

A brief introduction into evolutionary game theory

Gerhard Jäger

gerhard.jaeger@uni-tuebingen.de

October 12, 2009

Uni Tübingen

“The formation of different languages and of distinct species, and the proofs that both have been developed through a gradual process, are curiously parallel. . . . Max Müller has well remarked: ‘A struggle for life is constantly going on amongst the words and grammatical forms in each language. The better, the shorter, the easier forms are constantly gaining the upper hand, and they owe their success to their inherent virtue.’ To these important causes of the survival of certain words, mere novelty and fashion may be added; for there is in the mind of man a strong love for slight changes in all things. The survival or preservation of certain favoured words in the struggle for existence is natural selection.” (Darwin)

Historical remarks

- GT developed by John von Neumann and Oskar Morgenstern (1944: “Theory of Games and Economic Behavior”)
- meta-theory for economy and political strategy (cold war)
- standard tool in economics (Nobel prize for economics 1994 for Nash, Harsanyi and Selten, and 2006 for Aumann and Schelling)
- since early 1970s application in biology to model Darwinian natural selection (1973, John Maynard Smith and George Price, “The logic of animal conflict”, 1982: John Maynard Smith, “Evolution and the Theory of Games”)
- connections to epistemic logic (Stalnaker, Spohn)
- application in pragmatics/philosophy of language
 - David Lewis (1969: “Conventions”)
 - growing body of work in recent years (Parikh, Merin, van Rooij, ...)

Definition

A **strategic game** consists of

- a set of **players**
 - for each player, a set of **actions**
 - for each player, **preferences** over the set of action profiles
-
- A *action profile* is an assignment of an action to each player.
 - Preferences are expressed as **utilities** (real numbers):

$$u(a) > u(b)$$

if and only if the decision maker prefers profile a over profile b .

Prisoner's dilemma

“Two suspects in a major crime are held in separate cells. There is enough evidence to convict each of them of a minor offense, but not enough evidence to convict either of them of the major crime unless one of them acts as an informer against the other (finks). If they both stay quiet, each will be convicted of the minor offense and spend one year in prison. If one and only one of them finks, she will be freed and used as a witness against the other, who will spend four years in prison. If they both fink, each will spend three years in prison.” (Osborne, p. 14)

Players: The two suspects.

Actions: Each player's set of actions is $\{Quiet, Fink\}$

Preferences: Each player wants to spend as little time in prison as possible.

- Preferences can be expressed as **utility matrix**:
 - each dimension corresponds to one player
 - each row/column(/layer/...) corresponds to one strategy
 - each cell corresponds to one profile
 - each cell contains n numbers, one utility for each player

Prisoner's dilemma

Utility matrix

		Suspect 2	
		<i>Quiet</i>	<i>Fink</i>
Suspect 1	<i>Quiet</i>	2,2	0,3
	<i>Fink</i>	3,0	1,1

Utility matrix of two-person games

- In two-person games, the first number is by convention the row player's utility, and the second number the column player's

General format for two-player utility matrix

	C_1	C_2
R_1	$u_R(R_1, C_1), u_C(R_1, C_1)$	$u_R(R_1, C_2), u_C(R_1, C_2)$
R_2	$u_R(R_2, C_1), u_C(R_2, C_1)$	$u_R(R_2, C_2), u_C(R_2, C_2)$

Two people want to go out together. There is a concert with music by Bach, and one with music by Stravinsky. One of them loves Bach and the other Stravinsky, but they both prefer going out together over going to their favorite concert alone.

Bach or Stravinsky

Two people want to go out together. There is a concert with music by Bach, and one with music by Stravinsky. One of them loves Bach and the other Stravinsky, but they both prefer going out together over going to their favorite concert alone.

