Game Theory and Evolution

Toban Wiebe

University of Manitoba

July 2012

Why does evolution need game theory?

In what sense does evolution optimize organisms?

- Example: the height of trees in the rainforest
- From an engineer's perspective, this height is wasteful

Source: http://hdw.eweb4.com

Why does evolution need game theory?

Or consider the peacock's tail:

Source: http://piterxpippin.deviantart.com

Why does evolution need game theory?

From a game theoretical perspective, however, trees and peacocks are playing optimal strategies given that they are all competing amongst themselves for sunlight and mates, respectively.

Competition is everywhere in nature, e.g., for resources, for mates, between predator and prey, between parasite and host, etc.

Evolution is driven by competition. For this reason, the engineering conception of optimality is inadequate. Game theory, however, is well suited for modeling evolution.

The Hawk-Dove Game

Two players contest a resource of value v. Each can play one of two strategies:

- Hawk (H): fight
- Dove (D): don't fight

The probability of winning a fight is $\frac{1}{2}$ and the cost of defeat is c. If two doves enter a contest, they share the resource equally. The *payoffs* can be represented in matrix form:

$$\begin{array}{ccc} & H & D \\ H & \frac{1}{2}(v-c) & v \\ D & 0 & \frac{1}{2}v \end{array}$$

Definitions

A *game* is defined by:

- Players
- Structure
- Actions
- Payoffs: each player has a real-valued payoff function π defined over the set of all possible outcomes.

Definitions (cont'd)

A *pure strategy* is a complete game plan; it specifies an action for each possible decision node for a given player.

A *mixed strategy* is a probability distribution over a strategy set. In other words, a player using a mixed strategy chooses a pure strategy at random, according to some probability distribution.

An *outcome* is a vector composed of the strategies chosen by each player.

Nash Equilibrium

A *Nash Equilibrium* is an outcome such that no player could have obtained a greater payoff by unilaterally switching strategies, or equivalently, each player's strategy is a 'best response' to the strategies of all the other players.

Formally, (σ_1^*, σ_2^*) is a Nash Equilibrium iff for $i, j = 1, 2,\ i \neq j$ and for every mixed strategy σ_i,

$$\pi_i(\sigma_i^*, \sigma_j^*) \geq \pi_i(\sigma_i, \sigma_j^*).$$

Note that Nash Equilibria are not necessarily unique. For example, the following game has 3 equilibria:

	L	R
L	1	0
R	0	1

The Prisoners' Dilemma

Two partners in crime are arrested and interrogated simultaneously in separate rooms. Each is offered the following deal:

- Both confess: 2 year sentence
- Both don't confess: 1 year
- If one confesses and the other doesn't: confessor goes free and the other gets a 3 year sentence

	C	NC
C	-2	0
NC	-3	-1

The optimal strategy is to confess, and so the only Nash Equilibrium is (C, C). Why? Notice that no matter what the other player does, one is always better off confessing.

Solving the Hawk-Dove Game

	H	D
H	$\frac{1}{2}(v-c)$	v
D	0	$\frac{1}{2}v$

If $v > c$, this becomes equivalent to a Prisoners' Dilemma and (H, H) is the equilibrium outcome.

If $v < c$, there are two pure strategy equilibria: (H, D) and (D, H). There is also a mixed strategy equilibrium. These can be solved for by finding the *best response correspondences*.

Solving the Hawk-Dove Game (cont'd)

	H	D
H	$\frac{1}{2}(v-c)$	v
D	0	$\frac{1}{2}v$

For $p_j \in [0,1]$, let $\sigma_j = (p_j, 1-p_j)$ be an arbitrary mixed strategy for player j, where p_j denotes the probability of playing H. Then player i's best response can be determined by comparing payoffs:

$$\begin{aligned} \pi_i(H,\sigma_j) &\gtreqless \pi_i(D,\sigma_j) \\ \Leftrightarrow p_j\frac{1}{2}(v-c) + (1-p_j)v &\gtreqless (1-p_j)\frac{1}{2}v \\ \Leftrightarrow p_j &\lesseqgtr \frac{v}{c}. \end{aligned}$$

Solving the Hawk-Dove Game (cont'd)

Given p_j, the best response correspondence for player i is

$$BR_i(p_j) = \begin{cases} \{0\}, & \text{if } p_j > \frac{v}{c} \\ [0,1], & \text{if } p_j = \frac{v}{c} \\ \{1\}, & \text{if } p_j < \frac{v}{c} \end{cases}.$$

(σ_1, σ_2) is a Nash Equilibrium iff

$$p_1 \in BR_1(p_2) \text{ and } p_2 \in BR_2(p_1).$$

There are 3 such solutions: $(0, 1)$, $(1, 0)$, $(\frac{v}{c}, \frac{v}{c})$.

