An Invitation to Game Theory

R. Ramanujam
The Institute of Mathematical Sciences,
Chennai, India

jam@imsc.res.in

Conway's Game of Life

- A square grid of *cells*. Every cell has eight neighbours, and is alive or dead in any configuration.
- A dead cell with exactly three live neighbours comes alive.
- A live cell with either two or three live neighbours remains alive.
- Every other cell dies or stays dead.
- It is hard to predict whether an initial configuration will lead to all cells dying eventually, or grow forever or form a stable population. This is an undecidable problem.

Games in extensive form

- A game is given by a tuple $G = (I, T, \lambda, u)$ where I is a finite set of players, T is a finite rooted tree, $\lambda : N_T \rightarrow I$ is a labelling on N_T, the non-leaf nodes of the tree, by players and $u : (I \times L_T) \rightarrow \mathcal{R}$ is the utility (or payoff) function that specifies, for each player, the outcome at each leaf node of the tree.
- λ specifies which player's turn it is to move (or play) at any given node.
- A zero-sum game is one in which for every leaf node n, $\Sigma_{i \in I} u(i, n) = 0$. A win-lose game is one in which $u(i, n) \in \{0, 1\}$ for all $i \in I$ and leaf nodes n.

Strategies and winning strategies

In a two-person zero-sum game, exactly one player wins any play.

- A *strategy* for player i in game G is a subtree σ of G, which contains the root of G; for every $n \in \sigma$ such that $\lambda(n) = i$, there exists a unique G-successor of n in σ; for every $n \in \sigma$ such that $\lambda(n) \neq i$, every G-successor of n is in σ.
- Note that a *strategy profile*, or a tuple of strategies $\sigma = (\sigma_1, \ldots, \sigma_n)$, where σ_i is a strategy of player i in game G, generates a unique play in G (a maximal path from root to leaf).
- A strategy τ for player i in game G is said to be a winning strategy, if for every strategy profile σ such that $\sigma_i = \tau$, the resulting play is a win for player i.

Determinacy in two-person zero-sum games

In a two-person zero-sum game, exactly one player wins any play.

- A two-person zero-sum game is said to be determined if one of the two players has a winning strategy.
- **Theorem (Zermelo 1913)**: Every finite game two-person zero-sum game (of perfect information) is determined.
- The proof is by backward induction. Leaf nodes are labelled by winners. Given a labelling of children at a non-leaf node n with $\lambda(n) = i$, the node n gets label i there exists a child labelled i; otherwise all children are labelled by the other player j and n gets label j.
- Solving a game amounts to showing that the game is determined, finding who has a winning strategy, and presenting a winning strategy.

The Nim game

An example of a combinatorial game.

- Matchsticks are placed in heaps. Two players take turns, each player removing some non-zero number of sticks from one heap. The game goes on until all heaps are exhausted, and the last player to move wins the game.
- A characterization: Express the number of matches in each heap in binary. Form column sums. Call a configuration good if *all* column sums are even. We can show that any player moving from a good configuration makes it bad, and that any player moving from a bad configuration has a move so that the resulting one is good.
- Then we show that player II (I) has a winning strategy if the initial configuration is good (bad), solving the Nim game.

The normal form abstraction

In general, players have preferences on outcomes, not simply winning or losing. Decisions of players involve expected utility.

- Economists try to predict outcomes of games played by rational agents, without reference to how the game is actually played to achieve such outcomes.
- Consider a two-player game (of perfect information) in extensive form. This is a finite tree, and hence the set of possible strategies for a player (subtrees) is finite as well.
- Suppose player 1 has m strategies and 2 has n strategies. Then the game can be represented by an $m \times n$ matrix whose (i, j) entries are given by pairs (x, y): consider the unique play associated with the profile where 1 plays according to strategy i and 2 according to j. Then x is the payoff for player 1 and y for player 2 after that play.
- This can be done for any game of perfect information, such a presentation is said to be in normal form or strategic form.

