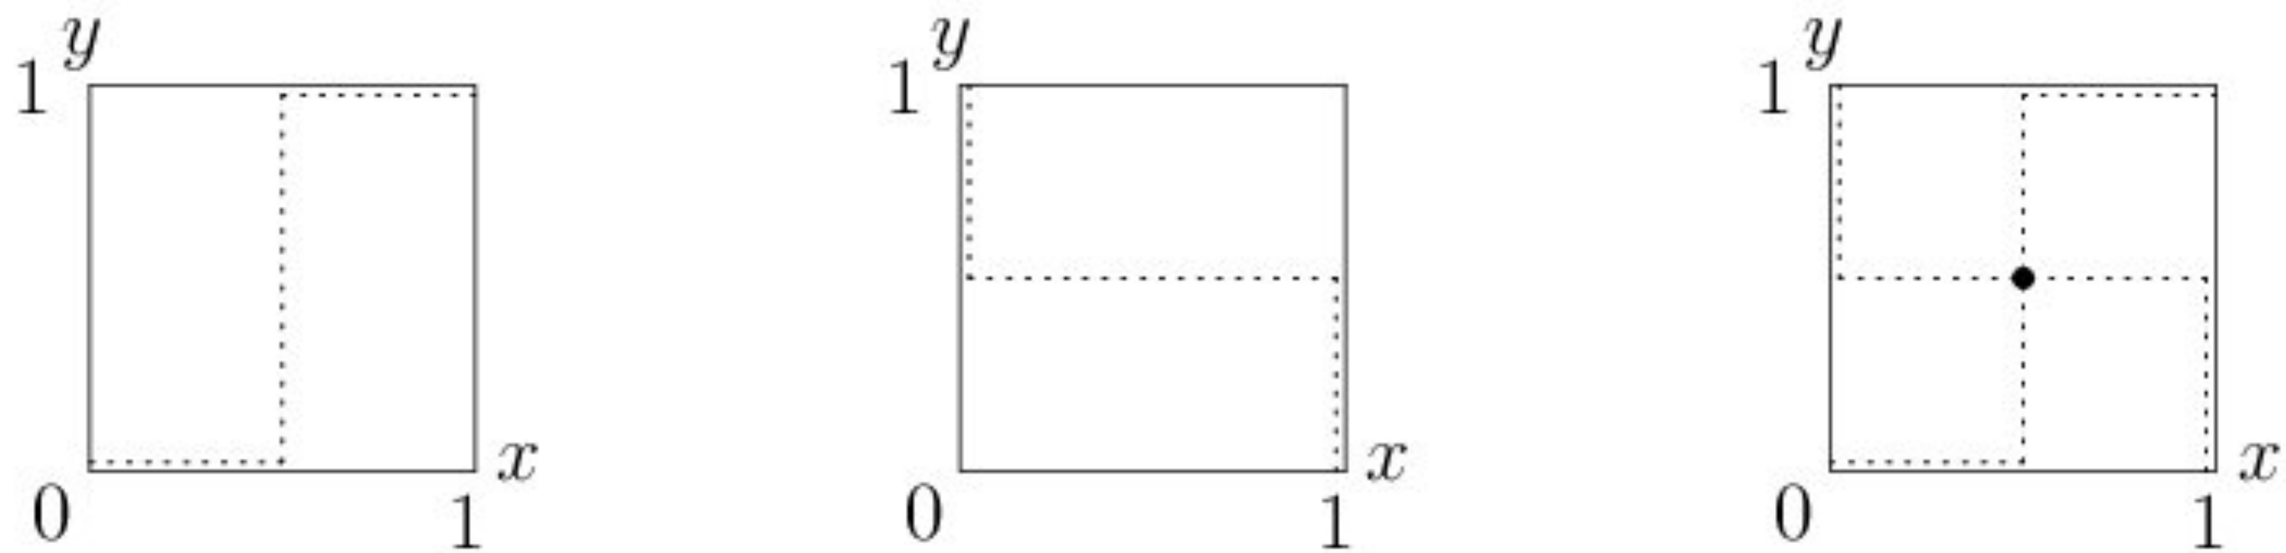

Challenges of games

- The need to remain analytically tractable makes it difficult to incorporate certain aspects of real world circumstances into games:
    - Temporal evolution of populations
    - Stochasticity
    - Space and interaction topology
    - Explorations of heterogeneity
    - Multiple or no equilibria situations

Challenges of games

- There have been many successful analytical approaches to tackling some of these issues (e.g. evolutionary game theory, etc.)

