Evolution of Cooperation II:
Replicator Dynamics and Dynamic Networks

Models and Simulations in Philosophy
May 7th, 2013

Last Week:

- Repeated Prisoners' Dilemmas
- Boundedly Rational Learning in Networks

Today: Did we dismiss population level explanations too early?

Today

Today: Did we dismiss population level explanations too early?

What happens in our ABMs if the interaction network is not fixed?

Outline

Outline

1. From Rational Choice Theory to Biological Evolution and Back
2. Replicator Dynamics
   - One-Shot PDs
   - Repeated PDs

Outline

1. From Rational Choice Theory to Biological Evolution and Back
2. Replicator Dynamics
    - One-Shot PDs
    - Repeated PDs
3. Dynamic Networks

Outline

1. From Rational Choice Theory to Biological Evolution and Back
2. Replicator Dynamics
   - One-Shot PDs
   - Repeated PDs
3. Dynamic Networks
4. NetLogo

Outline

- One common "population level" model for explaining the emergence of norms is called the replicator dynamics.

- One common "population level" model for explaining the emergence of norms is called the replicator dynamics.
- The model was originally introduced in biology.

- One common "population level" model for explaining the emergence of norms is called the replicator dynamics.
- The model was originally introduced in biology.
- So before discussing the replicator dynamics, it will be helpful to discuss a common and fruitful way of thinking that shows similarities between
  - Models of natural selection
  - Rational Choice Theory (i.e. decision and game theory)

Decision Theory in Biology

How are decision (and game theory) relevant for models of natural selection?

An Informal Argument:

1. Actions = Phenotypes (e.g., traits and behaviors)

An Informal Argument:

1. Actions = Phenotypes (e.g., traits and behaviors)
2. Payoffs = Offspring

An Informal Argument:

1. Actions = Phenotypes (e.g., traits and behaviors)
2. Payoffs = Offspring
3. By definition, organisms that have the highest actual payoffs (offspring) will become more prevalent in the population.

An Informal Argument:

1. Actions = Phenotypes (e.g., traits and behaviors)
2. Payoffs = Offspring
3. By definition, organisms that have the highest actual payoffs (offspring) will become more prevalent in the population.
4. So intuitively, actions (i.e. phenotypes) that have the highest expected payoffs (offspring) will become more prevalent.
   - The expected number of offspring of an organism with a given phenotype, given the current distribution of phenotypes in the population, is often called the fitness of the phenotype.

Conclusion: Nature can modeled as choosing organisms with particular phenotypes so as to maximize expected utility, where utility is number of offspring.

There exist deep and interesting connections, both thematic and formal, between evolutionary theory and the theory of rational choice . . . In rational choice theory, agents are assumed to make choices that maximize their utility, while in evolutionary theory, natural selection 'chooses' between alternative phenotypes, or genes, according to the criterion of fitness maximization. As a result, evolve organisms often exhibit behavioral choices that appear designed to maximize their fitness, which suggests the principles of rational choice might be applicable to them.

Okasha and Binmore [2012].

Three Observations

There are at least three aspects of this argument that deserve attention . . .

- **Observation 1:** The argument does not mention the interaction among organisms at all: that is, it seems equally applicable in either
  - "Game Theory" - An organism's reproductive success depends upon the behavior of conspecifics (i.e. organisms of the same species), predators, etc.
    - E.g., Conspecifics competing for a resource
  - "Decision Theory" - An organism's reproductive success is largely a function of its traits and environment minus competing organisms.
    - E.g., Thickness of polar bear fur in response to cold.

1. **Observation 2:** Consider the inference:
    - From: Organisms with highest actual payoffs will become more prevalent
    - To: Organisms with highest expected payoffs will become more prevalent
2. It is precisely the same type of inference that seemed to fail in the models of the prisoner's dilemmas on networks we studied last time. Why?
    - Since defecting is the dominant action in the prisoner's dilemma, it also maximizes SEU.

1. **Observation 2:** Consider the inference:
    - From: Organisms with highest actual payoffs will become more prevalent
    - To: Organisms with highest expected payoffs will become more prevalent
2. It is precisely the same type of inference that seemed to fail in the models of the prisoner's dilemmas on networks we studied last time. Why?
    - Since defecting is the dominant action in the prisoner's dilemma, it also maximizes SEU.
    - But non-cooperative behavior did not emerge in all the network models we discussed last class.