Finite strategic form games

- $N = \{1, 2, \ldots, n\}$, the set of players.
- for each $i \in N$, a finite set $S_i = \{1, \ldots, m_i\}$ of pure strategies. Let $S = S_1 \times S_2 \times \cdots \times S_n$ be the set of possible combinations of pure strategies.
- for each $i \in N$, a payoff (utility) function $u_i : S \rightarrow \mathcal{R}$, describes the payoff $u_i(s_1, \ldots, s_n)$ to player i under each combination of strategies.

Iterated elimination of dominated strategies

A dominant strategy is one that is best, no matter what the opponent is doing.

	Left	Centre	Right
Up	(4,3)	(2,7)	(0,4)
Down	(5,5)	(5,-1)	(-1,-2)

- Centre dominates Right.
- On removal, Down dominates Up.
- The solution is (Down, Left).

The Prisoners' Dilemma

Example of Dominant Strategy Equilibrium

Jack John	Confess	Deny
Confess	(-10,-10)	(0,-20)
Deny	(-20,0)	(-1,-1)

- Both must confess and get 10 years in jail each.
- Source of much angst in game theory.

Disarmament Dilemma

Replace prisoners by countries and you see a familiar situation.

Country X Country Y	Arm	Disarm
Arm	(1,1)	(3,0)
Disarm	(0,3)	(2,2)

- Rationality leads to armament as the inevitable choice.
- Note that the ordering is relevant, not the numbers.

Hawks and Doves

V, the value of the resource, and C the cost of escalation.

X Y	Hawk	Dove
Hawk	$((V-C)/2, (V-C)/2)$	$(V,0)$
Dove	$(0,V)$	$(V/2, V/2)$

- If $V > C$, Hawk is the dominant strategy.
- If $V < C$, there is no dominant strategy. So dominant strategies do not always exist.

Nash Equilibrium

In game theory, every player not only optimizes, but also steps into the shoes of the opponent and solves her optimization problem as well.

Anu Babu	Cricket	Concert
Cricket	(2,1)	(0,0)
Concert	(0,0)	(1,2)

- There is no dominant strategy, choice depends on what the other player might do.
- Two Nash equilibria.

No Nash equilibrium

NE does not always exist.

Moriarty Holmes	Canterbury	Paris
Canterbury	(1,-1)	(-1,1)
Paris	(-1,1)	(1,-1)

- There is no NE: whichever player loses, should anticipate losing and hence choose a different strategy.
- Each player can toss a coin and decide to play one of the two moves with probability $1/2$. The other would do the same and this is in fact a Nash equilibrium in mixed strategies.

Mixed strategies

Randomized strategies.
A **mixed (randomized)** strategy x_i for player i, with $S_i = \{1, \ldots, m_i\}$ is a probability distribution over S_i: a vector $x_i = (x_1(1), \ldots, x_i(m_i))$ such that for $1 \leq j \leq m_i$, $x_i(j) \geq 0$, and:

$$x_i(1) + x_i(2) + \ldots + x_i(m_i) = 1$$

Player i uses randomness to decide which strategy to play, based on the probabilities in x_i.
Let X_i be the set of mixed strategies for player i. For an n-player game, let $X = X_1 \times \ldots \times X_n$.
X denotes the set of all possible combinations, or profiles of mixed strategies.

Expected payoffs

Let $x = (x_1, \ldots, x_n) \in X$ be a profile of mixed strategies.
For $s = (s_1, \ldots, s_n) \in S$, a combination of pure strategies, let:

$$x(s) = \Pi_{j=1}^{n} x_j(s_j)$$

be the probability of the combination s under mixed profile x.
We are assuming that the players make their random choices independently.
The expected payoff for player i under a profile $x = (x_1, \ldots, x_n) \in X$ is:

$$U_i(x) = \Sigma_{s \in S} x(s) * u_i(s)$$

Expected payoffs

The expected payoff for player i under a profile $x = (x_1, \ldots, x_n) \in X$ is:

$$U_i(x) = \Sigma_{s \in S} x(s) * u_i(s)$$

It is the weighted average of what player i can win under each pure combination s, weighted by the probability of that combination.
Key assumption: Every player's goal is to maximize her own payoff.
`Statutory Warning`: this assumption is often dubious.

