

3. Evolution of Cooperation

Modelling Social Interaction in Information Systems (MSIIS) www.davidhales.com/msiis

David Hales, University of Szeged
dave@davidhales.com

Cooperation

- Cooperation is a broad range of phenomena
- It is social because involves more than one agent working together for mutual benefit
- Seen widely in human and animal societies
- It could be argued that many kinds of social interaction support the function of cooperation

Cooperation

- Biologists are interested in cooperation:
 - How animals come to help each other (cleaning fish)
 - How single-cell organisms form societies (slime molds)
 - Other biological phenomena (multi-cell organisms?)
- Social scientists are interested in cooperation:
 - How people form collective institutions
 - How people often spontaneously help each other
 - Other social phenomena (morality – Matt Ridley 1997)

Ridley, M (1997) *The Origins of Virtue*. Penguin Books.

“Cooperative” Cooperation

- In human and information systems, many ways of producing cooperation:
 - Shared goals and plans
 - Language, contracts, agreements
 - Law and policing
- These approaches assume:
 - Agents have shared goals and can recognise them
 - Can formulate and communicate plans
 - Can enforce contracts and agreements

Forward link: These kinds of assumptions link to what is often called “cooperative games” within Game Theory. Also they link to what has been termed a “Functionalist” view of society in Sociology (we will talk about these ideas later).

“Non-cooperative” Cooperation

- Such assumptions do not seem to hold:
 - In animal societies
 - In open information systems (P2P)
 - Where agents do not share the same goals
 - Where agents do not have access to fair and working enforcement mechanisms such as law
 - Where agents do not know they have mutual interests in the first place

Forward link: These kinds of assumptions link to “non-cooperative games” in Game Theory. They also link to what might be called a “Conflict” of society in Sociology (we will talk about these ideas later.)

Non-cooperative cooperation

- We will focus here on the non-cooperative cooperation situation
- In such situations there is often what is termed a “social dilemma”
- In a social dilemma it may be in everyone's collective interest to do one thing but everyone's individual interest to do another thing
- Such situations seem to occur in many social, biological and information systems scenarios
- It appears that in many situations agents seem to address them through forms of social interaction

Individualism v. Collectivism

- Consider a P2P file sharing system:
 - It is in the *collective interest* for all to upload to others so everyone gets the file quickly
 - But it is in the *individual interest* to save bandwidth by only downloading and hence free-riding on others
 - Free-riding (or free-loading) is a perennial problem in P2P file-sharing systems (Adar & Huberman 2000)
 - Any efficient system needs to tackle it in some way

Adar & Huberman (2000) Free riding on Gnutella. First Monday, vol. 5 no.10

The tragedy of the commons

- These kinds of situations have been termed “commons dilemmas” or “common pool resource dilemmas”
- Called “dilemmas” because we would all be better off if we “did the right thing” but there is an individual *incentive* to do the wrong thing
- Hardin (1968) summarized the issue in his famous paper: “*The Tragedy of the Commons*”
- These kinds of situations occur in P2P file-sharing systems like *BitTorrent*

Garrett Hardin (1968) The Tragedy of the Commons. Science 162, 1243-1248.

How to avoid a commons tragedy?

- Central enforcement of cooperative behaviour
 - A strong leader imposes cooperation on all (the State)
 - Agents accept / create this because it's good for them
 - Social contract theory: mutual coercion mutually agreed
 - Thomas Hobbes (1651) Leviathan – still important
- Decentralised “bottom-up” cooperation
 - Social interactions constrain individual agent behaviour
 - Producing cooperation emergently
 - Without a third party or enforceable contract
 - Peter Kropotkin (1902) Mutual Aid – now considered biologically flawed

Quote:

"during the time men live without a common power to keep them all in awe, they are in that condition which is called war; and such a war as is of every man against every man."

Thomas Hobbes (English philosopher 1588-1679)

Quote

“In the animal world we have seen that the vast majority of species live in societies, and that they find in association the best arms for the struggle for life: understood, of course, in its wide Darwinian sense — not as a struggle for the sheer means of existence, but as a struggle against all natural conditions unfavourable to the species.”

Peter Kropotkin (Russian anarchist, 1842–1921)

A simple cooperation scenario

- We will now consider a simple scenario that claims to capture the dilemma of cooperation called the *Prisoner's Dilemma*
- Derived from game theory formulations
- With a fascinating history in itself
- Some claim it may give insight into complex ideas like morality, wars and life itself
- It has been studied by philosophers, political scientists, military strategists, computer scientists, biologists...