1. **Observation 2:** Consider the inference:
   - From: Organisms with highest actual payoffs will become more prevalent
   - To: Organisms with highest expected payoffs will become more prevalent
2. It is precisely the same type of inference that seemed to fail in the models of the prisoner's dilemmas on networks we studied last time. Why?
   - Since defecting is the dominant action in the prisoner's dilemma, it also maximizes SEU.
   - But non-cooperative behavior did not emerge in all the network models we discussed last class.
   - So what assumptions does the above informal argument (about the emergence of expected offspring) make that are false in the network models we considered last time? [Think about this.]

Moral: Rational choice theory can be useful in helping us understand evolution.

What about the reverse?

- **Observation 3:** The argument discusses phenotypes, i.e realized behaviors or traits: organisms' genotypes matter only insofar as they produce traits or behavior that affect survival and reproduction.

- **Observation 3:** The argument discusses phenotypes, i.e realized behaviors or traits: organisms' genotypes matter only insofar as they produce traits or behavior that affect survival and reproduction.
- Human culture is a collection of behaviors and artifacts: it is not primarily a genetic phenomenon.

- **Observation 3:** The argument discusses phenotypes, i.e realized behaviors or traits: organisms' genotypes matter only insofar as they produce traits or behavior that affect survival and reproduction.
- Human culture is a collection of behaviors and artifacts: it is not primarily a genetic phenomenon.
- So models of natural selection might be applicable to modeling cultural evolution as well.

- **Observation 3:** The argument discusses phenotypes, i.e realized behaviors or traits: organisms' genotypes matter only insofar as they produce traits or behavior that affect survival and reproduction.
- Human culture is a collection of behaviors and artifacts: it is not primarily a genetic phenomenon.
- So models of natural selection might be applicable to modeling cultural evolution as well.
- In biology, there is a mechanism by which traits are passed from parent to offspring: genes. What is the corresponding mechanism for culture?

Because of their common informational and evolutionary character, there are strong parallels between genetic and cultural modeling [Mesoudi et al., 2006]. Like biological transmission, culture is transmitted from parents to offspring, and like cultural transmission, so in microbes and many plant species, genes are regularly transferred across lineage boundaries [Abbott et al., 2003, Jablonka and Lamb, 1995, Rivera and Lake, 2004]. Moreover, anthropologists reconstruct the history of social groups by analyzing homologous and analogous traits, much as biologists reconstruct the evolution of species by the analysis of shared characters and homologous DNA [Mace et al., 1994]. Indeed, the same programs biological systematists are used by cultural anthropologists [Holden, 2002, Holden and Mace, 2003].

Gintis [2012]. pp 216-217.

Outline

Let's see how the above argument can be made a bit more precise …

Imagine two conspecifics are competing for a similar resource:

	Hawk	Dove
Hawk	0, 0	4, 1
Dove	1, 4	2, 2

Hawk Dove

Imagine two conspecifics are competing for a similar resource:

	Hawk	Dove
Hawk	0, 0	4, 1
Dove	1, 4	2, 2

Story: Individuals either fight ("Hawk") or not for the resource.

- If they don't fight, they share the value of the resource.
- If one fights and the other doesn't, then the one who fights gets the resource.
- However, the cost of the fighting is higher than the value of the resource.

Imagine two conspecifics are competing for a similar resource:

	Hawk	Dove
Hawk	0, 0	4, 1
Dove	1, 4	2, 2

Imagine two conspecifics are competing for a similar resource:

	Hawk	Dove
Hawk	0, 0	4, 1
Dove	1, 4	2, 2

What is/are the Nash equilibria of this game?

- Imagine the payoffs indicate offspring: conspecifics with more resources reproduce more often.

- Imagine the payoffs indicate offspring: conspecifics with more resources reproduce more often.
- Now imagine that conspecifics, when searching for resources, encounter other **random** members of the population.

- Imagine the payoffs indicate offspring: conspecifics with more resources reproduce more often.
- Now imagine that conspecifics, when searching for resources, encounter other **random** members of the population.
- What are their expected payoffs?

- Let p denote the proportion of the population that exhibit Hawk behavior, and $1 - p$ that Dove behavior.

- Let p denote the proportion of the population that exhibit Hawk behavior, and $1 - p$ that Dove behavior.
- Then, as the conspecifics encounters other random members of the population, the fitness (i.e. expected utility) of Hawk and Dove respectively are:

$$\begin{aligned} F(Hawk) &= p \cdot 0 + (1-p) \cdot 4 \\ F(Dove) &= p \cdot 2 + (1-p) \cdot 1 \end{aligned}$$

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

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

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

$$F(\text{AVE}) = p \cdot F(\textit{Hawk}) + (1 - p) \cdot F(\textit{Dove})$$

- 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.