Some notation

A mixed strategy $x_i \in X$ is said to be pure if for some $j \in S_i$, $x_i(j) = 1$ and for all $j' \neq j$, $x_i(j') = 0$. Let π_{ij} denote such a pure strategy.

The "mixed" strategy π_{ij} does not randomize at all: it picks (with probability 1) exactly one strategy, j, from the set of pure strategies for player i.

Given a profile $x = (x_1, \ldots, x_n) \in X$, let:

$$x_{-i} = (x_1, \ldots, x_{i-1}, empty, x_{i+1}, \ldots, x_n)$$

x_{-i} denotes everybody's strategy except that of player i.

Some notation

x_{-i} denotes everybody's strategy except that of player i.
By abuse of notation, for $y_i \in X_i$, let $(x_{-i}; y_i)$ denote the new profile:

$$(x_1, \ldots, x_{i-1}, y_i, x_{i+1}, \ldots, x_n)$$

$(x_{-i}; y_i)$ is the new profile where everybody's strategy remains the same as in x, except for player i, who switches from mixed strategy x_i to mixed strategy y_i.

Best responses

A mixed strategy $z_i \in X$ is a best response for player i to x_{-i} if, for all $y_i \in X_i$,

$$U_i((x_{-i}; z_i)) \geq U_i((x_{-i}; y_i))$$

Clearly, if any player were given the opportunity to "cheat" and look at what other players have done, she would switch her strategy to a best response. By the rules of the game, all players pick their strategies simultaneously.

Suppose, somehow, that players "arrive" at a profile where everybody's strategy is a best response to everybody else's. Then no one has any incentive to change the situation. Then we are in a stable, or equilibrium situation, which we call Nash equilibrium.

Nash equilibrium

For a strategic game Γ, a strategy profile $x = (x_1, \dots, x_n) \in X$ is a mixed Nash equilibrium if, for every player i, x_i is a best response to x_{-i}.
That is, for $1 \leq i \leq n$, and for every $y_i \in X_i$,

$$U_i(x_{-i}; x_i) \geq U_i(x_{-i}; y_i)$$

In other words, no player can improve her own payoff by unilaterally deviating from the mixed strategy profile $x = (x_1, \dots, x_n)$.
x is called a pure Nash equilibrium if every x_i is a pure strategy π_{ij} for some $j \in S_i$.
(There are many interpretations of a Nash equilibrium).

Nash's Theorem

This can, with some justification, be called The Fundamental Theorem of Game Theory.

Theorem (Nash, 1951): Every finite *n*-person strategic game has a mixed Nash equilibrium.

He used a fundamental result from topology to prove the theorem, namely the Brouwer fixed point theorem.

The crumpled sheet experiment

- Take two identical rectangular sheets of paper.
- Make sure neither sheet has any hole in it, and that the sides are straight.
- "Name" each point on both sheets by its (x, y) coordinates.
- Crumple one of the two sheets any way you like, but *make sure you don't tear it in the process*.
- Place the crumpled sheet completely on top of the other flat sheet.

The crumpled sheet theorem

There must be a point named (a, b) on the crumpled sheet that is directly above the same point (a, b) on the flat sheet.

This fact, in its more formal and general form, is the key to why every finite game has a mixed Nash equilibrium.

The Brouwer fixed point theorem

Theorem (Brouwer, 1909): Every continuous function $f : D \rightarrow D$ mapping a compact, convex and non-empty subset $D \subset \mathcal{R}^m$ to itself has a "fixed point", i.e. there exists $x^* \in D$ such that $f(x^*) = x^*$.

The definitions

- $D \subset \mathcal{R}^m$ is convex, if for all $x, y \in D$ and all $\lambda \in [0, 1]$, $\lambda x + (1 - \lambda)y \in D$.
- $D \subset \mathcal{R}^m$ is compact iff it is closed and bounded.
- $D \subset \mathcal{R}^m$ is bounded if there exists a non-negative integer K such that $D \subseteq [-K, K]^m$. (i.e. D "fits inside" a finite m-dimensional box.)
- $D \subset \mathcal{R}^m$ is closed iff for all sequences $x_0, x_1, \ldots$, where for all $i \geq 0$, $x_i \in D$, if $\exists x \in \mathcal{R}^m$ such that $x = lim_i x_i$ then $x \in D$. (i.e. if a sequence of points is in D and the sequence has a limit, then the limit is also in D).