William Poundstone (1993) *Prisoner's Dilemma: John Von Neumann, Game Theory and the Puzzle of the Bomb* [interesting historical background]

What is game theory?

- Way to mathematically analyse games assuming we know:
 - number of players
 - possible moves they can make (strategies)
 - outcome of game based on players moves (pay-offs)
 - desirability of game outcomes for each player (utility)
 - the players are “rational” (want to maximise utility)

Future link: We will come back to a more detailed look at game theory later.

What is game theory?

- Developed as a response to the problem of the cold war within RAND corporation (post WWII)
- Generally assumes extremely selfish and non-communicating agents (not always true)
- Extremely intelligent and well informed agents who wish to maximise their individual “utility”
- Nice solution concepts elegant mathematics but must be applied with care and understanding!
- Has been applied to human, biological and information systems

John von Neumann and Oskar Morgenstern: Theory of Games and Economic Behavior, Princeton University Press (1944)

What game are you playing?

Games can be categorised into two types:

1) Zero-sum games

- when one player wins another loses
- summing the final utilities of players = 0
- e.g. poker, chess, monopoly etc.

2) Non-zero-sum games

- utilities do not always sum to zero
- both players may lose or both may win (“win-win”)
- considered to capture many social / economic realities
- e.g. tragedy of the commons examples

Capturing a commons tragedy with a simple game

- Consider a game composed of two players:
 - each player:
 - has choice of one move (C or D)
 - makes a single move then the game ends
 - does not know how the other will move
 - gets a payoff (or utility) based on how they moved and how the other player moved
 - for certain payoff values this game can, minimally, capture a form of commons tragedy (or dilemma)
 - a classic such game is called the *Prisoner's Dilemma*

The Prisoner's Dilemma - "payoff matrix"

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

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

The Prisoner's Dilemma - example games

Players =>	P1	P2	P1	P2	P1	P2	P1	P2
Moves =>	C	C	C	D	D	C	D	D
Payoffs =>	R	R	S	T	T	S	P	P
Values =>	3	3	0	5	5	0	1	1
Total =>	6		5		5		2	

A contradiction between collective and individual interests

Story behind the PD game

- Two people are arrested for robbing a bank. The police interrogator can not prove it was them. He puts each in a different room and offers them a deal: “say the other did it and you go free. Keep quiet and you will be charged with a lesser crime”
- Keeping quiet is playing C (cooperate)
- Betraying the other is playing D (defect)
- P (punishment), R (reward), T (temptation) S (sucker)

Game theory says defect!

- Game theory assumes players are:
 - rational - attempt to maximise *their* utility
 - selfish - don't care about the other guy
 - knowledgeable - have complete information
 - clever - have unlimited computational time
- Given these assumptions it can be proved:
 - agents will select equilibria where no player will improve by changing strategy unilaterally
 - many games have such equilibria - by the famous John Nash (so-called *Nash Equilibrium* - NE)
 - the NE for the PD is DD (all defect)

Nash, John (1950) Equilibrium points in n-person games. Proceedings of the National Academy of Sciences 36(1):48-49.

However..

- In real world situations like this people and animals don't seem to always defect
- Why?
- Obviously they are not acting rationally by the definition of game theory and / or...
- In the real world other assumptions might apply

Iterated Prisoner's Dilemma

- Previous example is a “one-shot” PD but:
 - real world interactions often repeated
 - might meet the guy you just ripped-off in the future
 - allows for more complex sequence of strategies based on past interactions with others
 - can punish someone tomorrow for defecting against you today - “the shadow of the future”
- The so-called Iterated PD (IPD) captures this
- This maps well onto P2P file-sharing protocols like Bittorrent (as we will see later)

Iterated PD – example game showing 6 iterations

Iterations	1	2	3	4	5	6	
Player 1	C	D	C	C	D	D	
Player 2	D	C	C	D	D	D	
Payoff P1	0	5	3	0	1	1	total = 10
Payoff P2	5	0	3	5	1	1	total = 15

In general in an IPD it is assumed each iteration is played one after the other and that the players play forever (or don't know when the game will end). Each player knows only what moves were made in the past. Based on the past moves it needs to select the next move. Hence many more strategies which can be more sophisticated.

