

Santa Fe Institute, Complex Systems Summer School 2008

Evolutionary Game Theory Tutorial¹

Part 1

Joshua Reyes

6 JUNE 2008

1. CLASSICAL GAME THEORY

Game theory is the study of interactive decision making. The classical theory studies the systematic, strategic play by so-called *rational* agents; that is, we assume that the players in classical context seek to optimize a quantity known as utility. Often utility can be thought in terms of money, but it need not be. Because optimization depends not only on what one agent chooses to do, but on what all the other players decide to do, game theoretic problems are, typically, hard to analyze. Before going to far, however, it's important to make precise just what a game is.

Rational agent

A *game* is a collection of rules and strategies. The rules specify who can do what and when, whereas the strategies describe what do to and when. In general strategies can be very complicated, or entirely arbitrary. In the select-a-meal game, you might employ a strategy that consists of spinning around three times, closing your eyes, and standing on one leg. According to this strategy, if the time it takes for you to lose balance is less than twenty seconds you take the comfort food option and the vegetarian offering otherwise. This strategy is both elaborate and arbitrary, and by the definition, perfectly valid. Strategies can be deterministic or probabilistic. As long as it completely specifies an legal action for each possible state of the game at hand, it's valid.

Game
Strategy

In order to study the structure of games, it is necessary to be able to write it down. Since a game really reduces to a series of decisions, it makes sense to write each possible choice down, one at a time, in the order dictated by the rules. This sort of representation lends itself nicely in the form of a (graph theoretic) tree.

Normally we think of the tree as being rooted at some node that represents what happens just before the game begins. The first move of the game made by a player is usually decided by an "act of God." Often the internal nodes carry labels with either the name or number of the player whose turn it is to move during that state of the game. Directed edges point to the state of the game that results in the action taken at the node. Terminal nodes carry additional information. The value of *payoff* at the end of the game along with the name of the recipient(s) to receive the prize. A walk from the root to a terminal node is called a *play* of the game. You could think of a play as simply a single match (or instantiation or whatever makes most sense).

Well, there's nothing that really requires the structure of an arbitrary game to be a tree. In fact, it need not even be an acyclic directed graph. Loops are fine. But it doesn't hurt to think of games essentially as trees.

Payoff
Play

The resultant graph structure is the *extensive form* of the game. While the extensive form of the game encodes the entire history, even for small games the size of the tree can be unwieldy or even deceptive. For this reason, people have come up with an alternative description of games.

Extensive form

Any new, more compact representation must include the following information:

- (1) The number of players in the game,

¹All of the notes from this tutorial have been lifted almost entirely from lectures given by Tim Killingback at the University of Massachusetts Boston during a course on game theory, ecology, and evolution in the spring of 2008.

The strategy set of one player may not be the same as the others. In a play of Marco Polo, for example, only one player can shout “Marco!” whereas everyone else yells “Polo!”

Bimatrix Normal form

You can pretend that each strategy has been meticulously drawn up and housed in a very safe place, so that if someone asked to you play the game in real life you could. But since we’re modeling behavior, and playing the game for real takes a lot of time, it’s best left someone where else.

Asymmetric game

Exercise. Fill in the corresponding matrix with the appropriate payoffs. Remember that $\pi_1(x, y)$ is the payoff to PLAYER I for playing x when PLAYER 2 chooses strategy y , and similarly for $\pi_2(x, y)$.

	C	D	E
A	*	*	*
B	*	*	*

Symmetric game

Yes, for this to work the sets of strategies available to each player have to be identical.

- (2) A complete description of all the strategies available to each player, and
- (3) A catalogue of payoffs that accrue to each player using each strategies against any other player using any of his strategies.

For two-player games, and those are the only kind we’re going to consider in this tutorial, we can package everything together in a *bimatrix*. In this representation, the *normal* or *strategic form* fo the game, the rows list the strategies of PLAYER I, and the columns gives the strategies of PLAYER II.

