Evolution of Signaling: Part I

Models and Simulations in Philosophy
December 2nd, 2013

This Month: The Platonic Puzzle concerning Meaning

The puzzle consists of three different questions:

This Month: The Platonic Puzzle concerning Meaning

The puzzle consists of three different questions:

- Definition: What makes a signal meaningful?

This Month: The Platonic Puzzle concerning Meaning

The puzzle consists of three different questions:

- Definition: What makes a signal meaningful?
- Evolution: How did meaningful signals evolve?

This Month: The Platonic Puzzle concerning Meaning

The puzzle consists of three different questions:

- Definition: What makes a signal meaningful?
- Evolution: How did meaningful signals evolve?
- Stability: Why are signals stable?

Preview

Definition:

- Lewis [2008]'s definitions of **convention** and **signaling system** are an attempt to characterize, in game theoretic terms, when individuals' signals have acquired meaning.

Definition:

- Lewis [2008]'s definitions of **convention** and **signaling system** are an attempt to characterize, in game theoretic terms, when individuals' signals have acquired meaning.
- In Chapter 3, Skyrms [2010] gives an information-theoretic account of when signals convey information.

Evolution:

- Lewis [2008] argues that common knowledge allows agents to solve coordination problems.

Evolution:

- Lewis [2008] argues that common knowledge allows agents to solve coordination problems.
- Skyrms [2004] and Skyrms [2010] will provide models that show how agents might learn to play a signaling system.

Stability: Both Lewis' definition and Skyrms' models provide explanations as to why conventional meanings might persist.

Last Class

Last Class: Lewis' definition of coordination equilibrium and convention

Today's Class

Today:

- Evolution and Stability of Convention
- Signaling games and signaling systems
- Modeling the evolution of signaling

Outline

Outline

Outline

Outline

1. Review and Preview
2. Convention
   - Stability
   - Evolution/Selection
3. Signaling Games
4. Evolution of Signaling
   - Replicator Dynamics
   - ABMs

Outline

Outline

1. Review and Preview
2. Convention
   - Stability
   - Evolution/Selection
3. Signaling Games
4. Evolution of Signaling
   - Replicator Dynamics
   - ABMs
5. Netlogo
6. References

Nash Equilibria vs. Coordination Equilibria

You tell me!

What is a coordination equilibrium? How does it differ from a Nash equilibrium?

A strategic profile s is a Nash equilibrium if no player strictly prefers to change **her own** action if other players actions are held fixed.

A strategic profile s is a Nash equilibrium if no player strictly prefers to change **her own** action if other players actions are held fixed.

Lewis: A strategic profile s is a coordination equilibrium if if no player strictly prefers that **some** player (potentially himself) changes actions, if all others' actions are held fixed.

Lewis' Rough Definition: A convention is a regularity R in the behavior of members of population P when they are in recurrent situation S if, in any instance S among members of P,

- Everyone conforms to R
- Everyone expects everyone else to conform to R
- Everyone prefers that everyone conform to R on condition that others do in S, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S.

Lewis' Rough Definition: A convention is a regularity R in the behavior of members of population P when they are in recurrent situation S if, in any instance S among members of P,

- Everyone conforms to R
- Everyone expects everyone else to conform to R
- Everyone prefers that everyone conform to R on condition that others do in S, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S.

We'll come back to modifications to this definition in a moment . . .

There can be more than one possible convention:

There can be more than one possible convention:

- In fact, generally something can count as "conventional" only if another convention could have been adopted.

There can be more than one possible convention:

- In fact, generally something can count as "conventional" only if another convention could have been adopted.
- Lewis wants to define a convention in terms of some type of strategic profile that may not be unique.

There can be more than one possible convention:

- In fact, generally something can count as "conventional" only if another convention could have been adopted.
- Lewis wants to define a convention in terms of some type of strategic profile that may not be unique.
- The Stag Hunt that there may be multiple coordination equilibria in a game: this is what Lewis wants!

But the non-uniqueness of coordination equilibria raises the other two other questions about convention at the outset:

- Evolution:: How does a convention come about? That is, why does one convention rather than another come to be?
- Stability: How is it maintained? That is, why does one convention persist if another could be adopted?