The definitions

A function $f : D \subset \mathcal{R}^m \rightarrow \mathcal{R}^m$ is continuous at a point $x \in D$ if for all $\epsilon > 0$, there exists $\delta > 0$ such that for all $y \in D$: if $dist(x, y) < \delta$ then $dist(f(x), f(y)) < \epsilon$. f is called continuous if it is continuous at every point $x \in D$.

The idea

Consider the interval $[0, 1]$. It is compact and convex.
More generally, $[0, 1]^n$ is compact and convex.
The set of profiles $X = X_1 \times \ldots X_n$ is a compact and convex subset of $\mathcal{R}^m$, where $m = \Sigma_{i=1}^n m_i$, and $m_i = |S_i|$.
We will define a continuous function $f : X \to X$ and show that if $f(x^*) = x^*$ then $x^* = (x_1^*, \ldots, x_n^*)$ must be a Nash equilibrium.
By Brouwer's theorem, we are done.
In fact, it turns out that x^* is a Nash equilibrium iff x^* is a fixed point of f.

A claim

Claim: A profile $x^* = (x_1^*, \ldots, x_n^*)$ is a Nash equilibrium iff for every player i, and every pure strategy $\pi_{ij}, j \in S_i$,

$$U_i(x^*) \geq U_i(x_{-i}^*; \pi_{ij})$$

Proof of claim: If x^* is a NE then the inequality is obvious. For the other direction, first note that, for any $x_i \in X_i$,

$$U_i(x_{-i}^*; x_i) = \Sigma_{j=1}^{m_i} x_i(j) * U_i(x_{-i}^*; \pi_{ij})$$

That is, the payoffs of player i is the "weighted average" of the payoffs of each of her pure strategies, weighted by the probability of that strategy.

Proof (continued)

We have:

$$U_i(x^*_{-i}; x_i) = \Sigma_{j=1}^{m_i} x_i(j) * U_i(x^*_{-i}; \pi_{ij})$$

By assumption, for all $j \in S_i$,

$$U_i(x^*) \geq U_i(x^*_{-i}; \pi_{ij})$$

So clearly, $U_i(x^*) \geq U_i(x^*_{-i}; x_i)$, for any $x_i \in X_i$, since a "weighted average" no bigger than $U_i(x^*)$ cannot exceed $U_i(x^*)$.
Hence each x_i^* is a best response strategy to x^*_{-i}, and hence x^* is a NE.

What we want

Thus we wish to find x^* such that for all i, and for all $j \in S_i$,

$$U_i(x^*) \geq U_i(x^*_{-i}; \pi_{ij})$$

That is, for all $i, j \in S_i$,

$$U_i(x^*_{-i}; \pi_{ij}) - U_i(x^*) \leq 0$$

For a mixed profile $x = (x_1, \ldots, x_n) \in X$ define:

$$\phi_{ij}(x) = max\{0, U_i(x_{-i}; \pi_{ij}) - U_i(x)\}$$

Intuitively, $\phi_{ij}(x)$ measures "how much better off" player i would be if she picked π_{ij} instead of x_i, when everone else's strategy is fixed.

A function

Define $f : X \to X$ as follows. For $x = (x_1, \ldots, x_n) \in X$, let:

$$f(x) = (x'_1, x'_2, \ldots, x'_n)$$

where, for all i, $1 \leq j \leq m_i$,

$$x'_i(j) = (x_i(j) + \phi_{ij}(x))/(1 + \Sigma_{k=1}^{m_i} \phi_{ik}(x))$$

If $x \in X$, then $f(x) \in X$.
f is continuous.
Thus, by Brouwer's theorem, f has a fixed point, x^*.
We need to show that x^* is a NE.

x^* is NE

For each i, $1 \leq j \leq m_i$,

$$x_i^*(j) = (x_i^*(j) + \phi_{ij}(x^*))/(1 + \Sigma_{k=1}^{m_i} \phi_{ik}(x^*))$$

Thus,

$$x_i^*(j)(1 + \Sigma_{k=1}^{m_i} \phi_{ik}(x^*)) = (x_i^*(j) + \phi_{ij}(x^*))$$

Hence,

$$x_i^*(j) \Sigma_{k=1}^{m_i} \phi_{ik}(x^*)) = \phi_{ij}(x^*)$$

We show that this in fact implies: RHS $= 0$, for all j.

