# Evolution of Cooperation

Models and Simulations in Philosophy
November 11th, 2013

## Review

**Last Week:**

- Classical Theory of Rationality in Economics
- Bounded Rationality
- Networks and their Properties
- Boundedly Rational Learning in Networks

## Today

Today: Can the concepts of bounded rationality and networks help to explain how cooperative behavior might evolve?

First, let's see why this question is sometimes thought to be difficult from a classical perspective in economics . . .

## Outline

## Three Decision Rules

Name three decision rules:

- Maximize (subjective) expected utility (SEU)
- Dominance
- Minimax

## Dominance

Recall the following two observations about dominance ...

## Dominant Actions Maximize Utility

**Observation 1:**

**Theorem**

*Suppose a is a dominant action. Then a is a minimax action and also maximizes subjective expected utility.*

## Dominant Actions Maximize Utility

**Observation 2:** Dominance is a well-defined decision rule even if

- Agents' preferences are only qualitative, in the sense that they are not represented by numerical utilities.
- Agents' assessment of likelihood are only qualitative, in the sense that they are not represented by numerical probabilities.

## Dominance in One Shot Prisoners' Dilemmas

A classical prisoner's dilemma has the following structure:

| | Cooperate | Defect |
|---|---|---|
| Cooperate | 2,2 | 0,3 |
| Defect | 3,0 | 1,1 |

From the standpoint of row player, is there a dominant action? What about column player?

## Dominance in One Shot Prisoners' Dilemmas

- In a Prisoner's Dilemma, the dominant action is to defect.
- By the previous, it is also an SEU maximizing action.
- So according to the classical economic view, rational actors will defect in a prisoner's dilemma.
- By the second remark, rational actors will defect **regardless** of the
  - Numerical payoffs in the outcomes and
  - Likelihood that their opponent employs a particular strategy.

## Explaining Cooperation with Classical Economics

Using classical economics, it is difficult, therefore, to explain the existence of cooperation . . .

## Explaining Cooperation with Population Models

Perhaps we can find cooperative behavior if agents play prisoner's dilemmas repeatedly in a large population.

Let's use the only population model we've discussed thus far: the replicator dynamics.

## Replicator Dynamics: One-Shot Prisoners' Dilemmas

Suppose we have a prisoner's dilemma with the following payoffs.

| | Cooperate | Defect |
|---|---|---|
| Cooperate | 2,2 | 0,3 |
| Defect | 3,0 | 1,1 |

Suppose $p$ percent of the population is cooperating, and $1-p$ are defecting.

What are the fitnesses of cooperating and defecting respectively?

## Replicator Dynamics: One-Shot Prisoners' Dilemmas

- Let $p$ denote the proportion of the population that cooperates, and $1-p$ that which defects.
- Then, as the conspecifics encounters other random members of the population, the fitness (i.e. expected utility) of cooperate and defect respectively are:

$$\begin{aligned} F(\mathit{Cooperate}) &= p \cdot 2 + (1-p) \cdot 0 \\ F(\mathit{Defect}) &= p \cdot 3 + (1-p) \cdot 1 \end{aligned}$$

## Replicator Dynamics in Prisoner's Dilemmas

So which has the higher fitness?

Just as dominant actions have higher expected utility regardless of one's subjective probabilities, so do dominant actions have higher fitness regardless of the current composition of the population.

## Replicator Dynamics

- Let $F(\textsc{ave})$ denote the average fitness of all phenotypes in the population. In this case,

$$F(\textsc{ave}) = p \cdot F(\mathit{Cooperate}) + (1-p) \cdot F(\mathit{Defect})$$

- In large populations, after one round of play the actual number of offspring for each phenotype will (with high probability) be close to the expected value, i.e. to the fitness of the phenotype.

## Replicator Dynamics

In large populations, therefore, one can show that proportion of Defect(ing) players change as follows:

$$p' - p = p(F(\mathit{Defect}) - F(\textsc{ave}))$$

where $p'$ is the proportion of defectors after one round of play. This equation is called the replicator dynamics.

**Question:** What happens to the proportions of defectors and cooperators in the population?

## Replicator Dynamics

**Moral:** The replicator dynamics predicts that defectors overtake the population eventually; how long it takes depends upon the payoffs.

Let me show you some simulations . . .

## Replicator Dynamics

Defection spreads because

- interactions among agents are random and take place on a global scale.
- the population is large.

## PDs on Networks

**Idea:** Perhaps cooperation can evolve if agents interact locally?

## Networks

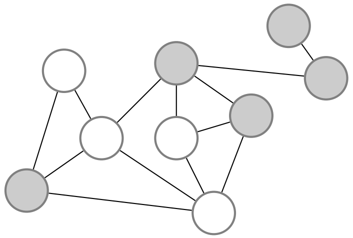

Nodes = Agents
Edges = Indicate which agents "interact"
Colors = Current strategy of the agent

## Four Types of Networks

Name four "types" of networks discussed by [Alexander, 2007]:

1. Lattice
2. Bounded degree
3. Small worlds
4. Dynamic - This really isn't specific, as a network might be dynamic in many different ways.
    - We'll talk more about dynamic networks at the end of class.

