Evol Comp II: Strategies
The role of game theory in theorising and the evolution of cooperation

Overview of lecture

- Part 1: Classical tools to understand evolution
  - Classical Game theory
  - Evolutionary game theory
- Part 2: Evolution of cooperation
  - Altruism doesn't work
  - Prisoner's dilemma
  - Iterated prisoner's dilemma (IPD)

Background

- Axelrod, R. (1984). *The evolution of co-operation*. Basic Books, NY.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.
- Ridley, M. (1996). *The origins of virtue: Human Instincts and the evolution of cooperation*. Viking.
- Pinker, S. (1997) *How the Mind Works.* Penguin Press.

Review: Rethinking Competition in Evolutionary theory

- In evolution, what is "competition"?
  - Not just tooth and claw…
- When do animals cooperate?
  - Not just kin…

What factors should be included in developing an evolutionary model?

- What aspects of the simulation are relevant?
- What level of detail is required?
  For example,
  - To explain the form of a bird's wing, relevant factors include atmospheric conditions, lift and drag on wing shape, the properties of feathers, etc. A successful wing shape is one that optimizes flight with respect to the physical world. An appropriate formalism would be *optimization theory*.
  - By contrast, to explain *variability* within a species, one needs to take into account not just the physical terrain or predators. Success depends on what other members of the species are doing, s.a., finding mates, avoiding competition for resources, joint protection against predators. Optimization theory doesn't produce useful models because the things being optimized are constantly changing. An appropriate formalism is *game theory*, which takes actions of other agents into account.

Classical Game theory

- Specification of a "game" can include player characteristics, rules, informational assumptions and payoffs.
- *Classical game theory* was developed to model human economic behaviour (~1950s by von Neumann and Morgenstern).
- Its basic assumptions are that players behave rationally and in their own self-interest.
- These assumptions don't hold in an evolutionary context (rationality requires cognitive abilities beyond most species, and that given such abilities, they also have sufficient knowledge of events and consequences to make appropriate decisions).

Evolutionary game theory

- *Evolutionary game theory* applied ideas from classical game theory to evolutionary situations.
- EGT is a way of thinking about evolution at the level of the phenotype, when fitness depends on the frequencies of other phenotypes in the population.
- Its basic assumptions are that populations have characteristic dynamics and stability and that agents optimize their Darwinian fitness.

Game theory ideas have been applied to biology in two main ways:

- Initially by Lewontin (1961) who modeled problems of species against nature. L. assumed that species minimize their probability of extinction (species level reasoning is fraught with issues).
- Secondly, by Maynard Smith and colleagues (~1973), who modeled members of a population playing against each other, and studied the population dynamics and equilibrium states.

- EGT models pair-wise contests between individuals or in more advanced models, individual interactions with a group. Usually a finite set of strategies is assumed.

Defn: *Evolutionarily stable strategy (ESS)*

- A "*strategy*" is a behavioural phenotype. That is, it specifies what an individual will do in any situation.
- An *ESS* is a strategy such that if all members of a population adopt it, then no mutant strategy could invade the population under the influence of natural selection.

Defn: *Evolutionarily stable <u>states</u>*

- A state is a set of strategies.
- A stable state is one where the relative frequencies of the strategies do not change over time.
  - This includes mixed strategies (An individual can adopt a mixed strategy (do X 20% of the time and Y 80% of the time), which could be a mixed ESS.
  - Or, within a population, each individual might adopt a pure strategy. Stable ratios of such strategies would constitute stable polymorphic states.)
  - e.g., percentage of sons vs. daughters is stable over time.

Strengths of EGT

- Simple models can have heuristic value, even though they are not testable in an empirical way.
  - For example, studying the dynamics of predator and prey species lead to the discovery that the population numbers will oscillate.
- Applications of EGT include:
  - Animal distribution (dispersal)
  - Sex ratio theory
  - Contest behaviour
  - Recipricol altruism
  - Parental investment
  - Polygyny in social animals (eg mammals)
- EGT is now used in biology, economics, anthropology, political science, sociology, and social psychology.

Weaknesses of EGT

- Emphasis on equilibrium states
- Games may not have an ESS
- Trajectory of change may be important
- Questions about individual behaviour cannot be answered with EGT

- Caveat emptor: Simple models can indicate possible behaviours in a variety of situations. However, when applied to real world problems, testability is an essential requirement, and heuristic value is not enough.