Let's start with the stability question.

Nash equilibria, in general, are not necessarily stable.

- Suppose we will play a stag hunt.
- In the past, you and I have both played stag (creating a Nash equilibrium), but,
- I have good reason to suspect you will play hare next.

What should I do?

Moral: If I don't expect you to play part of a Nash equilibrium, it might be rational for me to do something else.

Suppose, now, I expect you to play Stag.

What should I do?

In general, suppose we are playing some game and I expect you to play your half of a Nash equilibrium.

What should I do?

In general, suppose we are playing some game and I expect you to play your half of a Nash equilibrium.

What should I do?

Answer: Play the other half. By definition, a Nash equilibrium is one in which each player performs a best response to all others.

Definition of Convention

Lewis' Rough Definition: A convention is a regularity R in the behavior of members of population P when they are in recurrent situation S if, in any instance S among members of P,

- Everyone conforms to R
- Everyone expects everyone else to conform to R
- Everyone prefers that everyone conform to R on condition that others do in S, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S.

Evolution of Convention

But how do conventions come about?

In economists' terms, how is the equilibrium **selected**?

> *Agreement, salience, or precedent, we have seen, can solve a coordination problem by producing a system of concordant first- and higher-order mutual expectations.*

Lewis [2008], pp. 52.

Watch the "poison scene" from the Princess Bride.

Remember this game from a few weeks back?

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

Remember this game from a few weeks back?

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

I asked you questions like

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

Moral: Higher-order knowledge helped agents to select an equilibrium.

In previous classes, our theorems only assumed common-knowledge of the rationality of players, Moral: Higher-order knowledge helped agents to select an equilibrium.

In previous classes, our theorems only assumed common-knowledge of the rationality of players,

But there might also be common knowledge of facts about which actions (several of which might be rational) that agents may choose.

Definition of Convention

Lewis' Definition: A convention is a regularity R in the behavior of members of population P when they are in recurrent situation S if, in any instance S among members of P, it is common knowledge that

- Everyone conforms to R
- Everyone expects everyone else to conform to R
- Everyone prefers that everyone conform to R on condition that others do in S, since S is a coordination problem and uniform conformity to R is a proper coordination equilibrium in S.

Where does common knowledge come from?

Lewis gives three sources:

- Agreement
- Salience
- Precedent

Can these three sources explain the establishment of linguistic convention?

Skyrms' criticism of the three sources of common knowledge

- Agreement:

Skyrms' criticism of the three sources of common knowledge

- Agreement: Lewis admits this cannot be used to explain the origin of language.

Skyrms' criticism of the three sources of common knowledge

- Agreement: Lewis admits this cannot be used to explain the origin of language.
- Precedent: Puts the cart before the horse

Skyrms' criticism of the three sources of common knowledge

- Agreement: Lewis admits this cannot be used to explain the origin of language.
- Precedent: Puts the cart before the horse
- Salience:
  - The coordination equilibria representing meaningful communication are not salient.

Skyrms' criticism of the three sources of common knowledge

- Agreement: Lewis admits this cannot be used to explain the origin of language.
- Precedent: Puts the cart before the horse
- Salience:
    - The coordination equilibria representing meaningful communication are not salient.
    - In any event, salience is not necessary for equilibrium selection if agents can learn.

Outline

Example: Vervet signals

Example: Vervet signal

- Sender: A vervet monkey who sees a predator. He or she can
    - "Cough"
    - "Chutter", or
    - "Bark"

Example: Vervet signal

- Sender: A vervet monkey who sees a predator. He or she can
    - "Cough"
    - "Chutter", or
    - "Bark"
- Receiver: Another vervet monkey, who has not seen the predator yet.

Example: Vervet signal

- Sender: A vervet monkey who sees a predator. He or she can
  - "Cough"
  - "Chutter", or
  - "Bark"
- Receiver: Another vervet monkey, who has not seen the predator yet.
- They both want each other to evade the predator, but . . .