## Lattice Networks

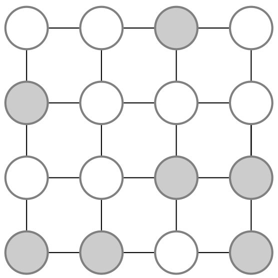

A Lattice Network

## Lattice Networks

Lattice networks

- Were some of the first studied in ABMs, likely because they are easy to program
- Exhibit a number of formal properties (e.g. regularity) that are uncommon in social networks.
- Nonetheless, provide an easy starting point to experiment.

## Common Features of Social Networks

**Question:** What happens when rational agents play a one-shot PD on lattice networks?

**Answer:** Defection spreads because agents are rational, and rational agents immediately choose the dominant action.

## Common Features of Social Networks

**Question:** What happens when agents employ a boundedly rational strategy - say "Imitate the Best Average" - in a PD on a lattice network?

**Answer:** It depends upon the payoff structure. Let's run some simulations.

## Common Features of Social Networks

**Question:** What happens in more complex networks?

**Answer:** It depends upon the payoff structure, the learning rule, and the network structure. Let's run some simulations.

## Repeated PDs

Perhaps we short-changed classical economics' ability to explain cooperation, as we only considered one-shot prisoner's dilemmas.

## Repeated PDs

What happens if a prisoner's dilemma is repeated?

For concreteness, let's assume its repeated five times, and the total payoff to a player is the **sum** of his payoffs of each play.

## Repeated PDs

- The strategy space is now much larger for players.
- One strategy is to defect all the time; one is to cooperate always.
- But a player's actions may also depend upon previous moves by his opponent. E.g.,
  - GRIM: Cooperate until one's opponent does not. Defect on every subsequent stage.

## Repeated PDs

Is always defecting the dominant strategy in a repeated prisoner's dilemma?

Generally not.

## Repeated PDs

Suppose the PD is repeated five times:

| | Cooperate | Defect |
|---|---|---|
| Cooperate | 2,2 | 0,3 |
| Defect | 3,0 | 1,1 |

Always cooperate vs. GRIM $\Rightarrow 5 \cdot 2 = 10$
Always defect vs. GRIM $\Rightarrow 3 + (4 \cdot 1) = 7$

## Repeated PDs

So always defecting is not a dominant action in some repeated PDs!

Have we saved ourselves from a pessimistic conclusion about rational actors in PDs?

Not so fast.

## Iterated Elimination of Dominated Strategies

| | Left | Center | Right |
|---|---|---|---|
| Top | 0,2 | 3,1 | 2,3 |
| Middle | 1,4 | 2,2 | 4,1 |
| Bottom | 2,1 | 4,4 | 3,2 |

Is there any action that is dominated? (Hint: Look at row player first, whose payoffs are on the left).

## Iterated Elimination of Dominated Strategies

| | Left | Center | Right |
|---|---|---|---|
| Top | 0,2 | 3,1 | 2,3 |
| Middle | 1,4 | 2,2 | 4,1 |
| Bottom | 2,1 | 4,4 | 3,2 |

For Row, Bottom dominates Top. So if Row is rational, then Row won't choose Top.

Suppose Column knows Row is rational. What outcomes will Column consider?

## Iterated Elimination of Dominated Strategies

| | Left | Center | Right |
|---|---|---|---|
| Middle | 1,4 | 2,2 | 4,1 |
| Bottom | 2,1 | 4,4 | 3,2 |

Column knows Row won't play Top. So she considers the above game matrix.

Are there any actions that are dominated from Column's perspective?

Yes. Center dominates Right. So if Column is rational, she won't play Right.

## Iterated Elimination of Dominated Strategies

| | Left | Center | Right |
|---|---|---|---|
| Middle | 1,4 | 2,2 | 4,1 |
| Bottom | 2,1 | 4,4 | 3,2 |

- Suppose Row knows that Column is rational.
- **And** Row knows that Column knows that Row is rational.
- Then what outcomes will Row consider?

## Iterated Elimination of Dominated Strategies

| | Left | Center |
|---|---|---|
| Middle | 1,4 | 2,2 |
| Bottom | 2,1 | 4,4 |

Row's game matrix now looks like this.
Repeating this reasoning we get . . .

## Iterated Elimination of Dominated Strategies

| | Center | Right |
|---|---|---|
| Bottom | 4,4 | 3,2 |

As Bottom dominates Middle.

## Iterated Elimination of Dominated Strategies

| | Center |
|---|---|
| Bottom | 4,4 |

As Center dominates Right.