Example: The Hawk-Dove game

- Two animals are contesting a resource of value, V (eg nesting site).
  - For example, V=2 if it represents the difference between a more favourable (5 offspring per season) and a less favourable habitat (3 offspring per season).
- During the contest, an animal has the options:
  - display – threaten but do not injure the opponent
  - escalate – fight
  - retreat – abandon resource to the opponent
- Assume that individuals in a given contest adopt one of two "strategies", and they always behave in the same way:
  - "Hawk" escalate and continue until injured or opponent retreats
  - "Dove" display; retreat at once if opponent escalates.
- If two contestants both escalate, it is assumed that sooner or later one is injured and forced to retreat.
  - Injury reduces fitness by a cost, C.

Payoffs for the Hawk-Dove game

	H	D
H	½(V-C)	V
D	0	V/2

- Hawk vs Hawk:
  - Each contestant has a 50% chance of injuring its opponent and obtaining the resource, V, and a 50% chance of being injured.
- Hawk vs Dove:
  - Hawk obtains the resource and Dove retreats before being injured. NB. Doves do not have zero fitness, just that this encounter does not change their fitness.
- Dove vs Dove:
  - The resource is shared equally by the two contestants

More issues in the Hawk-Dove game …

- Game theory studies the frequency of the strategies in the population.
  - D is not an ESS, since E(D,D)<E(H,D).
  - H is an ESS if V>C.
  - If V<C neither H nor D is an ESS.
- However, let I be a mixed strategy, where an individual plays H with probability P and D with probability (1-P).
- I can be an ESS, where E(H,I) = E(D,I)
- Solve for the values above gives P=V/C.
- Some further calculation shows that a genetic polymorphism with a frequency of p=V/C of pure Hawk is also stable.

- Maynard Smith extended the Hawk-Dove to include "retaliator" strategies and a whole variety of extensions.

Summary: Main points in game theory and evolutionary game theory

- Defns: ESS and evolutionary stable states
- Strengths and weaknesses
- Example: Hawk-Dove game

The chocolate game

Rules of the Chocolate Game

- A game consists of several rounds
  - the number is known to the administrator, but not to the participants beforehand
- In each round, two names, A and B, are drawn at random from the participants.
  - A offers B a number of points, 0 <=*p*<=10.
  - B can accept or reject the offer.
    - If B accepts, B receives *p* points, and A receives (10-*p*) points.
    - If B rejects, neither receive any points.
  - Points that are not accepted lapse
    - i.e., they are not recycled into extra rounds of the game.
- Points from all rounds are summed and at the end of all rounds,
  - every 7 points are worth 1 freddo. Less than 7 points is worth nothing.
  - For example,
    - 16 points = 2 freddos
    - 7 points = 1 freddo
    - 6 points = nothing
- A and B are not allowed to negotiate or communicate in any way.
  - If they do, both lose all points for that round.
- All decisions of the administrator are final.

Part 2. Evolution of cooperation
How is cooperation integrated into theories in EC?

- Altruism is an evolutionary impossibility
- But cooperation is an evolutionary fact.

Genes that benefit other genes at their own expense are an evolutionary impossibility

- Defn: Altruism
  - Altruism is consistent behaviour by one individual that benefits unrelated individuals to the net detriment of the altruist.
- Altruists (as defined above) may be good for the species, but their genes can't last in a world of competition
- Consider and instinctively sharing society
  - Fixed amount of resources
  - More food -> bigger, stronger individuals -> more offspring
  - As soon as a selfish mutant arises, those genes will have an advantage.
  - The populations will eventually be dominated by selfish genes

Cooperation is an evolutionary fact

- Nice guy genes do seem to compete very effectively – why?
- The puzzle of cooperation is that we can see that sharers can survive hard times better – selfish individuals are not necessarily fitter.
  - Cooperation exists everywhere, and animals, not just humans are "nice" to each other, even when there is seemingly little benefit to themselves.
  - In part this can be explained in terms of kin selection (its worth making sacrifices for those you are related to, since you are favouring copies of your own genes), but kin selection doesn't explain cooperation with non-kin.
- Something is missing in the classic account of competition.