- Given concerns with things like **heterogeneity, space, interaction topology, simplistic actors, adaptation,** and **temporal dynamics**, seems like computational modeling may be useful in understanding these dynamics as well…

Game theory & ABMs

- Basic mapping:

Player Level

Game theory	ABM
Players →	*Agents*
Strategies and Actions →	*Agent rules*
Player Payoff →	*Agent variable*

Game Level

Game theory	ABM
Payoff Structure →	*Global variable called during agent interaction*
Iterations →	*Time Steps*
Number of Players →	*Number of agents involved in an interaction*

Game theory & ABMs

- Capture bounded rationality with agents using simple behavioral rulesets based only on local information
    - **Bounded rationality** - individuals have limited information, cognitive limitations, and finite time to make a decision
- Capture rudimentary “learning” through incorporation of agent memory in behavioral rules
- Introduce interaction topologies to determine who plays (interacts) with whom

Population dynamics & game theory

- Link agent payoffs to fitness and begin with a heterogeneous mix of agents imbued with different strategies

- Can use a tournament (i.e. multiple rounds of interaction) to assess robustness of different strategies or go further and link payoffs to reproduction in next rounds

- Investigate strategy evolution through allowing strategy “mutations” during reproduction (genetic algorithms, evolutionary game theory)

Population dynamics & evolutionary game theory

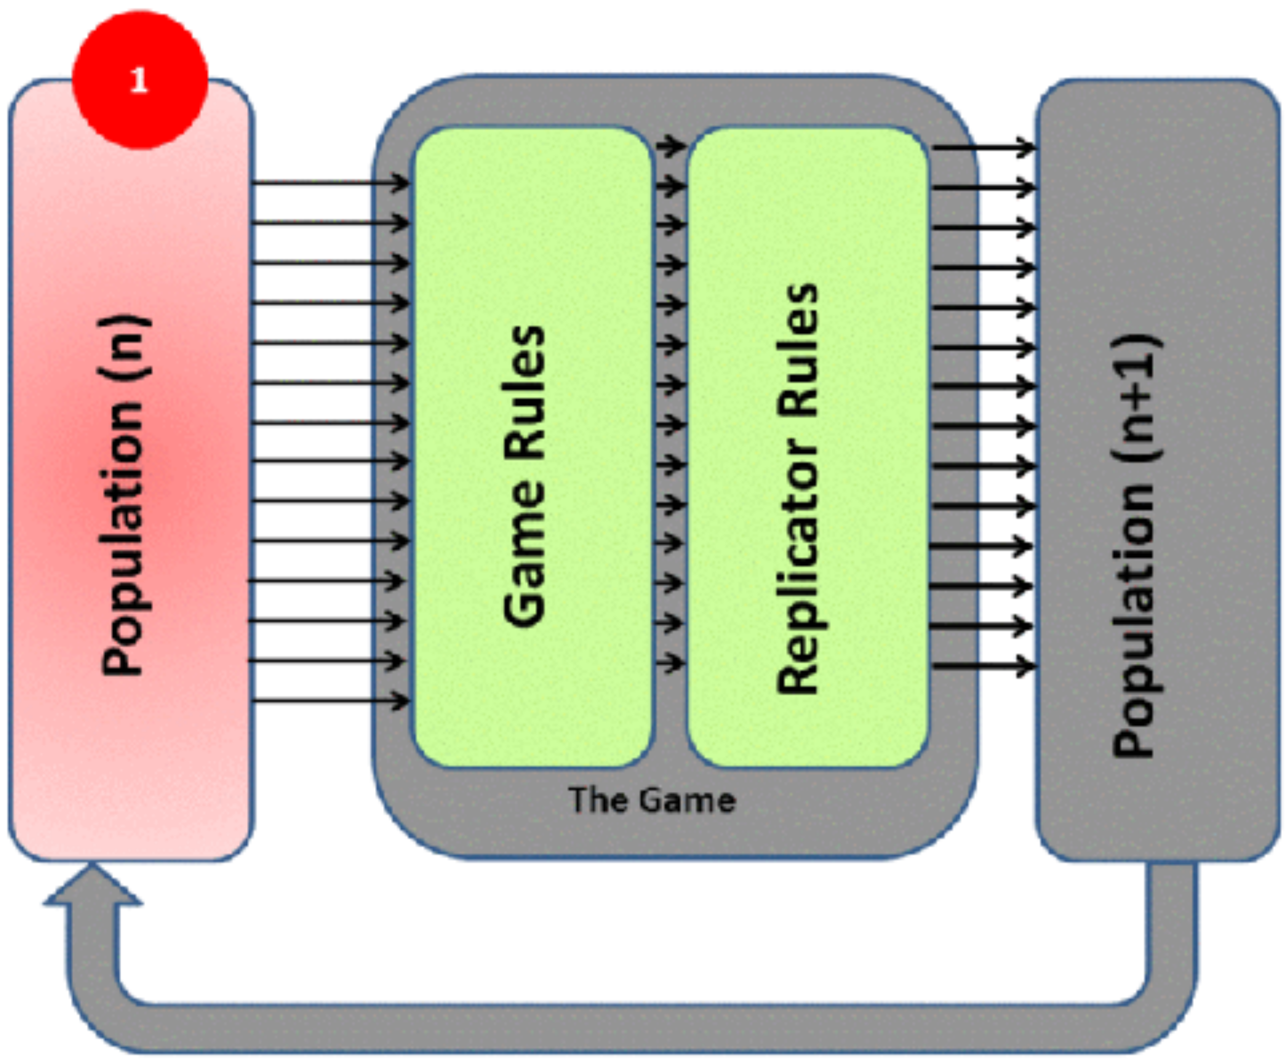

https://en.wikipedia.org/wiki/Evolutionary_game_theory

Evolutionary game theory

- An evolutionary game describes interactions at a single point in time.
- Evolutionary dynamics describe how traits change over time.
- Replicator equation: Every generation, suppose traits increase in prevalence proportional to the difference between their fitness and the average fitness in the population.