## Iterated Elimination of Dominated Strategies

Moral: In a game with rational players who knew each other to be rational, contestants will not choose strategies that can be eliminated by considerations of dominance in this manner.

## Iterated Elimination of Dominated Strategies

Let's apply this reasoning to a repeated PD.

## Backwards Induction in PDs

Suppose a PD is repeated five times.

- Take any strategy you like (E.g. always cooperate, GRIM, etc.). Call it $s$.
- Suppose $s$ cooperates in round five.
- Define a strategy $s^*$ that is just like $s$, except that $s$ defects on the fifth stage, regardless of what has happened previously in the game.
- What is the relationship between $s$ and $s^*$?

## Backwards Induction in PDs

I claim $s^*$ dominates $s$.

- Suppose your opponent plays strategy $t$.
- Then $s$ and $s^*$ earn the same payoffs up to stage five versus $t$.
- So the difference between $s$ and $s^*$ comes down to the last stage.
- Remember: defecting is dominant in a one shot game.
- So regardless of the opponent's strategy $t$, the strategy $s^*$ will have better outcomes than $s$ on the last stage.

## Backwards Induction in PDs

- In other words, **regardless** of the strategy $t$ employed by one's opponent, $s^*$ is at least as good as $s$ on the first four stages, and it is strictly better on the last.
- So $s^*$ dominates $s$.

## Backwards Induction in PDs

- Rational players, who know each other to be rational, will defect on the last stage of a repeated prisoner's dilemma.
- What about the second to last stage?

## Backwards Induction in PDs

Repeat the same reasoning.

- Let $s$ be any strategy that defects on the last stage.
- Suppose $s$ cooperates on the second to last stage.
- Define $s^*$ to be just like $s$, except that $s^*$ defects on the second to last stage.
- By the same reasoning as before, $s^*$ dominates $s$ against strategies that defect on the last stage.

## Backwards Induction in PDs

Moral 1: In a repeated PD, the only strategy that survives the repeated elimination of dominated strategies is to defect always.

Again, this argument did not depend upon agents making judgments of probability.

It also does not depend upon payoffs being numerical, but I don't want to state the assumption that is necessary . . .

## Iterated Elimination of Dominated Strategies

Moral 2: If agents are rational in the classical sense, it seems hard to explain how cooperation might emerge in repeated prisoners' dilemmas.

## Repeated PDs and Replicator Dynamics

Does cooperative survive in the replicator dynamics if each agents plays a **repeated** prisoners' dilemma with a random member of the population?

## Repeated PDs and Replicator Dynamics

Inspired by Axelrod [2006]'s PD tournament, [Alexander, 2007]

- Randomly assigns each agent in a large population a strategy for a repeated PD.
- Let's the population evolve according to the replicator dynamics

## Repeated PDs and Replicator Dynamics

**Result:** Strategies that sometimes cooperate and sometimes defect were the ones left after many stages of evolution.

## Repeated PDs on Networks

**Question:** According to Alexander, what happens if, on each stage of evolution, agents play repeated prisoners' dilemmas on the various types of networks?

It's a trick question. [Alexander, 2007] develops no models of this sort and runs no simulations.

## Purpose of Models

**Question:** What does all this tell us about cooperation, especially if the models give different results?

In the last few classes, we'll talk about the purposes of modeling, the pitfalls, the advantages, and the disadvantages.

## Purpose of Models

**Answers for Today:**

- How possible stories vs. How so
  - Given the problems with classical economic explanations, we are often just interested in explaining how it is **possible** that cooperation evolved.
- Provides motivation and framework for particular empirical investigations:
  - Many social scientists have characterized properties of real social networks.
  - Biologists can sometimes quantify the energy spent by organisms in acting; that is, they can measure the payoff structure.
  - Both biologists and social scientists study learning rules employed by organisms.

## How Possible

**Question:** If we were just interested in "how possible" stories for the evolution of cooperation, then why consider so many models? Isn't one sufficient?

## Robustness

**Potential Answer:** Robustness.

- "How possible" stories are not convincing if they are fragile, i.e., if slight changes to the model cause drastic changes in behavior.
- If many different models behave similarly, however, then "how possible" explanations become more convincing. Such behavior is said to be robust.
- Different models are more-or-less realistic in different ways and so may provide different reasons to believe a "how-possible" story.

## Robustness

For discussions of robustness, see [Muldoon, 2007] and [Parker, 2011]; the former defends the value of robust models and the latter questions it.

## Topics

Topics we'll discuss today:

- World Commands
- Agents: Turtles, Patches, and Links
- Agent Sets

## References I

Alexander, J. M. (2007). *The structural evolution of morality*. Cambridge University Press Cambridge.

Axelrod, R. (2006). The evolution of cooperation: revised edition.

Muldoon, R. (2007). Robust simulations. *Philosophy of Science*, 74(5):873–883.

Parker, W. S. (2011). When climate models agree: The significance of robust model predictions. *Philosophy of Science*, 78(4):579–600.