Example: Vervet signal

- Sender: A vervet monkey who sees a predator. He or she can
    - "Cough"
    - "Chutter", or
    - "Bark"
- Receiver: Another vervet monkey, who has not seen the predator yet.
- They both want each other to evade the predator, but ...
- Clearly, vervets did not schedule a meeting in which they decided that "cough" means that an eagle is approaching.

Signaling games also have two players: sender and receiver.

- Sender observes some state of the world (e.g., bsnake, eagle, or leopard).
- She then sends a signal to receiver (e.g., cough, chutter, or bark).
- The receiver then chooses an action (e.g. stand tall and back away, take cover in underbrush, scale a tree)
- The payoff that both receive depends upon the world and the receiver's action, e.g.,
    - Snakes are evaded by standing tall and backing away
    - Leopards are evaded by climbing trees, and
    - Eagles are evaded by hiding in the underbrush

Liste to Vervet monkey calls . . .

Formally, in cooperative signaling games:

- There are finite sets of states of the world W, a finite number of signals S, and finitely many actions A.
- Nature's "Actions": Probability distributions over worlds W
- Sender's actions: A function from worlds W to signals S.
- Receiver's actions: Functions from signals to acts.
- The payoffs to sender and receiver are the same, and they are determined by the state of the world and the action taken by the receiver.

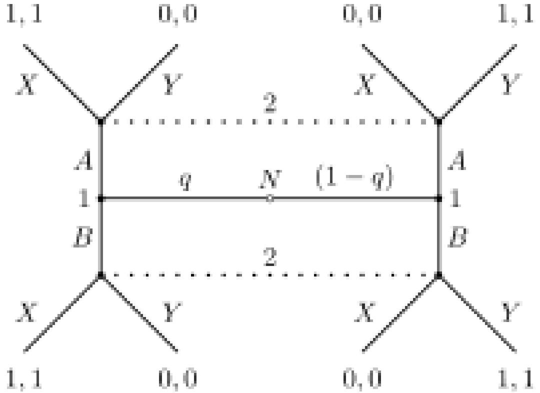
1, 1
0, 0
0, 0
1, 1
X
Y
X
Y
2
A
A
q
N
(1 − q)
1
1
B
B
2
X
Y
X
Y
1, 1
0, 0
0, 0
1, 1

- Assumption: For each state of the world w, there is at least one action a_w that is optimal.
- A signaling system is a pair of strategies $\langle f, g\rangle$ such that $g(f(w)) = a_w$ for all worlds w.
- **Question:** Is a signaling system a Nash equilibrium? A coordination equilibrium?

A signaling system is a strategic profile in a Lewis signaling game representing meaningful exchange of information

Just as cooperation, trust, etc. were represented by particular strategic profiles in prisoners' dilemmas, stag hunts, etc.

Just as we developed ABMs to model the evolution of cooperation, trust, etc., we can now do the same with signaling,

- **Population Model:**

Just as we developed ABMs to model the evolution of cooperation, trust, etc., we can now do the same with signaling,

- **Population Model:** Replicator Dynamics
- **Network Models:**

Just as we developed ABMs to model the evolution of cooperation, trust, etc., we can now do the same with signaling,

- **Population Model:** Replicator Dynamics
- **Network Models:**
  - Lattice, Small-Worlds, Bounded Degree, and Dynamic

Outline

1. Review and Preview
2. Convention
   - Stability
   - Evolution/Selection
3. Signaling Games
4. Evolution of Signaling
   - Replicator Dynamics
   - ABMs
5. Netlogo
6. References

What is the replicator dynamics?

Describe a model that produces the changes in population frequencies described by the replicator dynamics.

What is the replicator dynamics?

Answer: It's the following equation.

- Let $p_t(s)$ be the proportion of individuals in the population employing strategy s at time t.
- Let $F_t(s)$ be the fitness of the strategy, which is (recall) its expected payoff when paired with a random other agent from the population.
- Let $F_t(\textsc{ave})$ be the average fitness of all strategies.

$$p_{t+1}(s) = p_t(s) \cdot \frac{F_t(s)}{F_t(\textsc{ave})}$$

Describe a model that produces the changes in population frequencies described by the replicator dynamics.

Answer: Consider a very large population of individuals, each with one of several possible strategies.