Utility matrix

	<i>Bach</i>	<i>Stravinsky</i>
<i>Bach</i>	2,1	0,0
<i>Stravinsky</i>	0,0	1,2

Stag hunt

(from Rousseau's "Discourse on the origin and foundations of inequality among men") *A group of people want to hunt together. If they stay together and coordinate, they will be able to catch a stag. If only one of them defects, they will get nothing. Each of them has a good chance to hunt a hare if he goes hunting by himself. A stag is better than a hare, which is still better than nothing.*

Stag hunt

(from Rousseau's "Discourse on the origin and foundations of inequality among men") *A group of people want to hunt together. If they stay together and coordinate, they will be able to catch a stag. If only one of them defects, they will get nothing. Each of them has a good chance to hunt a hare if he goes hunting by himself. A stag is better than a hare, which is still better than nothing.*

Utility matrix

	<i>Stag</i>	<i>Hare</i>
<i>Stag</i>	2,2	0,1
<i>Hare</i>	1,0	1,1

Mixed strategies: motivation

- players may choose to randomize their action
- games may involve random pairing from a population
- I may have incomplete knowledge about the actions of the other players, but enough knowledge to quantify my ignorance, i.e., to assign probabilities

In these cases, a rational decision has to be based on the **expected utility**, taking probabilities into account.

Definition

A **mixed strategy** of a player in a strategic game is a probability distribution over the player's action.

If the other players play mixed strategies, my utility for each of my possible actions becomes a random variable. I don't know its value in advance, but I can calculate its expected value. Also, if I play a mixed strategy myself, my utility is a random variable.

Definition (Expected utility)

Let α be a mixed strategy profile, and α_j be the mixed strategy of player j in profile α .

The **expected utility** for player i in the mixed profile α is defined as

$$u_i(\alpha) = \sum_a (\prod_j \alpha_j(a_j)) u_i(a)$$

where a ranges over pure strategy profiles.

- some more notation:

Profiles

Let α be a (possibly mixed) action profile and i a player.

- α_i is the strategy of player i in the profile α .
- α_{-i} is the profile of actions that all players **except** i play in α .

In a two-person game, α_{-i} is simply the action of the other player in α .

Definition (Strict domination)

In a strategic game, player i 's action α_i'' **strictly dominates** her action α_i' if

$$u_i(\alpha_i'', \alpha_{-i}) > u_i(\alpha_i', \alpha_{-i})$$

for every distribution α_{-i} of the other players' actions.

An example

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

- no rational player would ever play a strictly dominated strategy
- therefore they can be left out of consideration
- if a mixed strategy is strictly dominated, all pure strategy in its support are strictly dominated as well — so we only eliminate pure strategies
- note that a pure strategy may be dominated by a mixed strategy (plays no role in this example)
- this procedure can be iterated

An example

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

- no rational player would ever play a strictly dominated strategy
- therefore they can be left out of consideration
- if a mixed strategy is strictly dominated, all pure strategy in its support are strictly dominated as well — so we only eliminate pure strategies
- note that a pure strategy may be dominated by a mixed strategy (plays no role in this example)
- this procedure can be iterated

An example

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

- no rational player would ever play a strictly dominated strategy
- therefore they can be left out of consideration
- if a mixed strategy is strictly dominated, all pure strategy in its support are strictly dominated as well — so we only eliminate pure strategies
- note that a pure strategy may be dominated by a mixed strategy (plays no role in this example)
- this procedure can be iterated

An example

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

- no rational player would ever play a strictly dominated strategy
- therefore they can be left out of consideration
- if a mixed strategy is strictly dominated, all pure strategy in its support are strictly dominated as well — so we only eliminate pure strategies
- note that a pure strategy may be dominated by a mixed strategy (plays no role in this example)
- this procedure can be iterated

An example

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

- no rational player would ever play a strictly dominated strategy
- therefore they can be left out of consideration
- if a mixed strategy is strictly dominated, all pure strategy in its support are strictly dominated as well — so we only eliminate pure strategies
- note that a pure strategy may be dominated by a mixed strategy (plays no role in this example)
- this procedure can be iterated

Order of iterated elimination does not matter

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

Order of iterated elimination does not matter

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

Order of iterated elimination does not matter

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

Order of iterated elimination does not matter

8,3	6,4	15,0
9,1	5,2	6,3
3,2	4,3	5,4
2,9	3,10	4,8

Theorem

In a finite game, a unique set of action profiles survives iterated elimination of strictly dominated actions.