In large populations, therefore, one can show that proportion of Hawk players change as follows:

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

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

Suppose conspecifics play a **one-shot** prisoner's dilemma, rather than Hawk/Dove, and their proportions evolve according to the replicator dynamics.

What happens?

Remember, dominant actions maximize expected utility:

	Sun	Rain
Frisbee	5	-1
Biergarten	4	-2

Remember, 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

Dominance in PDs

And recall, defecting is dominant in a prisoner's dilemma:

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

And recall, defecting is dominant in a prisoner's dilemma:

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

Suppose one believes that one's opponent will Defect with probability p.

$$\begin{aligned} \text{SEU}(\textit{Defect}) &= (1 \cdot p) + (3 \cdot (1 - p)) \\ \text{SEU}(\textit{Cooperate}) &= (0 \cdot p) + (1 \cdot (1 - p)) \end{aligned}$$

Each term in the sum of Defect is bigger than the corresponding term for Cooperate.

Defecting in Evolutionary Settings

In the evolutionary setting:

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

In the evolutionary setting:

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

Suppose the **proportion of the population** defecting is p. Then the **fitness** of each phenotype is:

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

Defect has higher fitness than Cooperate.

Replicator Dynamics for One-Shot PDs

Suppose the population consists of both Defectors and Cooperators: so both p and $1-p$ are greater than zero.

- Since Defect has higher fitness than Cooperate:

$$F(\mathit{Defect}) > F(\textsc{ave})$$

- Hence, according to the replicator dynamics:

$$p' - p = p(F(\mathit{Defect}) - F(\textsc{ave})) > 0$$

- In other words, the proportion of defectors increases.

Replicator Dynamics for One-Shot PDs

Suppose the population consists of both Defectors and Cooperators: so both p and $1-p$ are greater than zero.

- Since Defect has higher fitness than Cooperate:

$$F(\textit{Defect}) > F(\textsc{ave})$$

- Hence, according to the replicator dynamics:

$$p' - p = p(F(\textit{Defect}) - F(\textsc{ave})) > 0$$

- In other words, the proportion of defectors increases.
- Since this happens every stage, cooperation will be driven to extinction!

Replicator Dynamics for Repeated PDs

What happens if fitness depends upon a repeated prisoner's dilemma?

- That is, one each stage, two random conspecifics meet and play a prisoner's dilemma some fixed number of times.

- That is, one each stage, two random conspecifics meet and play a prisoner's dilemma some fixed number of times.
- Their actual number of offspring depend upon their payoffs in this repeated prisoner's dilemma.

- That is, one each stage, two random conspecifics meet and play a prisoner's dilemma some fixed number of times.
- Their actual number of offspring depend upon their payoffs in this repeated prisoner's dilemma.
- Then the replicator dynamics is applied to payoffs in this repeated game.

Replicator Dynamics for Repeated PDs

- That is, one each stage, two random conspecifics meet and play a prisoner's dilemma some fixed number of times.
- Their actual number of offspring depend upon their payoffs in this repeated prisoner's dilemma.
- Then the replicator dynamics is applied to payoffs in this repeated game.
- **Memory Test:** Are there are dominant strategies in a repeated prisoner's dilemma?

- No strategies are dominant in a repeated PD.

- No strategies are dominant in a repeated PD.
- So the above argument does not show that always defecting (in a repeated PD) will spread through a population.

- No strategies are dominant in a repeated PD.
- So the above argument does not show that always defecting (in a repeated PD) will spread through a population.
- In fact, Alexander [2007] simulation results show that cooperative can persist in populations in which strategies (with limited memory) for a repeated PD are chosen at random.

Replicator Dynamics vs. Network Models

Discussion Question: What are the differences between the replicator dynamics and the network models we discussed the last two classes?

Discussion Question: What are the differences between the replicator dynamics and the network models we discussed the last two classes?

Replicator	Networks
Random Interactions	Local Interactions
Large Population	Potentially Small

Recall, we assumed that actual offspring most approximate the expected number in the replicator dynamics, which may not be true in small populations

Outline

1. From Rational Choice Theory to Biological Evolution and Back
2. Replicator Dynamics
   - One-Shot PDs
   - Repeated PDs
3. Dynamic Networks
4. NetLogo
5. References

In the past, I've shown you networks like this:

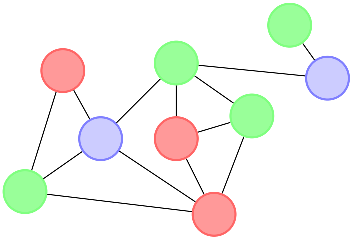

Nodes = Agents
Edges = Indicate which agents "interact"
Colors = "Type" of Agent

But real networks change . . .