Another Claim

Claim: For any mixed profile x, for player i, there exists j such that $x_i(j) > 0$ and $\phi_{ij}(x) = 0$.
Assume the claim. Then for such a j, $x_i^*(j) > 0$ and:

$$x_i^*(j)\Sigma_{k=1}^{m_i}\phi_{ik}(x^*)) = 0 = \phi_{ij}(x^*)$$

But all $\phi_{ik}(x^*) \geq 0$, hence for all k, $\phi_{ik}(x^*) = 0$.
Thus, for all i, $1 \leq j \leq m_i$,

$$U_i(x^*) \geq U_i(x^*_{-i}; \pi_{ij})$$

as required.

Proof of Claim

Claim: For any mixed profile x, for player i, there exists j such that $x_i(j) > 0$ and $\phi_{ij}(x) = 0$.

$$\phi_{ij}(x) = max\{0, U_i(x_{-i}; \pi_{ij}) - U_i(x)\}$$

Since $U_i(x)$ is the "weighted average" of the $U_i(x_{-i}; \pi_{ij})$'s, based on the weights in x_i, there exists j such that $x_i(j) > 0$ such that $U_i(x_{-i}; \pi_{ij})$ does not exceed the "weighted average".
That is,

$$U_i(x_{-i}; \pi_{ij}) - U_i(x) \leq 0$$

Hence for this j, $\phi_{ij}(x) = 0$, as required.

A Corollary

Since $U_i(x^*)$ is the "weighted average" of the $U_i(x^*_{-i}; \pi_{ij})$'s, we see that:

$U_i(x^*) = U_i(x^*_{-i}; \pi_{ij})$, whenever $x^*_i(j) > 0$.

That is, in any NE x^*, if $x^*_i(j) > 0$, then each such π_{ij} is itself a best response to x^*_{-i}.

(Useful in computing NE.)

Two-person zero-sum games

Note that $u_1(s) = -u_2(s)$, for every strategy profile s.
$u(s_1, s_2)$ can be conveniently viewed as an $m_1 \times m_2$ matrix A_1 where $A_2 = -A_1$.
Thus we are given only one matrix A, player 1 maximizing $u(s)$.
Let $A(i, j)$ denote the entry in the ith row and jth column.

Matrix view of 2p-zs games

Suppose players 1 and 2 choose mixed strategies x_1 and x_2 respectively.
Consider the product $x_1^T\ A\ x_2$.
We see that $x_1^T\ A\ x_2 = \Sigma_i \Sigma_j (x_1(i) * x_2(j)) * A(i,j)$.
But note that $(x_1(i) * x_2(j))$ is precisely the probability of the (pure) combination (i,j).
Thus, for the mized profile $x = (x_1, x_2)$,
$x_1^T\ A\ x_2 = U_1(x) = -U_2(x)$,
where $U_1(x)$ is the expected payoff (which player 1 is trying to maximize, 2 is trying to minimize).

Minimaximizing strategies

Suppose player 1 decides on x_1^* by trying to maximize the *worst that can happen*. That would be player 2 choosing x_2 that minimizes $(x_1^*)^T \ A \ x_2$.
Define x_1^* to be a minimaximizer for player 1 if

$$min_{x_2}(x_1^*)^T \ A \ x_2 = max_{x_1} min_{x_2} x_1^T \ A \ x_2$$

Similarly x_2^* is a maximinimizer for player 2 if

$$max_{x_1} x_1^T \ A \ x_2^* = min_{x_2} max_{x_1} x_1^T \ A \ x_2$$

Note that $min_{x_2}(x_1^*)^T \ A \ x_2 \leq (x_1^*)^T \ A \ x_2^* \leq max_{x_1} x_1^T \ A \ x_2^*$.
In 1928, von Neumann showed that equality holds.

