Games Digital Animals Play

Vinny Just

Mathematical Biosciences Institute and Department Mathematics, Ohio University

October 11, 2005

Examples of Games

Chess, bridge, the dating game.

Common features:

- two or more players *interact*
- benefits for each player depend on action of *all* players
- players may regret their actions with hindsight

Note that games may or may not involve chance events, and the interests of the players may be totally opposed (a zero-sum game), or in (partial or total) alignment.

Ingredients of a game

- a set of n *players*
- a set of possible *strategies* for each player
- a *payoff function* that assigns a payoff for each player to each n-tuple of strategies

Example: In the dating game, a set of strategies might be:

{show up for the date, don't show up}.

Note that players choose their strategies *simultaneously* and *independently.*

Rational Self-Interest

The basic assumption of game theory is that players only act out of *rational self-interest*, which means that their only goal is maximizing their own payoff. Thus niceness or spite are not considered in game theory, unless they are factored into the payoff function.

Nash Equilibria

An n-tuple of strategies in an n-player game is called a *Nash Equilibrium* if none of the players regrets their choice of strategy after learning about the strategies adopted by all other players.

Games in Economics

Payoffs are expected profits.

Examples of Strategies:

- buy, sell, or hold a stock
- enter a market or stay out
- start a prize war or collude

Games in Politics

The currency of payoffs is usually political power.

Examples of Strategies:

- invade or don't invade Iraq
- consult or ignore your allies
- respect or break the Geneva convention

Games in Animal Behavior

Payoffs are expected number of offspring.

Examples of Strategies:

- in a contest over a resource, attack your opponent, retreat, or display
- graze on your neighbor's beautiful green patch or forage elsewhere
- guard your offspring or seek new mating opportunities

How about Rational Self-Interest?

Animals (including *H. sapiens*) of course cannot be presumed to act out of *rational* self-interest. This, incidentally, somewhat limits the usefulness of game-theoretic models in economics and political science.

However, if payoffs are expected number of offspring, then biological evolution will produce behavior that is consistent with this assumption.

How about Nash Equilibria?

Evolution's way of showing regret is producing a mutant that does better than the wild type. In evolutionary game theory, we are looking not just for any Nash Equilibrium, but for a strategy that when adopted by an infinite population is immune to invasion by mutants. Such a strategy is called an *evolutionarily stable strategy,* or *ESS* (Maynard Smith and Price, 1973).

Example: The Hawk-Dove Game

This game due to Maynard Smith and Price (1973) models a contest over a resource (e.g., food, territory) of value V.

Strategies: {Hawk, Dove}

Payoffs: Two Hawks always fight for the resource and incur an expected net cost C due to the fight; two Doves share the resource for a net gain of $V/2$ each; a Dove will retreat if encountering a Hawk, receiving a payoff of 0 while the Hawk gains V.

Problem: A population of Doves can be invaded by a mutant Hawk, a population of Hawks can be invaded by a mutant Dove.

Mixed-Strategy ESS's

While neither Hawk nor Dove is a *pure-strategy* ESS in this game, it can be shown that a population of players that behave like Hawks with probability $\frac{V}{V+2C}$ and like Doves with probability $\frac{2C}{V+2C}$ is an ESS. Such an ESS is called a *mixed-strategy ESS.* A population that contains a proportion of $\frac{V}{V+2C}$ pure Hawks and $\frac{2C}{V+2C}$ pure Doves is also evolutionarily stable; in the latter case we say that the ESS is *realized by a polymorphism.*

Caveat: Additional strategies, such as Assessor, could invade this ESS.

Fighting Fish

In 1995, Morris *et al.* studied territorial contests of males in swordtail fish species *Xiphophorus nigrensis* and *X. multilineatus.* Contests always started with a display phase, followed either by a retreat of the smaller contestant, or by a fight that involved biting. Most, but not all, of the observed fights were won by the larger fish.

Question: Which fish usually delivered the first bite?

The Napoleon Complex

Answer: 78% of the observed fights were initiated by the smaller animal, and in 70% of the fights, the fish that delivered the first bite lost the contest.

What kind of model can explain these observations?

One (mildly complicated) model was presented in (Just and Morris, 2003, *Evolutionary Ecology*).

Can we develop a kind of "minimal model" for studying escalation?

Model Assumptions

This part of the talk describes joint work with Molly R. Morris and Xiaolu Sun.