What is the “rational” thing to do in the IPD according to game theory?

- Application of traditional game theory is tricky:
 - cooperative equilibria exist in infinitely repeated games but not in finite games of known length
 - many equilibria exist and it is not clear which one would be chosen by rational agents
 - defection on every round is still an equilibrium
- For these reasons *Robert Axelrod* (political scientist), in the late 70's, decided to find out what kinds of strategies worked well in the IPD by using computer tournaments
- Robert Axelrod (1984) *The Evolution of Cooperation*, Basic Books

Axelrod's Tournament - programs as strategies

- Axelrod organised an open IPD tournament:
 - Academics were asked to submit programs (BASIC or FORTRAN) that would play the IPD against each other
 - Nobody knew competitors code
 - The only input would be the on-going past history of the game (a string of C's and D's)
 - The aim was to get the highest score (utility) based on round-robin playoffs between all pairs of programs
 - Axelrod's aim was to see which programs did best against all the others and understand why

Axlerod's Tournament - what happened?

- Basic results were:
 - many strategies were submitted (complex and simple)
 - the one with the highest overall score turned out to be simple: *TIT-FOR-TAT* (TFT) or “look back”
 - starts playing C, then “looked back” at the last move made by opponent and copied that move (see P1 in our IPD example)
 - submitted by Psychologist Anatol Rapoport
 - didn't “win” against each strategy but did better overall on average against all strategies
 - TFT mechanism an example of “*reciprocal altruism*” (Robert Trivers (1971) The evolution of reciprocal altruism. Quarterly Review of Biology. 46: 35-57)

Axelrod's tournaments – details

- Ran two tournaments sending out requests for the second tournament with results from the first
- asked “game theorists” for programs (knew the PD)
- Included also a “RANDOM” strategy which just played either C or D at random
- Each round involved every submitted strategy (program) playing five IPD games against every other strategy including itself
- Results were the average payoff of each strategy over the entire round – hence one can rank the strategies
- Used payoffs we showed earlier: $R=3$, $P=1$, $T=5$, $S=0$

Axelrod's tournaments – details

- 1st round tournament: got 14 submitted programs. Each IPD game was 200 individual games of PD in length. TIT-FOR-TAT ranked 1st
- 2nd round tournament: got 62 submitted programs. Each IPD game was random length averaging 151 over the 5 runs. TFT won again!
- Axelrod states that using algorithms in this way allows disciplines to “communicate” in new kind of way!

TABLE 2
The Contestants: Round One

<i>Rank</i>	<i>Name</i>	<i>Discipline (if faculty)</i>	<i>Length of Program</i>	<i>Score</i>
1	Anatol Rapoport	Psychology	4	504.5
2	Nicholas Tideman & Paula Chieruzzi	Economics	41	500.4
3	Rudy Nydegger	Psychology	23	485.5
4	Bernard Grofman	Political Sci.	8	481.9
5	Martin Shubik	Economics	16	480.7
6	William Stein & Amnon Rapoport	Mathematics Psychology	50	477.8
7	James W. Friedman	Economics	13	473.4
8	Morton Davis	Mathematics	6	471.8
9	James Graaskamp		63	400.7
10	Leslie Downing	Psychology	33	390.6
11	Scott Feld	Sociology	6	327.6
12	Johann Joss	Mathematics	5	304.4
13	Gordon Tullock	Economics	18	300.5
14	Name withheld		77	282.2
15	RANDOM		5	276.3

Results of first round tournament.
 From Axelrod (1984) *Evolution of Cooperation*.