In this case we label the strategies available to each player with names like A or B, and leave the details of their exact specification out of the picture. For this description, only the fact that the strategies exist and that we know their payoffs matter. The payoff values, an ordered pair of numbers, make up the entries of the bimatrix. Within the pair, the payoff to PLAYER I comes first; the payoff to PLAYER II, second. For example, say the rows are labeled with strategies A and B for PLAYER I and the columns by C, D, E for PLAYER II. Then the pair $(\pi_1(A, D), \pi_2(A, D))$ consisting of the payoffs of an A-strategist (PLAYER I) against a D-strategist $\pi_1(A, D)$ and the payoff of a D-strategist (PLAYER II) against an A-strategist $\pi_2(A, D)$ appears in the Ath row and Dth column of the payoff matrix.

There is no reason to assume in advance that the two payoffs are equal. Generally the payoffs will not equal each other. That is, if PLAYER I adopts some strategy x and PLAYER II adopts another strategy y , then the payoffs the two players receive need not be equal:

$$\pi_1(x, y) \neq \pi_2(x, y).$$

Such games are called, surprisingly enough, *asymmetric games*.

2. TWO-PLAYER, TWO-STRATEGY GAMES

For the rest of this section, we are going to investigate the behavior of three classes of games parametrized by the following payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \left[\begin{array}{cc} (\frac{1}{2}(a-b), \frac{1}{2}(a-b)) & (a, 0) \\ (0, a) & (\frac{1}{2}(a+c), \frac{1}{2}(a+c)) \end{array} \right]. \end{array}$$

Here a , b , and c are just numbers that we’ll fill in as time goes by. You may notice that the payoffs, in this case, are symmetric. A games whose corresponding payoff matrix is symmetric is itself called a *symmetric game*. We’ll discuss some simplifications we can make for symmetric games in a bit more detail later on but for now let’s just plug in some values and play some games.

2.1. Prisoner’s Dilemma ($a = 4, b = 2, c = 0$). In our first example, plugging in the parameters yields the following pay-off matrix

$$\begin{array}{cc} & \begin{array}{cc} A & B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \left[\begin{array}{cc} (1, 1) & (4, 0) \\ (0, 4) & (2, 2) \end{array} \right]. \end{array}$$

Before you read on, go find someone else and play the game once. Remember that PLAYER I chooses a strategy x from one of the rows, and PLAYER II selects a strategy y from one of the columns. Also, the payoff to PLAYER I is the first number in the pair where the row x and column y meet—denoted by the symbol $\pi_1(x, y)$ —and the payoff to PLAYER II is the second number, denoted by $\pi_2(x, y)$.

In this case strategy A *dominates* B; that is, no matter what your opponent chooses, you obtain a higher payoff if you choose A. So it makes sense for both players to choose A. If classical game theory were my lawyer advising me to plead strategy A or B in a court of law, it'd surely tell me to go with A. We'll say some more about that a little bit later.

2.1.1. *The Story.* All of these games come with a funny little narrative. This one is called Prisoner's Dilemma, and its story goes something like this:²

Flávia and Petr several years ago had planned to rob a bank. Unfortunately for them, the robbery was not as clean as they would have hoped, and quickly found themselves engaged in a high-speed chase covering an excess of 400 miles over the rural roads of Nebraska. After several hours and considerable damage to the local corn fields, the Nebraska state police, in a singular act of courage and quick maneuverability, surrounded the car that contained the would-be robbers, and arrested them. Back at the police station, Flávia and Petr were placed in two separate interrogation rooms and made an identical offer.

"Flávia, here's the offer that we are making to both you and Petr. If you both hold out on us, and don't confess to bank robbery, then we admit that we don't have enough proof to convict you. However, we *will* be able to jail you both for one year, for reckless driving and endangerment of corn. If you turn state's witness and help us convict Petr (assuming he doesn't confess), then you will go free, and Petr will get twenty years in prison. On the other hand, if you don't confess and Petr does, then *he* will go free and *you* will get twenty years."