- On each stage t_1, t_2, and so on, each individual in the population plays exactly one other player.
  - Pairs of players are chosen at random.
- They produce a number of offspring equal to their payoffs in the game.

But how are the roles of sender and receiver chosen in the replicator dynamics for signaling games?

But how are the roles of sender and receiver chosen in the replicator dynamics for signaling games?

Answer: It doesn't really matter in simple cases:

> *Evolutionary dynamics could operate on one population of senders and another of receivers as in some cases of interspecies communication, or it could operate on a single population, where individuals sometimes find themselves in the role of sender and sometimes in the role of receiver.*

[Skyrms, 2010], pp. 10.

Simple case: two equally probable states, two Signals $\Rightarrow$ The population moves towards a signaling system.

- We can also use ABMs to study the evolution of signaling.

Two Types of Learning

- We can also use ABMs to study the evolution of signaling.
- However, an important distinction arises because in signaling games, players do not choose simultaneously.

Two Types of Learning

- We can also use ABMs to study the evolution of signaling.
- However, an important distinction arises because in signaling games, players do not choose simultaneously.
- This is why Skyrms' distinguishes between "learning actions" and "learning strategies."

Recall, a (receiver's) strategy in a signaling game is plan (i.e. function) consisting of conditionals of the form "If I see signal *s*, I will choose act *a*" for each possible state *s*.

If I am unsuccessful when I chose action a after seeing signal s, I could change either

- **Learning Actions:** Only the part of my plan about how I should respond to signal s.

Two Types of Learning

If I am unsuccessful when I chose action a after seeing signal s, I could change either

- **Learning Actions:** Only the part of my plan about how I should respond to signal s.
- **Learning Strategies:** Several parts of my plan, including how I might respond to signals other than s.

Two Types of Learning

The two ways of learning are plausible (or implausible) in different contexts.

Example: Bacteria likely only change how they respond to chemicals one at a time. They learn actions.

A second example in which learning strategies is implausible: Imitation rules.

- When agents played simultaneous games like a Prisoners' dilemma, imitating one's neighbor meant imitating "Cooperate" or "Defect." One could easily imitate an entire strategy.

Two Types of Learning

A second example in which learning strategies is implausible: Imitation rules.

- When agents played simultaneous games like a Prisoners' dilemma, imitating one's neighbor meant imitating "Cooperate" or "Defect." One could easily imitate an entire strategy.
- Consider now a signaling game on a network.

A second example in which learning strategies is implausible: Imitation rules.

- When agents played simultaneous games like a Prisoners' dilemma, imitating one's neighbor meant imitating "Cooperate" or "Defect." One could easily imitate an entire strategy.
- Consider now a signaling game on a network.
- Suppose I see you (my neighbor) play a signaling game successfully in which you respond to signal s with action a.

Two Types of Learning

A second example in which learning strategies is implausible: Imitation rules.

- When agents played simultaneous games like a Prisoners' dilemma, imitating one's neighbor meant imitating "Cooperate" or "Defect." One could easily imitate an entire strategy.
- Consider now a signaling game on a network.
- Suppose I see you (my neighbor) play a signaling game successfully in which you respond to signal s with action a.
- It seems that I should only be able to imitate how you respond to signal s; I cannot imitate your entire strategy because I may not have seen how you behave in other circumstances!

Two Types of Learning

Nonetheless, there are circumstances in which humans clearly "learn strategies," in Skyrms' sense.

- Suppose you learn that the correct answer to (the signal) "Is $5 > 3$?" is (the act of asserting) "Yes."
- Then you'll likely update your disposition to answer (the signal) "Is $5 < 3$?" with the answer "No."
- In this case, you've updated your response to one signal given your response to another different signal.

Future Classes: We'll talk about how individuals in ABMs might learn actions and learn strategies.

Programming Concepts: Debugging and the Behaviorspace

References I

Lewis, D. (2008). *Convention: A philosophical study*. Wiley-Blackwell.

Skyrms, B. (2004). *The stag hunt and the evolution of social structure*. Cambridge University Press.

Skyrms, B. (2010). *Signals: Evolution, learning, & information*. Oxford University Press.