The Minimax theorem

For a two-person zero-sum game given by A, there exists a unique value $v^* \in \mathcal{R}$ such that, for some $x^* = (x_1^*, x_2^*)$:

1. For all j, $((x_1^*)^T\ A)_j \geq v^*$.
2. For all j, $(A\ x_2^*)_j \leq v^*$.
3. Thus, $v^* = (x_1^*)^T\ A\ x_2^*$, and

$$max_{x_1} min_{x_2} x_1^T\ A\ x_2 = v^* = min_{x_2} max_{x_1} x_1^T\ A\ x_2$$

4. The conditions above hold precisely when (x_1^*, x_2^*) is a Nash equilibrium.

x_1^* is a minimaximizer for player 1 and x_2^* is a maximinimizer for player 2.

Remarks

x_1^* guarantees player 1 at least expected profit v^*, and x_2^* guarantees player 2 at most expected loss v^*.
We call any such (x_1^*, x_2^*) a mini-max profile.
We call the unique value v^* the mini-max value of the game.
Now it is obvious that max profit guaranteed for player 1 is $\leq$ the minimum loss guaranteed for player 2. The theorem asserts the non-trivial converse.

Proof

The mini-max theorem follows easily from Nash's theorem.
Let (x_1^*, x_2^*) be an NE. Let $v^* = (x_1^*)^T\ A\ x_2^* = U_1(x^*) = -U_2(x^*)$.
Since x_i^* is the best response to the other, we have:
$U_i(x_{-i}^*; \pi_{i,j}) \leq U_i(x^*)$.
But $U_1(x_{-1}^*; \pi_{1,j}) = (Ax_2^*)_j$. Thus, $(Ax_2^*)_j \leq v^* = U_1(x^*)$.
$U_2(x_{-2}^*; \pi_{2,j}) = -((x_1^*)^T A)_j$. Thus, $((x_1^*)^T A)_j \geq v^* = -U_2^(x^*)$.

Proof (contd)

$max_{x_1} x_1^T \ A \ x_2^* \leq v^*$, because $x_1^T \ A \ x_2^*$ is a weighted average of $(Ax_2^*)_j$'s.
Similarly. $v^* \leq min_{x_2} (x_1^*)^T \ A \ x_2$.
Thus $max_{x_1} x_1^T \ A \ x_2^* \leq min_{x_2} (x_1^*)^T \ A \ x_2$.
Thus $min_{x_2} max_{x_1} x_1^T \ A \ x_2 \leq max_{x_1} min_{x_2} x_1^T \ A \ x_2$.
We did not assume anything about the NE chosen. So for every NE x^*, if $v' = (x_1^*)^T \ A \ x_2^*$, we get $v' = v^*$.
It is easy to see that any x^* satisfying conditions 1-3 is an NE.

Remarks

Over 2p-zs games, NE and minimax profiles are the same.
Moreover, when $x^* = (x_1^*, x_2^*)$ is a minimax profile and $x_1^*(j) > 0$, we have: $((x_1^*)^T A)_j = v^* = (x_1^*)^T\ A\ x_2^*$.
Similarly if $x_2^*(j) > 0$, $(A\ x_2^*)_j = v^* = (x_1^*)^T\ A\ x_2^*$.
Note: we have, as yet, no clue how to compute the minimax value and a minimax profile.
We are trying to maximize v subject to $(x_1^T\ A)_j \geq v$, for all j.
Optimizing a linear objective subject to linear constraints.

Critique of NE-1

When there are multiple equilibria, it is not clear which one the players would or should try for.

Rose Colin	A	B
A	(1,1)	(2,5)
B	(5,2)	(-1,-1)

- (A,B) and (B,A) are non-equivalent and non-interchangeable pure Nash equilibria.
- But (A,B) is better for Colin and (B,A) is better for Rose. If both try for their favourite equilibrium, they will end up with (B,B) which is not an equilibrium and indeed the worst outcome possible.