"What happens if both Petr and I confess?" asked Flávia.

"Then you both get ten years," said the interrogator.

Flávia, who had been a participant at the 2008 Santa Fe Institute Complex Summer School, reasoned this way: "Suppose Petr intends to confess. Then if I don't confess, I'll get twenty years, but if I do confess, I'll only get ten years. On the other hand, suppose Petr intends to hold out on the cops. Then if I don't confess, I'll go to jail for a year, but if I do confess, I'll go free. So no matter what Petr intends to do, I am better off confessing than holding out. So I'd better confess."

Naturally, Petr employed the very same reasoning. Both criminals confessed, and both went to jail for ten years. (Actually, they didn't go to jail. When they were in court, and heard that they had both turned state's witness, they strangled each other. But that's another story.) The police, of course, were triumphant, since the criminals would have been free in a year had both remained silent.

Our payoff matrix doesn't measure in terms of jail time, instead it uses some other utility. Even still the structure is the same as in the story: strategy A corresponds to FINK and strategy B corresponds to STAY MUM.

A = DEFECT
B = COOPERATE

2.2. Nash Equilibrium. The pair of strategy (A, A) is in some sense the only non-stupid move (that is, of course, unless you have some extra information about your opponent). Once both players choose this strategy neither can do better by *unilaterally* switching to some other strategy. This is likely true of any "good" pair of strategies that we would want to look for.

²This version of the Prisoner's Dilemma has been shameless adapted from a problem set that accompanies the *Structure and Interpretation of Programming Languages* by Sussman and Abelson on the book website <http://mitpress.mit.edu/sicp/psets/ps4prs/readme.html>; at least it did at 11:13AM, 7 June 2008.

A general pair of strategies (s, t) is good if

- (1) Strategy s is an optimal choice for PLAYER I given that PLAYER II will choose t ; s is a best response to t .
- (2) Strategy t is an optimal choice for PLAYER II given that PLAYER I will choose s ; t is a best response to s .

Nash equilibrium

If both properties hold simultaneously, then the pair (s, t) is called a *Nash equilibrium* (NE). We can cast our observations in a more mathematical language by noting that for all strategies \tilde{s} and \tilde{t} ,

- (1) s is a best response to t if and only if $\pi_1(s, t) \geq \pi_1(\tilde{s}, t)$,
- (2) t is a best response to s if and only if $\pi_2(s, t) \geq \pi_2(s, \tilde{t})$.

\tilde{s} is pronounced s-twiddles.

Strict Nash equilibrium

If the inequalities are strict, then the corresponding pair is called a *strict Nash equilibrium*. An arbitrary two-person game with finite strategies may have more than one NE. In that case, classical game theory does not give a clear-cut way of selecting among them. On the other hand, strict Nash equilibria are unique. (Why?)

Pure strategy

Mixed strategy

2.2.1. *Mixed Strategies.* Up until now we have been considering *pure strategies*, those strategies labeled on the sides of the payoff matrix. A *mixed strategy* σ chooses pure strategies s_1, \dots, s_n with probabilities p_1, \dots, p_n . Since the p_i really are probabilities we require that $p_i \geq 0$ for all i and that they sum to one. With these definitions, the pure strategy s_i is also the mixed strategy with $p_i = 1$ and all other $p_j = 0$.

Let $s = (s_1 \cdots s_n)^T$ and $p = (p_1 \cdots p_n)^T$, then their inner product is $\sigma = \langle s, p \rangle$.

It is common to write a mixed strategy as a linear combination of pure strategies

$$\sigma = p_1 s_1 + \cdots + p_n s_n = \sum_{i=1}^n p_i s_i.$$

Payoff for mixed strategies

We need to expand our notion of payoff to accommodate mixed strategies. To do so, we'll sum over the the payoffs of the pure strategies, because we already know how to do that, but weight them according to how often two strategies are (probably) played against each other. So, let $\sigma = \sum_{i=1}^n p_i s_i$ and $\tau = \sum_{j=1}^m q_j t_j$, then the payoff to PLAYER I in the game of σ against τ is

The product $p_i q_j$ measures about how often s_i plays against t_j .