Population dynamics & evolutionary game theory

- Continuous time version:

$$\frac{df_i}{dt} = f_i[\phi_i(f) - \bar{\phi}(f)]$$

- where $\bar{\phi}(f)$ is the average fitness

- Wide range of approaches to looking at these issues (replicator-mutator, imitation dynamics, etc.)

- But we can also look at this with agents!

Population dynamics & games

- Note that these don't necessarily have to refer to evolution in a biological sense—evolutionary game theory and similar approaches are often used to understand many different systems, e.g.:
  - Infectious diseases (e.g. disease/behavior feedback loops—consider social distancing, vaccination, etc.)
  - Voting patterns, communication
  - And many other systems where behavior may change over time

Evolution of Cooperation

- Perhaps the most famous example of incorporating game theoretic model into an ABM context comes in classic study of the Evolution of Cooperation (Axelrod and Hamilton, 1981; Axelrod 1984)

- Begins with a persistent problem in both the social sciences and biology:

    - **How can cooperative behavior in groups arise and persist?**

Cooperative behavior and altruism

- Humans, many different animals (bats, etc.)
- However, cheating would often seem to gain higher payoff, why does cooperation and altruism persist?
- Historically, there has been a major debate on how individually costly behavior that benefits the group can arise and sustain within populations.
- Let's look at this with prisoner's dilemma

Cooperation and the Prisoner's Dilemma

- Being in a group of cooperators is good, but being a defector in a group of cooperators is even better.

- Holds true for biology as well: if payoffs are linked to reproduction, who will produce the most offspring?

		Prisoner A	
		Silent	Betray
Prisoner B	Silent	-1 -1	0 -3
	Betray	-3 0	-2 -2

Evolution of cooperation

- Axelrod’s Insight:
- In a one-shot PD game, “Always Defect” [All-D] always wins at both the individual and population levels (anything else can always be “invaded” by a newcomers playing of [All-D])
- In an iterated PD game with an uncertain time horizon and a basic ability to remember prior interactions, however, **other strategies may also be potentially stable**

Evolution of cooperation

- Are there simple strategies relying on simple memory that can allow cooperative group behavior to succeed in situations of on-going interaction?
- Success Criteria
  - Robustness: thrive in mixed population of strategies
  - Stability: once established can resist “invasion”
  - Initial viability: can establish in the midst of a lot of Defectors

Axelrod’s tournament

- **Agents**
  - Agents are assigned to play one of 14 extremely simple to somewhat more elaborate strategies submitted by a set of experts
  - Strategies also include [All D], [All C], and [Random]

Axelrod’s tournament

- **Model Setup**
  - Round-robin tournament of one-to-one matchups of all strategy pairs
  - Each matchup goes for 200 iterations (but agents don’t know that)
- **Model Outcome Assessment**: see which strategy had the highest average payoff across whole tournament

Axelrod’s tournament

- The Winner:
    - **Tit-for-Tat [TFT]**
- Even though extremely simple and involving only a very short memory, [TFT], that involves basic “nice” reciprocal cooperation, won out over everything else – including [ALL D]!

Let’s play!

- Netlogo iterated prisoner’s dilemma

Axelrod’s tournament (Round 2)

- **Agents**
  - 64 more strategies submitted from experts in a large number of fields (including Game Theory)
- **Model Setup**
  - Same round-robin tournament of one-to-one matchups of all strategy pairs
- Also looked at an “ecological” variant where populations for the next tournament were proportional to success in prior tournament (generated a time path of strategy distributions)

Axelrod’s tournament (Round 2)

- The Winner:
  - **Tit-for-Tat [TFT]** (Again)
- Here too, this basic strategy dominated both in terms of average success AND by completely taking over the population distribution in the “ecological variant”

Take-home messages

- Given a set of extremely plausible assumptions (like some initial clustering of cooperatively inclined individuals in a population), the basic principle of reciprocal cooperation can outperform an “All Defection” approach

- Without any appeals to “group selection,” can explain from “the bottom-up” emergence and persistence of cooperative behavior

- Given importance of bounded rationality, heterogeneity, and temporal evolution of populations in this analysis, very unlikely we could have gotten these results without availability of computational modeling

For next time…

- The evolution of trust: https://ncase.me/trust/
- The Evolution of Cooperation, Robert Axelrod; William D. Hamilton. Science, New Series, Vol. 211, No. 4489. (Mar. 27, 1981), pp. 1390-1396.