- Individuals find new friends and ditch old ones on Facebook.
- Computers in computer networks break and are sometimes replaced.
- Airports in airport networks are abandoned or shut down particular flights.
- Authors on the WWW add new pages, destroy old hyperlinks, etc.
- And so on.

Question: If we were just interested in "how possible" stories for the evolution of cooperation, then why consider dynamic networks at all?

We already have how possible stories using population models (e.g., the replicator dynamics) and static network models.

Further, dynamic network models will likely also be too simple and idealized to provide "how so" explanations.

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.

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.

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.
- Dynamic network models need not even be "more realistic" than the static ones to accomplish this goal.

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.
- Dynamic network models need not even be "more realistic" than the static ones to accomplish this goal.
- Of course, adding more realism does provide additional confidence.

Two ways to change a network:

1. Add and delete agents.
2. Add and delete edges.

Alexander [2007] considers only modifications of the second type.

There are several different ways of changing edges:

1. In the game-theoretic setting: form links with those with whom you earned higher payoffs in the past.
    - This is the model Alexander describes.
    - Perhaps unsurprisingly, cooperators stop interacting with defectors in PDs.
    - So cooperation can be sustained in a population, which self segregates according to strategy.

There are several different ways of changing edges:

1. In the game-theoretic setting: form links with those with whom you earned higher payoffs in the past.
   - This is the model Alexander describes.
   - Perhaps unsurprisingly, cooperators stop interacting with defectors in PDs.
   - So cooperation can be sustained in a population, which self segregates according to strategy.
2. But there are lots of other methods for changing networks. See Bilgin and Yener [2006] for a survey.

The most common model for dynamic networks is called preferential attachment: agents form new link to agents that have many existing neighbors.

Preferential Attachment

The most common model for dynamic networks is called preferential attachment: agents form new link to agents that have many existing neighbors.

1. The idea is that edges represent status, and agents try to gain status by forming links with those who have it.
2. Think of co-authorship among scientists: writing a paper with a famous scientist makes you look good.

Preferential Attachment

The most common model for dynamic networks is called preferential attachment: agents form new link to agents that have many existing neighbors.

1. The idea is that edges represent status, and agents try to gain status by forming links with those who have it.
2. Think of co-authorship among scientists: writing a paper with a famous scientist makes you look good.
3. Preferential attachment models evolve to produce power law degree distributions, which lends them some measure of empirical support for certain social networks.

Topics we'll discuss today:

- Nested Loops and If-Statements
- Procedures and Reporters
- Global vs. Local Variables
- Writing Pseudocode

Abbott, R. J., James, J. K., Milne, R. I., and Gillies, A. C. (2003). Plant introductions, hybridization and gene flow. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 358(1434):1123–1132.

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

Bilgin, C. C. and Yener, B. (2006). Dynamic network evolution: Models, clustering, anomaly detection. *IEEE Networks*.

Gintis, H. (2012). An evolutionary perspective on the univfication of the behavioral sciences. In *Evolution and Rationality: Decisions, Cooperation, and Strategic Behavior*. Cambridge University Press.

Holden, C. J. (2002). Bantu language trees reflect the spread of farming across sub-saharan africa: a maximum-parsimony analysis. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 269(1493):793–799.

Holden, C. J. and Mace, R. (2003). Spread of cattle led to the loss of matrilineal descent in africa: a coevolutionary analysis. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 270(1532):2425–2433.

Jablonka, E. and Lamb, M. J. (1995). *Epigenetic inheritance and evolution: the Lamarckian dimension*. Oxford University Press.

Mace, R., Pagel, M., Bowen, J. R., Otterbein, K. F., Ridley, M., Schweizer, T., and Voland, E. (1994). The comparative method in anthropology [and comments and reply]. *Current Anthropology*, 35(5):549–564.

Mesoudi, A., Whiten, A., and Laland, K. N. (2006). Towards a unified science of cultural evolution. *Behavioral and Brain Sciences*, 29(4):329346.

Okasha, S. and Binmore, K. (2012). *Evolution and Rationality: Decisions, Co-operation, and Strategic Behaviour*. Cambridge University Press.

Rivera, M. C. and Lake, J. A. (2004). The ring of life provides evidence for a genome fusion origin of eukaryotes. *Nature*, 431(7005):152–155.