I've turned the payoff bimatrix in one just for PLAYER I with the payoff function π_1 in that second equality.

$$\pi_1(\sigma, \tau) = \sum_{i=1}^n \sum_{j=1}^m p_i q_j \pi_1(s_i, t_j) = \langle p, \pi_1(s, t) \cdot q \rangle,$$

and similarly for PLAYER II.

Not only do we need to revise our concept of payoff, but also of Nash equilibrium.

Nash equilibrium for mixed strategies

We say that a pair of mixed strategies (σ, τ) is a *Nash equilibrium* if for any other (mixed) strategies $\tilde{\sigma}$ and $\tilde{\tau}$,

$$\pi_1(\sigma, \tau) \geq \pi_1(\tilde{\sigma}, \tau) \text{ and } \pi_2(\sigma, \tau) \geq \pi_2(\sigma, \tilde{\tau}).$$

A strict NE over mixed strategies is defined similarly.

The Fundamental Theorem of Classical Game Theory, sometimes known as Nash's Existence Theorem, promises that for any two-person game with a finite number of pure strategies, at least one Nash equilibrium exists in (possibly) mixed strategies.

Proposition 2.1. *The pair (A, A) is a Nash equilibrium of the Prisoner's Dilemma with the payoff matrix given above.*

Proof. We want to show that A is a best response to any other mixed strategy σ for both of the first player and the second player. (So there are two things to show here.) So let $\sigma = pA + (1-p)B$ for some probability $p \in [0, 1]$. If PLAYER I adopts σ then the payoff to her is

$$\begin{aligned}\pi_1(\sigma, A) &= p\pi_1(A, A) + (1-p)\pi_1(B, A) \\ &= p \cdot 1 + (1-p) \cdot 0 \\ &= p.\end{aligned}$$

The second equality comes directly from the payoff matrix.

$$\begin{array}{cc} & \begin{array}{c} A \\ B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{bmatrix} (1, 1) & (4, 0) \\ (0, 4) & (2, 2) \end{bmatrix} \end{array}$$

But since $p < 1$ and $\pi_1(A, A)$ equals 1, we've shown that A is a best response for PLAYER I against A. Identical reasoning establishes the equality for PLAYER II. \square

Because the inequalities were strict, we actually proved something much stronger.

Proposition 2.2. *The pair (A, A) is the unique Nash equilibrium in the Prisoner's Dilemma.*

2.3. Hawk-Dove Game. (a = 2, b = 4, c = 0). Again, go find a partner and have a shot at the game described by the payoff matrix

$$\begin{array}{cc} & \begin{array}{c} A \\ B \end{array} \\ \begin{array}{c} A \\ B \end{array} & \begin{bmatrix} (-1, -1) & (2, 0) \\ (0, 2) & (1, 1) \end{bmatrix} \end{array}$$

In this game there are (count'em, one, two) three Nash equilibria. Two of them include only pure strategies, (A, B) and (B, A). To find the third we need to consider mixed strategies as well. Indeed, once we do we discover that the strategy $(\frac{1}{2}, \frac{1}{2})$ also does the trick.

I'm being a little sloppy. A NE consists of a pair of strategies. But $(\frac{1}{2}, \frac{1}{2})$ describes on a single albeit mixed strategy. Call it σ^* . Then the real NE is (σ^*, σ^*) .

This class of games goes by a few different names. Sometimes it is called Chicken, the Snowdrift Game, and the Hawk-Dove Game. Each name reflects a different interpretation.