Critique of NE-2

Nash equilibria need not be Pareto optimal.

Rose Colin	A	B
A	(3,3)	(-1,5)
B	(5,-1)	(0,0)

- (B,B) is the unique (pure) Nash equilibrium.
- But (A,A) is a better outcome for both players.

Solvable in the strict sense

A better solution concept for non-zero sum games

- There is at least one equilibrium outcome that is Pareto optimal, and
- All Pareto optimal outcomes are equivalent and interchangeable.
- Below, (B,B) and (A,C) are equilibria, but the latter is the only Pareto optimal outcome.

Rose Colin	A	B	C
A	(0,-1)	(0,2)	(2,3)
B	(0,0)	(2,1)	(1,-1)
C	(2,2)	(1,4)	(1,-1)

Natural selection rather than rationality

Payoff is *fitness points*: increased probability of passing along genes to the next generation.

P1 P2	Hawk	Dove
Hawk	(-25,-25)	(50,0)
Dove	(0,50)	(15,15)

- A Hawk fights for a resource; a dove merely postures.
- Winner of resource gets 50 points. A losing Hawk gets -100, and wasting time (for doves) gets -10 points.
- Each player is genetically determined to always play hawk or dove.

Population dynamics

What kind of populations are stable ?

- Suppose that the population starts off with almost entirely doves.
- Doves meet mostly doves, so get 15 fitness points.
- An occasional hawk gets 50, and being genetically advantaged, the hawk population will begin to rise.
- A hawk minority (by mutation) would eventually invade the population.
- Thus a population of doves is not evolutionarily stable.

Mixed strategies

How about a mixed population, of $\frac{1}{4}$ hawks and $\frac{3}{4}$ doves?

Best solved by considering a *focal player* playing against an opponent playing a mixed strategy.

Focal Other	Hawk	Dove
Hawk	(-25,-25)	(50,0)
Dove	(0,50)	(15,15)

- Expected payoff for focal player playing hawk is $31\frac{1}{4}$ and dove is $11\frac{1}{4}$.
- It pays to be a hawk, and hence hawks will increase.

Evolutionary equilibria

When can we be sure that no mutant strategy will invade ?

- If there are few hawks then hawks will increase and if there are few doves, then doves will increase.
- There must be some proportion of hawks and doves where the two tendencies balance out.
- Consider a mixed population, of x hawks and $1 - x$ doves. Then this is the mixed strategy that the focal player is up against.
- Expected payoff for focal player playing hawk is $50 - 75x$ and dove is $15 - 15x$.
- Solving, $x = \frac{7}{12}$. Thus a population of $\frac{7}{12}$ hawks and $\frac{5}{12}$ doves would be evolutionarily stable.

Evolutionary game theory

Evolutionarily stable strategies.

- We say that a strategy S is evolutionarily stable if the following condition holds: Let T be any strategy. If almost everyone in the population plays S and a few play T. Then the expected payoff for playing S should be at least as much as the expected payoff for playing T.
- If a population has adopted S, no mutant strategy T can invade and prosper against S.
- We no longer need the assumption that each individual is either a pure hawk or pure dove. The same kind of stability obtains if all players played a mixed strategy of $\frac{7}{12}$ hawk and $\frac{5}{12}$ dove.
- Different individuals may play different strategies but averaging to $\frac{7}{12}$ hawks across the population.

ESS Examples

A pure ES strategy.

P1 P2	A	B
A	(1,1)	(2,3)
B	(3,2)	(4,4)

- Strategy B is an ESS and the only one.
- B strictly dominates A, and hence advantaged in any population.

ESS Examples

More than one ES strategy.

P1 P2	A	B
A	(3,3)	(1,2)
B	(2,1)	(4,4)

- Strategy A and B are both ESS.
- In any population of almost all B's, B would be best (by 4:1).
- In a population of almost all A's, A would be best (by 3:2).
- Whichever gets established first would persist.

ESS Examples

No pure ES strategy.