Evolutionary Game Theory

In an evolutionary game, players are no longer rational, but blindly follow inherited strategies. However, equilibria will still consist of optimal strategies because any suboptimal strategy would be eliminated through natural selection.

Payoffs can be interpreted as *fitness* in the biological sense.

The term *mutant strategy* will refer to a strategy that differs from the equilibrium strategy which is adopted by a small subpopulation. In other words, this mutant population is a group of organisms that have inherited a genetic mutation.

Evolutionarily Stable Strategies

Intuitively, an *Evolutionarily Stable Strategy (ESS)* is a strategy such that, if the entire population adopts it, then no mutant strategy could invade the population through natural selection.

A *population profile* $\mathbf{x}$ is a vector of probabilities with which each pure strategy is played in the population. $\mathbf{x}_\varepsilon$ will denote a population profile in which the entire population uses a strategy σ^* except for a small proportion ε using a mutant strategy σ.

Formally, a mixed strategy σ^* is an *ESS* iff for sufficiently small $\varepsilon > 0$ and $\forall \sigma \neq \sigma^*$

$$\pi(\sigma^*, \mathbf{x}_\varepsilon) > \pi(\sigma, \mathbf{x}_\varepsilon).$$

A Sex Ratio Game

Assumptions:

- The proportion of males in the population is μ.
- Females mate once, produce k offspring, and die.
- As a result, males mate on average $\frac{1-\mu}{\mu}$ times, and die.
- The sex of the offspring is determined by the female's genetics from two pure strategies: produce only males (M) or only females (F).

Since a population profile $\mathbf{x} = (x, 1-x)$ produces a sex ratio $\mu = x$, it can be written as $\mathbf{x} = (\mu, 1-\mu)$.

Payoffs are given by the number of grandchildren:

$$\pi(M, \mathbf{x}) = k^2 \left(\frac{1-\mu}{\mu} \right)$$
$$\pi(F, \mathbf{x}) = k^2.$$

A Sex Ratio Game (Cont'd)

Comparing payoffs yields one Nash Equilibrium: $\sigma^* = (\frac{1}{2}, \frac{1}{2})$.

To check if σ^* is an *ESS*, consider a mutant mixed strategy $\sigma = (p, 1-p)$, $p \neq \frac{1}{2}$. Then

$$\begin{aligned} \mathbf{x}_\varepsilon &= (1-\varepsilon)\sigma^* + \varepsilon\sigma; \\ \text{and so, } \mu_\varepsilon &= (1-\varepsilon)\frac{1}{2} + \varepsilon p \\ &= \frac{1}{2} + \varepsilon\left(p - \frac{1}{2}\right). \end{aligned}$$

$$\begin{aligned} \pi(\sigma^*, \mathbf{x}_\varepsilon) &= k^2\left[\frac{1}{2}\left(\frac{1-\mu_\varepsilon}{\mu_\varepsilon}\right) + \frac{1}{2}\right] \\ \pi(\sigma, \mathbf{x}_\varepsilon) &= k^2\left[p\left(\frac{1-\mu_\varepsilon}{\mu_\varepsilon}\right) + (1-p)\right] \end{aligned}$$

A Sex Ratio Game (Cont'd)

The *ESS* condition is that $\pi(\sigma^*, \mathbf{x}_\varepsilon) - \pi(\sigma, \mathbf{x}_\varepsilon) > 0$ for small ε.

Simplifying this difference yields

$$\pi(\sigma^*, \mathbf{x}_\varepsilon) - \pi(\sigma, \mathbf{x}_\varepsilon) = k^2 \left(\frac{1}{2} - p\right)\left(\frac{1 - 2\mu_\varepsilon}{\mu_\varepsilon}\right).$$

- $p < \frac{1}{2}$ produces $\mu_\varepsilon < \frac{1}{2}$
- $p > \frac{1}{2}$ produces $\mu_\varepsilon > \frac{1}{2}$

The *ESS* condition holds in both cases.