The Prisoner's dilemma

(A tool for reasoning about the dilemma of cooperation)

- Two partners-in-crime have been nabbed.
- If neither confess, the police have enough evidence to send both away for 1 year.
- If one confesses, he will be let off on good behaviour and the other will get 20 years.
- If both confess they will each get 10 years.
- They are held incommunicado. What should they do?

Payoff Matrix

	Cooperate	*Defect*
Cooperate	*3* **R=3**	*5* **S=0**
Defect	*0* **T=5**	*1* **P=1**

- Temptation = 5, Reward = 3, Sucker = 0, Penalty = 1
- T>R>P>S and R>(S+T)/2

Why call it a dilemma?

- The dilemma is that
  - if the other person defects, your best option is to defect (1 point vs. none),
  - but even if the other person cooperates, your best option is still to defect (you get 5 points rather than 3).
- So both reason that the other person will defect, and they both get 1. If they had both kept quite, they would both have had 3 points.

Examples of real world scenarios

- In prisoner's dilemmas, nice guys finish last
  - Tragedy of the commons
  - Cold war
  - Airport food
  - Species extinctions

Extending the prisoner's dilemma …

- What's different in real world situations?
- History of the Prisoner's Dilemma
  - 1950s – Flood and Dresher formalised
  - 1960s – defection is the only rational approach
  - 1970s – Maynard Smith evolutionary game theory
  - 1979 – Axelrod ran a computer competition

Iterated prisoner's dilemma (IPD)

- The PD Game is played repeatedly
- Individuals remember others and how they previously behaved and can base their decisions on past performance
- Outcomes:
  - Different strategies have different long term payoffs
  - E.g. always cooperate, or always defect, or guess what your opponent will do
- 1979: Axelrod's computer based IPD tournament
  - easily implemented as an evolutionary algorithm, and strategies can be evolved and tested computationally

Example contest

Round	1	2	3	4	
Individual A	C	C	D	D	
Individual B	C	D	D	D	
Payoff for A	3	0	1	1	total = 5
Payoff for B	3	5	1	1	total = 10

- If the number of offspring was proportional to this contest, then B would have twice as many offspring as A.

The course of evolution of strategies in the IPD

- 3 strategies are
  - nice (always cooperate),
  - nasty (always defect)
  - TfT (always do what your opponent did last time).
- Start with equal numbers of each
  - Initially nasty does well at the expense of nice, TfT holds its own
  - As nice strategies are wiped out, nasty begin to decline
  - TfT dominates

Tit For Tat

- TfT is an ESS in the standard game, but is not in a noisy environment
- Against itself it may either cooperate, or show repeated mutual defection, or oscillate
- Other strategies may do better in different environments

Strategies used in the lab exercise

from http://www.lifl.fr/IPD/ipd.html#cipd

- Here is a description of some of the basic strategies used in our simulations as well as in the literature:
  - **all_c**
    - Always cooperates. [c]*
  - **all_d**
    - Always defects. [d]*
  - **tit_for_tat**
    - The tit_for_tat strategy was introduced by Anatole Rapoport. It begins to cooperate, and then play what its opponent played in the last move.
  - **spiteful**
    - It cooperates until the opponent has defected, after that move it always defects.
  - **soft_majo**
    - Plays opponent's majority move, if equal then cooperates. First move is considered to be equality.
  - **per_ddc**
    - Plays periodically : [d,d,c]*
  - **per_ccd**
    - Plays periodically : [c,c,d]*
  - **mistrust**
    - Defects, then plays opponent's move.
  - **per_cd**
    - Plays periodically [c,d].
  - **pavlov**
    - The win-stay/lose-shift strategy was introduced by Martin Nowak and Karl Sigmund. It cooperates if and only if both players opted for the same choice in the previous move.
  - **tf2t**
    - Cooperates except if opponent has defected two consecutive times.
  - **hard_tft**
    - Cooperates except if opponent has defected at least one time in the two previous move.
  - **slow_tft**
    - Plays [c,c], then if opponent plays two consecutive time the same move plays its move.
  - **hard_majo**
    - Plays opponent's majority move, if equal then defects. First move is considered to be equality.
  - **random**
    - Cooperates with probability 1/2.