P1 P2	A	B
A	(1,1)	(4,2)
B	(2,4)	(3,3)

- The unique ESS is a mixed strategy.
- This game is ordinally equivalent to our original hawk dove game.

Expanding our horizons

EGT wants us to resist invasion by *any* other strategy T. How can we be sure that we have identified *all* feasible strategies ?

- Since it is advantageous to play dove against a hawk and to play hawk against a dove, why not play *conditional* strategies ?
- Bully: In any contest, show initial fight. Continue to fight if opponent does not fight back. If opponent fights, run away.
- When bullies meet, both run away but one runs faster and the other is left holding the prize.

P1 P2	Hawk	Dove	Bully
Hawk	(-25,-25)	(50,0)	(50,0)
Dove	(0,50)	(15,15)	(0,50)
Bully	(0,50)	(50,0)	(25,25)

Further expansion of horizons

What is the best way to deal with bullies?

P1 P2	Hawk	Dove	Bully
Hawk	(-25,-25)	(50,0)	(50,0)
Dove	(0,50)	(15,15)	(0,50)
Bully	(0,50)	(50,0)	(25,25)

- No pure strategy is an ESS.
- Bully dominates dove, so doves would eventually die out.
- The only ESS is half-hawk and half-bully.
- We can expect a lot of conflict and cowardice.

Dealing with bullies

The solution you used as a child.
Retaliator: In any contest, initially behave as a dove. However, if you are persistently attacked, fight back with all your strength.

P1 P2	Hawk	Dove	Bully	Retaliator
Hawk	(-25,-25)	(50,0)	(50,0)	(-25, -25)
Dove	(0,50)	(15,15)	(0,50)	(15, 15)
Bully	(0,50)	(50,0)	(25,25)	(0, 50)
Retaliator	(-25, -25)	(15, 15)	(50,0)	(15, 15)

- The pure strategy retaliator is an ESS. So is any mixture of retaliators and doves, with $< 30\%$ doves.
- A little paradoxical, as retaliators should have the worst of both worlds.
- Biologically interesting: in a population of mostly retaliators, there is not much conflict but much posturing. This is reminiscent of behaviour recorded by Konrad Lorentz.

Signalling

ESS can provide Pareto-inferior outcomes. This can be overcome by co-operation: signals telling a player to be a hawk etc.

Bourgeois: Be a hawk on home ground, and a dove elsewhere.

P1 P2	Hawk	Dove	Bully	Retaliator	Bourgeoi
Hawk	(-25,-25)	(50,0)	(50,0)	(-25, -25)	(12.5, -1
Dove	(0,50)	(15,15)	(0,50)	(15, 15)	(7.5, 32.
Bully	(0,50)	(50,0)	(25,25)	(0, 50)	(25, 25)
Retaliator	(-25, -25)	(15, 15)	(50,0)	(15, 15)	(-5, -5)
Bourgeois	(-12.5, 12.5)	(32.5, 7.5)	(25,25)	(-5, -5)	(25, 25)

Both pure strateges retaliator and bourgeois are ESS.

An invitation to Game Theory

- EGT can be used to study not only emergence of aggression and responses to it, but also sexual behaviour, altruism and cooperation.
- Studying repeated games leads to interesting notions like threats, promises and commitments.
- Placing limits on rationality assumptions is an important topic of research.
- We have not touched on games of imprefect information, where the notions are in general more complicated.
- We have also not spoken of co-operative game theory which studies behaviour of coalitions.
- The study of infinite games offers a rich mathematical theory with important implications for descriptive set theory.

Bibliography

- Robert Aumann and Sergiu Hart, Handbook of game theory, 1992.
- Robert Axelrod, The Evolution of Cooperation, 1984.
- Berlekamp, Conway and Guy, Winning ways for your mathematical plays, 1982.
- Antonia Jones, Mathematical models of conflict, 1983.
- Yiannis Moschovakis, Descriptive set theory, 1980.
- Noam Nisan et al, Algorithmic game theory, 2005.
- John Maynard Smith, Evolution and the theory of games, 1982.
- Anatol Rapoport, Mathematical models in the social and behavioural sciences, 1983.
- Philip Straffin, Game theory and strategy, 1993.