TABLE 3
Tournament Scores: Round One

Player	Other Players															Average Score
	TIT FOR TAT	TIDE AND CHIER	NYDEG GER	GROFMAN	SHUBIK	STEIN AND RAP	FRIEDMAN	DAVIS	GRAASKAMP	DOWNING	FELD	JOSS	TULLOCK	(Name Withheld)	RANDOM	
1. TIT FOR TAT (Anatol Rapoport)	600	595	600	600	600	595	600	600	597	597	280	225	279	359	441	504
2. TIDEMAN & CHIERUZZI	600	596	600	601	600	596	600	600	310	601	271	213	291	455	573	500
3. NYDEGGER	600	595	600	600	600	595	600	600	433	158	354	374	347	368	464	486
4. GROFMAN	600	595	600	600	600	594	600	600	376	309	280	236	305	426	507	482
5. SHUBIK	600	595	600	600	600	595	600	600	348	271	274	272	265	448	543	481
6. STEIN & RAPOPORT	600	596	600	602	600	596	600	600	319	200	252	249	280	480	592	478
7. FRIEDMAN	600	595	600	600	600	595	600	600	307	207	235	213	263	489	598	473
8. DAVIS	600	595	600	600	600	595	600	600	307	194	238	247	253	450	598	472
9. GRAASKAMP	597	305	462	375	348	314	302	302	588	625	268	238	274	466	548	401
10. DOWNING	597	591	398	289	261	215	202	239	555	202	436	540	243	487	604	391
11. FELD	285	272	426	286	297	255	235	239	274	704	246	236	272	420	467	328
12. JOSS	230	214	409	237	286	254	213	252	244	634	236	224	273	390	469	304
13. TULLOCK	284	287	415	293	318	271	243	229	278	193	271	260	273	416	478	301
14. (Name Withheld)	362	231	397	273	230	149	133	173	187	133	317	366	345	413	526	282
15. RANDOM	442	142	407	313	219	141	108	137	189	102	360	416	419	300	450	276

Results of first round tournament – detail.
From Axelrod (1984) Evolution of Cooperation.

Results of second round tournament. Only first 35 strategies are shown.
 From Axelrod (1984)
 Evolution of Cooperation.

TABLE 4
The Contestants: Round Two

Rank	Name	Country (if not U.S.)	Discipline (if faculty)	Language (FORTRAN or BASIC)	Length of Program ^a
1	Anatol Rapoport	Canada	Psychology	F	5
2	Danny C. Champion			F	16
3	Otto Borufsen	Norway		F	77
4	Rob Cave			F	20
5	William Adams			B	22
6	Jim Graaskamp & Ken Katzen			F	23
7	Herb Weiner			F	31
8	Paul D. Harrington			F	112
9	T. Nicolaus Tideman & P. Chieruzzi		Economics	F	38
10	Charles Kluepfel			B	59
11	Abraham Getzler			F	9
12	Francois Leyvraz	Switzerland		B	29
13	Edward White, Jr.			F	16
14	Graham Eatherley	Canada		F	12
15	Paul E. Black			F	22
16	Richard Hufford			F	45
17	Brian Yamauchi			B	32
18	John W. Colbert			F	63
19	Fred Mauk			F	63
20	Ray Mikkelson		Physics	B	27
21	Glenn Rowsam			F	36
22	Scott Appold			F	41
23	Gail Grisell			B	10
24	J. Maynard Smith	United Kingdom	Biology	F	9
25	Tom Almy			F	142
26	D. Ambuelh & K. Kickey			F	23
27	Craig Feathers			B	48
28	Bernard Grofman		Political Sci.	F	27
29	Johann Joss	Switzerland	Mathematics	B	74
30	Jonathan Pinkley			F	64
31	Rudy Nydegger		Psychology	F	23
32	Robert Pebley			B	13
33	Roger Falk & James Langsted			B	117
34	Nelson Weiderman		Computer Sci.	F	18 ²¹
35	Robert Adams			B	43

Observations and findings

- Axelrod observed that many of the strategies that did well were variants on TFT
- In general they are:
 - Nice: they don't defect first
 - Punish defection from other quickly
 - Forgiving: they stop defecting when the other does

An interpretation from Axelrod(1984)

- First world war fraternisation
 - When soldiers found themselves stuck in trenches for long periods of time on front line
 - Launching shells or not was a kind of cooperate / defect strategy
 - They were playing an IPD
 - Learned to stop shelling each other

What has this got to do with Evolution?

- If we view the utilities different strategies obtain by playing such games as fitness
- Then strategies that perform well will tend to reproduce in a population
- Replacing those that perform less well
- Hence an evolutionary process should select them
- We will revisit these ideas in later lectures

Readings and questions

- Books:
 - Flake book – chapter 17
 - Axelrod (1984) The evolution of cooperation. Basic Books.
- Questions:
 - Can you think of some situations (other than those discussed) in which social interactions appear to produce a PD?
 - Can you think of any other ways (than those discussed) that might produce cooperation in a PD?
 - How is “rational” behaviour defined or used in everyday language (as apposed to utility maximisation)?
 - Do you think the approaches discussed have anything to say about “morality”?
 - Can you think of an application for these approaches in info. systems?