Issues for simulations

- Two major decisions for a GA
  - Problem representation
  - Fitness function
- For IPD, fitness is straightforward
  - points accumulate over all games in a generation
- Problem representation can vary
  - Two bits allows four strategies
  - Other representations allow a wider range of strategies
  - Decision is always binary – cooperate / defect
- Implementation details may affect the simulation

Extensions

- How to model more sophisticated strategies
  - Degrees of cooperation
  - Longer memory
  - Allow evolving strategies, e.g., using a neural net
- More realistic environments
  - Adding uncertainly,
  - Noise or asynchrony
  - Spatial organisation
  - restricted interaction
- Reputation
- Refusing to play (optional PD) – discriminating altruists

Implications / Links to other ideas

(Important for fields such as Evol Psych)

- Reciprocity is rampant in human society
- Instinctive?
- Only works when people recognize each other (~150 individuals for humans)
- Humans have incredibly good cheater-detection reasoning (much better than reasoning about other factors)

Summary of IPD

- One shot prisoner's dilemma has a single optimal strategy – always defect.
- Does not correspond with the generally high level of cooperation observed in real life
- Iterated prisoner's dilemma favours the existence of "nice" strategies (although in most environments a degree of retaliation is essential).
- Suggests one mechanism for the evolution of altruistic acts.
- A good example of a problem for which EC is an essential tool.

Rethinking Cooperation and Competition in Evolutionary theory

- In evolution, what is "competition"?
  - Not just tooth and claw!
- When do animals cooperate?
  - Not just kin!

Turtles all the way down (1) Competition

- Competition is with
  - Other species
  - Unrelated members of own species
  - Siblings
  - Between parents
  - Within an individual, between their own genes
  - Possibly even different segments of DNA (parasitic DNA)

Turtles all the way down (2) Cooperation

- Cooperation is with
  - Other species (honeybee birds,)
  - Unrelated members of own species (guardians, trade)
  - Siblings (strongest allies)
  - Between parents
  - Within an individual, between their own genes

End of Lecture

- Additional notes follow

Extensions to IPD

- **Game Theory: When it pays not to have perfect knowledge**
- The prisoner's dilemma has some unexpected twists to it. Consider these variations:-
  - Variation 1: Two recent members of a gang are nabbed. They don't have a strong relationship between each other, but the rules of the gang are that if any member turns informer, the gang will meet, hold a vote, and if the crime is considered bad enough, will kill the informer. With all the gang warfare around, there's about a 50% chance that this would happen. What would you do?
  - Variation 2: [now for the interesting case] One of the prison priests is actually a trusted friend of both the prisoner's, and offers to hook up a radio so that they can both know whether the meeting has taken place. To make it fair, either both will know, or neither. He gives both them the option of vetoing the radio. What would you do?

Rules of the Chocolate Game

- A game consists of several rounds
  - the number is known to the administrator, but not to the participants beforehand
- In each round, two names, A and B, are drawn at random from the participants.
  - A offers B a number of points, $0 <= p <= 10$.
  - B can accept or reject the offer.
    - If B accepts, B receives *p* points, and A receives (10-*p*) points.
    - If B rejects, neither receive any points.
  - Points that are not accepted lapse
    - i.e., they are not recycled into extra rounds of the game.
- Points from all rounds are summed and at the end of all rounds,
  - every 7 points are worth 1 freddo. Less than 7 points is worth nothing.
  - For example,
    - 16 points = 2 freddos
    - 7 points = 1 freddo
    - 6 points = nothing
- A and B are not allowed to negotiate or communicate in any way.
  - If they do, both lose all points for that round.
- All decisions of the administrator are final.

Chocolate Game debrief

Games: studies of rational fools?

- The chocolate game is one example of a formal game. There are many such, studying different aspects of interaction between agents. Many focus on cooperation and competition, and human reactions in those situations.
- In the chocolate game, on purely rational grounds, it is sensible for B to accept any number of points above zero, since any number of points is useful.
- Hence, if A reasons that B is rational, and A is trying to maximize their own gain, it makes sense for A to offer just one point, and hence collect 9 for themselves.
- In its original formulation, the game is played with just one round, with $100 on offer.
- A and B are unknown to each other, separated in different rooms.
- A can offer any amount to B, from $0-$100.
- On purely rational grounds, B should accept any offer from $1 up.
- In actual experiments, offers much below 40% are almost universally rejected.
- Different subject populations behave quite differently, with students who have studied economic theory offering much less than other participants (and getting knocked back for their offers!).
- Contemporary studies aim to explain such differences