A = DON'T SWERVE
B = SWERVE

In the game of Chicken, Ari and Ben drive their cars straight on at each pedal-to-the-metal. If one of the players swerves, both escape with their lives but only one receives the glory. If both are bull-headed and refuse to swerve, then they collide at a fantastic speed and incur large medical bills as a result. If both of them swerve, then they split the glory evenly having showed equal amounts of courage (and stupidity) for playing the game at all.

The second interpretation also involves cars. Imagine that you live in a cold and snowy place. During the day the elements unleashed a terrific blizzard, blocking of the many small roads you use on your ride home from work with a snowdrift. Luckily, another tired worker on her way home meets the impasses at just about the same time but from the other side. Being aware of your environment both of you remembered to keep a shovel in your trunk for just such an emergency. You could go out and clear away the snowdrift, which would benefit both of you, or you could decide to stay inside, crank up the heat, and turn on the radio and wait for the other person to do it for you. If neither of you shovels, though, no one gets home and both of you are attacked by bears. (Did I mention there were bears?)

A = TURN UP THE HEAT
AND WAIT
B = SHOVEL

The last common interpretation of the game considers two types of fighters: hawks and doves. Biologist John Maynard Smith used this description to help describe the frequency of conventional contests in the animal world. Suppose there are two types of behaviors. The first one, HAWK, escalates conflict until injury or

A = ESCALATE THE FIGHT
B = POSTURE, BUT RUN
AWAY IF THREATENED

until its opponent flees. The second, DOVE, makes threatening displays but retreats once a fight begins. It's worth noting that these names are complete misnomers. First off, the game is supposed to take place between individuals in the same species. Hawks aren't doves, so we're already off to a misleading start. Also, in real life, doves do escalate the fight.

We'll return to the Hawk-Dove game again in an even more evolutionary context. This game has been used as a model of conflict, but also of group behavior and evolutionary branching and specialization.

2.4. Stag Hunt. ($\mathbf{a} = 4, \mathbf{b} = 0, \mathbf{c} = 6$). Once again, turn to your neighbor and play the game given by the payoff matrix

$$\begin{array}{cc} & \begin{array}{cc} \text{A} & \text{B} \end{array} \\ \begin{array}{c} \text{A} \\ \text{B} \end{array} & \begin{bmatrix} (2, 2) & (4, 0) \\ (0, 4) & (5, 5) \end{bmatrix} \end{array}$$

You may not be surprised to learn that this game also admits two pure and one mixed Nash equilibria. They are (A, A), (B, B) and $(\frac{1}{3}, \frac{2}{3})$. For those of you who attended the tutorial, you may remember that one-third of you voted for strategy A and about two-thirds of you voted for strategy B. How's that for external validation?

This game comes to us by way of Rousseau, who was interested in the tension between the goals of the individual and the goals of the group. In this example the group is as small as can be. Kathleen and Abby, bored by the lizards in the desert, venture on a big-game hunting trip in rugged Idaho. In the forest they happen upon a single stag and one rabbit. While neither has a sure enough shot to take down the stag on her own, they could easily do so together. But the two of them have been in the forest all day and are hungry. It occurs to both of them independently that it'd be easier to catch an eat a rabbit of their own instead. Going for the stag results in more food, but only if the other one coordinates her efforts. But if they both go for the rabbit, each enjoys a much smaller share.

This sort of game is also called a *coordination* game. The pure Nash equilibria in this game come in two different flavors: *risk dominant* and *payoff dominant*. Strategy A always ensures that the player adopting this strategy gets some food, and in some sense minimizes risk. A B-strategist could go hungry, but he could also end up with the stag, and thus maximizes his payoff. Again, classical game theory does not offer clear suggestions for choosing between these two equilibria. In evolutionary game theory, one could investigate the basins of attraction surrounding points.

Well, it's almost external validation. We had to look at a population of pure strategists to express the mixed strategy. It's impossible to know if someone really has adopted a mixed strategy in a one-shot game. After all, what does one-third of one game even mean? We'll return to this point later.

A = HUNT THE RABBIT
B = HUNT THE STAG

Coordination game

Risk dominance

Payoff dominance