The Evolutionary Hawk-Dove Game

It can be shown that the *ESS* condition is equivalent to following condition for a pairwise contest game:

$$\pi(\sigma^*, \sigma^*) > \pi(\sigma, \sigma^*), \text{ or}$$
$$\pi(\sigma^*, \sigma^*) = \pi(\sigma, \sigma^*) \text{ and } \pi(\sigma^*, \sigma) > \pi(\sigma, \sigma).$$

Recall that $\sigma^* = (\frac{v}{c}, 1 - \frac{v}{c})$ is the only strategy that produces a symmetric Nash Equilibrium (σ^*, σ^*). Since it is a mixed strategy,

$$\pi(\sigma^*, \sigma^*) = \pi(\sigma, \sigma^*).$$

Let $\sigma = (p, 1 - p)$, $p \neq \frac{v}{c}$. Then, simplifying the following difference yields the *ESS* condition:

$$\pi(\sigma^*, \sigma) - \pi(\sigma, \sigma) = \frac{c}{2}\left(\frac{v}{c} - p\right)^2 > 0.$$

Adding Dynamics to The Model

Evolutionary dynamics can be analyzed by adding differential equations to the model. Consider a system of differential equations where n_i denotes the number of organisms using the strategy s_i.

$$n_i' = n_i \pi(s_i, \mathbf{x})$$

This can be written equivalently in terms of population proportions x_i and the average payoff $\bar{\pi}(\mathbf{x}) = \sum x_i \pi(s_i, \mathbf{x})$:

$$x_i' = x_i[\pi(s_i, \mathbf{x}) - \bar{\pi}(\mathbf{x})]$$

These systems describe the evolution of the population. When there are multiple *ESS*, the one that is attained depends on the initial population configuration.

Dynamics and Evolutionary Stability

A population in equilibrium maintains constant proportions of each strategy: $x_i' = 0, \forall i$. However, this is not a sufficient condition for evolutionary stability (the analog of an *ESS*).

Evolutionary stability requires a further stability criterion: that any small deviation from the equilibrium state is eliminated over time.

Dynamics in The Hawk-Dove Game

A population profile $\mathbf{x} = (x_1, x_2)$ can be written $\mathbf{x} = (x, 1 - x)$, where $x \equiv x_1,\ x_2 = 1 - x$, and $x_2' = -x_1'$. So the dynamics are captured by a single differential equation

$$x' = x(\pi(H, \mathbf{x}) - \bar{\pi}(\mathbf{x})).$$

Substituting $\bar{\pi}(\mathbf{x}) = x\pi(H, \mathbf{x}) + (1 - x)\pi(D, \mathbf{x})$, the differential equation simplifies to

$$\begin{aligned} x' &= x(1 - x)(\pi(H, \mathbf{x}) - \pi(D, \mathbf{x})) \\ &= x(1 - x)(v - cx) \end{aligned}$$

The equilibria of this system are points satisfying $x' = 0$:

$$x^* = 0,\ x^* = 1,\ \text{and } x^* = \frac{v}{c}.$$

Dynamics in The Hawk-Dove Game (cont'd)

To check for the evolutionary stability of $x^* = 0$, let $\varepsilon > 0$ and let $x = x^* + \varepsilon = \varepsilon$. Then $x' = \varepsilon'$ and

$$\varepsilon' = \frac{1}{2}\varepsilon(1-\varepsilon)(v - c\varepsilon)$$

For small ε, $\varepsilon' \approx \frac{1}{2}v\varepsilon$. The solution $\varepsilon(t) = \varepsilon_0 e^{\frac{1}{2}vt} \nrightarrow 0,\ t \to \infty$, and so $x^* = 0$ is not evolutionarily stable (the mutant population grows).

Similarly, for $x^* = 1$, $\varepsilon(t) = \varepsilon_0 e^{-\frac{1}{2}(v-c)t} \nrightarrow 0,\ t \to \infty$. So $x^* = 1$ is not evolutionarily stable either (recall: $v < c$).

For $x^* = \frac{v}{c}$, let $\varepsilon \in \mathbb{R}$ to account for both H and D mutants. Then, for small ε, $\varepsilon(t) = \varepsilon_0 e^{-\frac{1}{2}v(1-\frac{v}{c})} \to 0,\ t \to \infty$. Thus, $\mathbf{x} = (\frac{v}{c}, 1 - \frac{v}{c})$ is evolutionarily stable.

Further Development and Application

Extending the model:

- Repeated games and spacial dynamics
- Payoff (fitness) functions: dependence on population density and resource scarcity
- Evolutionary stability: account for stability of population size, invasion by combinations of mutant strategies

Evolutionary game theory can apply to anything evolved: from viruses to humans.

Evolutionary game theory has been quite successful for understanding evolved strategies.

That's all, folks!

Email: tobanw@gmail.com

Download this presentation (PDF): `http://goo.gl/NhR82`

Recommended:

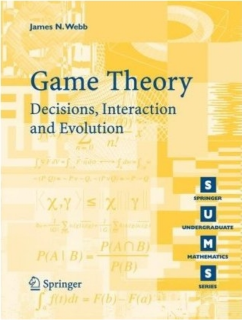

Game Theory: Decisions, Interaction and Evolution.
James Webb: Springer, 2007.