Contests have two stages, a “display” stage, and an optional “fight” stage. We want to model the decision to either retreat without a fight (R), escalate to the fighting stage (E), or continue displaying and leave the initiative to the opponent (D). Let p be the probability of winning a fight. We assume that the decision whether to escalate to fighting should depend only on p and the expected fitness gain G from a fight, where

$$G = p \cdot (ResourceValue) - CostOfFight(p).$$

Four Classes of p

The action of a player should only depend on whether p is

- *very small* ($G < 0$)
- *small* ($G > 0,\ p < 0.5$)
- *large* (opponent's $G > 0$, $p > 0.5$)
- *very large* (opponent's $G < 0$)

Note that if p is *very small,* then unilateral retreat R is the only sensible action; in the three other classes a fight is preferable to unilateral retreat.

The Strategies

A strategy specifies the action that a player will perform in each class of p. Thus strategy REDD calls for retreat if p is *very small,* for escalation if p is *small,* and for continued display if p is *large* or *very large.*

Note that in a population of REDD, most fights will be started by their eventual losers, in a population of RDED, most fights will be started by their eventual winners, and in a population of REED, about half of the fights will be started by their eventual losers.

Overall there are $3^4 = 81$ possible strategies. One should not restrict the model to the three most "plausible" strategies REDD, REED, and RDED.

The Missing Ingredient

In the model sketched so far, there are no ESS's.

Note that we have assumed that player's actions depend on their (real) winning probability p, but animals can only act on what they perceive.

We need an extra parameter q, which is the probability of misperception of the class of p for a neighboring one.

Parameters of the Model

V - value of the resource

K - cost of losing a fight

L - cost of engaging in a fight

a - representative value of *very small* p

b - representative value of *small* p

q - probability of misperception

d - penalty for displaying indefinitely

Sampling the Parameter Space

How typical are parameter settings where REDD or an REDD/REED mix is an ESS?

How typical are parameter settings where RDED or an RDED/REED mix is an ESS?

We examined this question by randomly sampling 100,000 parameter settings drawn from a fairly generic region of the parameter space and testing each setting for pure-strategy ESS's and mixed-strategy ESS's with at most two pure-strategy components.

Results

For 85.5% of the sampled settings, either pure REDD or an REDD/REED mix was found to be an ESS.

In contrast, for only 3.2% of the sampled settings, either pure RDED or an RDED/REED mix was found to be an ESS.

Interpretations

Naive version: Most animal fights will be started by their likely losers.

But most of the relatively few studies of escalation in various species show that fights are more often initiated by their likely winners.

Interpretations

Pessimistic version: The model must be wrong.

Interpretations

More cautious version: Our "minimal model" is clearly too simple to adequately account for the factors that drive escalation in any specific species. The model should be considered as a kind of "null hypothesis" that shows that escalation by likely winners is far from obvious and needs explanations. The model can be extended in various ways to incorporate biological details relevant to different species. By pinpointing the different ways in which our model is "wrong" (or too simplistic) with regard to any given species, we can learn something about the biology of this particular species.

The ESS Concept Revisited

An ESS is a strategy that when adopted by an infinite population is immune to invasion by mutants.

Note that this definition says nothing about whether an ESS could evolve in a population that is not at equilibrium, or whether it is likely to be lost in finite populations due to random drift.

One can study these issues with *simulated evolution.*

Our Simulations

Our digital animals are strings of four letters from the alphabet {R, E, D}. We let populations of 3,000 of them evolve over 25,000 mating seasons. Fitness after each mating season is determined by the outcomes of about 10 simulated contests per animal. Parents for the next generation are chosen with probability proportional to fitness; strategies are inherited with crossover and mutation.

We start from random populations. Do predicted ESS's evolve? How robust are they under genetic drift?

Results of Simulations

For a parameter setting for which pure REDD was predicted as an ESS, a population that consisted of over 90% REDD did consistently evolve from random populations, and the ESS was robust under random drift.

For several parameter settings for which a mix of REED/REDD was predicted as an ESS, populations whose great majority played either REED or REDD did consistently evolve from random populations or RDED populations, and this fact was robust under random drift. *However,* the proportions of REED and REDD in the mix fluctuated wildly.

Simulating other Games

Our digital animals were little strings of symbols that prescribe actions in certain situations. Strategies in many other games can be described in this way, but the number of possible "situations" is usually much bigger than four.

A Problem: The longer our genomes get, the more mutational meltdown do we observe.

A Toy Game

Now I describe joint work with Fang Zhu (2003, *GECCO Workshop on Representation*).

