Evolution of Cooperation:
Repeated Prisoner's Dilemmas and Network interactions

Models and Simulations in Philosophy
April 30th, 2013

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 emerge in communities?

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

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

Dominance

	Sun	Rain
Frisbee	5	-1
Biergarten	4	-2

	Sun	Rain
Frisbee	5	-1
Biergarten	4	-2

Dominance: Playing Frisbee is the dominant action: its payoffs are better in every state of the world than the corresponding ones for going to the Biergarten.

	Sun	Rain
Read	2	3
Biergarten	4	-2

	Sun	Rain
Read	2	3
Biergarten	4	-2

Minimax: Reading is the minimax action. It's worst-case payoff (2) is better than the worst-case payoff of going to the Biergarten (-2).

Suppose you fully believe the weather forecast, which claims the chance of rain is .5%.

	Sun	Rain
Read	2	3
Biergarten	4	-3

Expected Utility Maximization

Suppose you fully believe the weather forecast, which claims the chance of rain is .5%.

	Sun	Rain
Read	2	3
Biergarten	4	-3

Maximized Expected Utility: The subjective expected utility of going to the Biergarten is higher than that of reading.

$$\begin{aligned} \text{SEU}(\textit{Biergarten}) &= .995 \cdot 4 - .005 \cdot 3 = 3.965 \\ \text{SEU}(\textit{Read}) &= .995 \cdot 2 + .005 * 3 = 2.005 \end{aligned}$$

Three Decision Rules

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

Here are two simple observations about dominance:

Theorem

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

Dominant actions maximize expected utility:

	Sun	Rain
Frisbee	5	-1
Biergarten	4	-2

Dominant actions maximize expected utility:

	Sun	Rain
Frisbee	5	-1
Biergarten	4	-2

Suppose one believes the probability of rain is p. Then:

$$\begin{aligned} \textsc{seu}(\textit{Frisbee}) &= (5 \cdot p) + (-1 \cdot (1 - p)) \\ \textsc{seu}(\textit{Biergarten}) &= (4 \cdot p) + (-4 \cdot (1 - p)) \end{aligned}$$

Each term in the sum of Frisbee is bigger than the corresponding term for Biergarten

Here are two simple observations about dominance:

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

Three Decision Rules

Here are two simple observations about dominance:

1. **Fact:** Suppose a is a dominant action. Then a is a minimax action and also maximizes subjective expected utility.
2. Dominance is a well-defined decision rule even if
   - One does not assign states of the world probabilities; in fact, dominance does not even require qualitative comparison of the likelihood of outcomes.
   - One does not assign outcomes numerical payoffs; the decision rule makes sense even if outcomes can only be qualitatively compared.

Outline

1. Review
   - Three Decision Rules

Outline

1. Review
   - Three Decision Rules
2. Prisoners' Dilemmas
   - One Shot
   - Repeated
     - Iterated Elimination of Dominated Strategies
     - Backwards Induction in PDs

Outline

1. Review
   - Three Decision Rules
2. Prisoners' Dilemmas
   - One Shot
   - Repeated
     - Iterated Elimination of Dominated Strategies
     - Backwards Induction in PDs
3. PDs on Networks

Dominance in Game Theory

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

- In game theory, one player's strategies are his opponent's states of the world.

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

- In game theory, one player's strategies are his opponent's states of the world.
- So dominance says that, if the outcome of employing a_1 is better than that of a_2 for each possible strategy employed by one's opponent, then one should not choose a_2.

A classical prisoner's dilemma has the following structure:

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

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?

- In a Prisoner's Dilemma, the dominant action is to defect.

- In a Prisoner's Dilemma, the dominant action is to defect.
- By the first remark, it is also an SEU maximizing action.

Dominance in One Shot Prisoners' Dilemmas

- In a Prisoner's Dilemma, the dominant action is to defect.
- By the first remark, it is also an SEU maximizing action.
- So according to the classical economic view, rational actors will defect in a prisoner's dilemma.

- In a Prisoner's Dilemma, the dominant action is to defect.
- By the first remark, 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.

What happens if a prisoner's dilemma is repeated?

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.

- The strategy space is now much larger for players.

Repeated PDs

- The strategy space is now much larger for players.
- One strategy is to defect all the time; one is to cooperate always.

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.

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

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

Generally not.

Suppose the PD is repeated five times:

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

Always cooperate vs. GRIM $\Rightarrow$

Suppose the PD is repeated five times:

	Cooperate	Defect
Cooperate	2,2	0,3
Defect	3,0	1,1

Always cooperate vs. GRIM $\Rightarrow 5 \cdot 2 = 10$

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 $\Rightarrow$

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 $\Rightarrow 3 + (4 \cdot 1) = 7$

So always defecting is not a dominant action in some 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?

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.

	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).

	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 Bottom.

	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 Bottom.

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

	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.

	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?

	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.

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

- Suppose Row knows that Column is rational.

	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.

	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?

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

Row's game matrix now looks like this.

	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 . . .

	Center	Right
Bottom	4,4	3,2

As Bottom dominates Middle.

	Center
Bottom	4,4

As Center dominates Right.

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.

Let's apply this reasoning to a repeated PD.

Suppose a PD is repeated five times.

- Take any strategy you like (E.g. always cooperate, GRIM, etc.). Call it s.

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.

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.

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^*?

I claim s^* dominates s.

- Suppose your opponent plays strategy t.

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.

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.

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.

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.

- 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.

- 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.

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?

Repeat the same reasoning.

- Let s be any strategy that defects on the 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.

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.

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.

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

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 ...

Moral 2: If agents are rational in the classical sense, it seems hard to explain how cooperation might emerge in prisoner's dilemma like games.

Outline

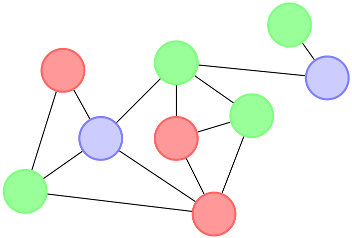

Nodes = Agents
Edges = Indicate which agents "interact"
Colors = "Type" of Agent

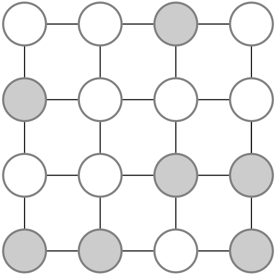

A Lattice Network

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.

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.

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.

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

Answers:

- How possible stories vs. How so

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

Answers:

- How possible stories vs. How so
- 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.

Topics

Topics we'll discuss today:

- Global vs. Local Variables
- If-then statements
- Loops

References I