Imagine a predator that feeds on n different prey species. Whenever the predator spots prey, he must decide whether to chase after it. The decision should depend on whether his probability of catching the prey is above a threshold t, where

$t =$ (energy lost in chase)/(energy gained from consuming prey)

Assume each prey species is characterized by one threshold. We want to study how evolution can optimize all these thresholds simultaneously.

A Toy Model

Imagine a game against Nature who randomly draws numbers i from $\{1, ..., n\}$ and probabilities p, and asks us whether we want to chase. If we don't chase, we get a payoff of zero. If we decide to chase, Nature will give us a net reward V_i with probability p and will make us pay a net cost C_i with probability $1 - p$.

A strategy in this game can be conceptualized as a string of probability thresholds t_i so that we decide to chase when Nature presents us with the pair (i, p) if, and only if, $p > t_i$.

A Toy Simulation

Fix n, let our digital animals be strings of n real numbers $g_i \in [0, 1]$ that represent candidates for the thresholds t_i. In each mating season, let them play a number of times against nature, choose parents with probability proportional to fitness, produce offspring by crossover and mutations (adding a little noise). Let populations of 1,000 of them evolve for 40,000 mating seasons. How close do we get (on average) to optimal fitness?

Sample Results

n	below optimal fitness
2	0.36%
3	0.47%
4	0.56%
5	0.63%
6	0.72%
20	2.42%

Conventional wisdom: Mutational meltdown is bad for evolution. The problem gets worse the larger the genome size and the higher the mutation rate.

The War of Attrition

The War of Attrition is a game-theoretic model of animal contests for a resource of value V that do *not* involve physical contact, but are decided by the length of time an animal is prepared to keep displaying.

Maynard Smith (1974) predicts a mixed-strategy ESS with probability density function of the maximal displaying times t given by:

$$g(t) = \frac{1}{V} e^{-t/V}.$$

For a population that adopts the ESS, the expected payoff from the game is zero.

Simulating the War of Attrition

Just and Zhu (2004, *Behavioural Processes*) simulated the evolution of strategies in this game. Our simulated populations consistently achieve an average **positive** payoff from the War of Attrition game, and this payoff is larger for larger mutation rates.

It appears that in this case **mutational meltdown is good!**

Mutational Meltdown Can Be Very Bad

But in other cases we found that it was difficult to achieve any meaningful evolution in the computer if our "genome sizes" increased. While this may be partially due to the small population sizes and high mutation rates needed in our simulations, simulation studies of the widely used *Sequential Assessment Game* show that it is highly dubious whether the predicted ESS in this game could evolve even under Nature's small mution rates and large population sizes (Just and Sun, 2004, *Proceedings of GECCO*).

How does Nature Deal with Mutational Meltdown?

- by inventing sex
- by keeping genome sizes small
- perhaps in other, currently not well-understood ways

Revisiting our Toy Simulation

We want to optimize n thresholds. Thus far we have represented each threshold t_i by a separate "gene" $g_i \in [0, 1]$. Let us call this "one-gene-one-trait" approach *separable coding.* What if we try to use only six "genes" and code twenty different thresholds t_i by taking (geometric) averages of combinations of three among these genes, *i.e.*

$$t_i = (g_{j_{i1}} g_{j_{i2}} g_{j_{i3}})^{1/3} \text{ ?}$$

Sample Results

n	coding	below optimal fitness
2	separable	0.36%
3	separable	0.47%
4	separable	0.56%
5	separable	0.63%
6	separable	0.72%
20	separable	2.42%
20	nonseparable	0.50%

Note that by reducing genome size to six, one might hope for results similar to the ones for separable coding of six thresholds with six genes, but we do much better!

Why?

Pleitropy

Conventional wisdom: Pleiotropy is a nuisance if we want to decode the workings of a genome. A human engineer would have designed organisms more neatly, on a one-gene-one-trait basis. But evolution is an opportunistic tinkerer who doesn't care about neat blueprints.

Moreover, pleiotropy may help Nature to keep genome sizes small.

But is this the whole story?

Understanding Genome Organization

Our experiments suggest that there may be other benefits of pleiotropy besides keeping the genome small.

Can we describe these benefits in terms of mathematical theorems?

If so, how do these translate into expectations that we should form about the genotype-phenotype map?

Can we study these questions by ignoring the higher levels of biological organization and concentrating on "phenotypes" that are properties of the organism's biochemical networks?