Rationality

A player is **rational** iff

- he holds consistent beliefs,
- he is logically omniscient,
- he knows the utility matrix (i.e. the preferences of the other players), and
- always chooses an action that maximizes the utility that he expects on the basis of his beliefs.

Rationalizability

An action profile a is **rationalizable** if there is a situation where

- each player is rational,
- it is common knowledge among the players that each player is rational
- each player i plays a_i .

Theorem

The action profiles that survive iterated elimination of strictly dominated actions are exactly those that are rationalizable.

How should a rational player play?

- rational people should play rationalizable actions
- Prisoner's dilemma: only one rationalizable profile (F, F)
- **but:** in Stag Hunt (and BoS etc.), all actions are rationalizable
- **Suppose** you know for sure what the other player does \Rightarrow simplifies the decision a lot

Definition (Best response)

Let α be a strategy profile. α_i is the **best response** of player i to the strategy profile α_{-i} of the other players iff

$$u_i(\alpha_i, \alpha_{-i}) \geq u_i(\alpha'_i, \alpha_{-i})$$

for any alternative strategies α'_i of player i .

If a rational player knows the actions of the other players, he will always play a best response.

Nash equilibria

- Suppose each player knows in advance what the others will do.
- If all players are rational, they will all play a best response to the actions of the others.
- Such a state is called **equilibrium**.
- First discovered by *John Nash*, therefore **Nash equilibrium**

Definition (Nash equilibrium)

The profile α is a **Nash equilibrium** if for each player i , α_i is a best response to α_{-i} .

Do the following games have Nash equilibria, and if yes, which ones?

- 1 Prisoner's dilemma
- 2 Bach or Stravinsky
- 3 Stag hunt
- 4 Hawks and Doves

Hawks and Doves

	<i>Hawk</i>	<i>Dove</i>
<i>Hawk</i>	1,1	7,2
<i>Dove</i>	2,7	3,3

Nash equilibria

Matching pennies

	<i>Head</i>	<i>Tail</i>
<i>Head</i>	1,-1	-1,1
<i>Tail</i>	-1,1	1,-1

Rock-Paper-Scissors

	<i>Rock</i>	<i>Paper</i>	<i>Scissor</i>
<i>Rock</i>	0,0	-1,1	1,-1
<i>Paper</i>	1,-1	0,0	-1,1
<i>Scissor</i>	-1,1	1,-1	0,0

Non-strict NEs

1,1	1,0	0,1
1,0	0,1	1,0

- one NE: (R_1, C_1)
- for R , it is not the unique best response to C_1

Theorem (Existence of mixed strategy Nash equilibrium in finite games)

Every strategic game in which each player has finitely many actions has a mixed strategy Nash equilibrium.

- Suppose you are the row player in BoS. The column player will play *Bach* with probability $\frac{1}{3}$ and *Stravinsky* with probability $\frac{2}{3}$. What is your expected utility for *Bach*? What for *Stravinsky*? What for the mixed strategy: playing *Bach* with probability p and *Stravinsky* with probability $1 - p$?
- Same problem for Stag hunt.
- What is your maximal expected utility that one can achieve in Matching Pennies, provided the other player knows your strategy and is rational?
- Same problem for Rock-Paper-Scissors.

- The following games have one mixed strategy equilibrium each:
 - Bach or Stravinsky
 - Stag hunt
 - Hawk and Dove
 - Matching Pennies
 - Rock-Paper-Scissors

Find them.

- if the “game” is a symmetric interaction between members of same population, players can swap places

Symmetric games

A two-person game is symmetric only if both players have the same set of strategies at their disposal, and the utility matrix is symmetric in the following sense:

$$u_R(R_n, C_m) = u_C(R_m, C_n)$$

for all strategies m and n .

- symmetric games (more precisely: games that can be conceived as symmetric):
 - Prisoner's dilemma
 - Stag hunt
 - Hawk and Dove
 - Rock-Paper-Scissors
- asymmetric games (more precisely: games that cannot be conceived as symmetric):
 - Bach or Stravinsky
 - Matching pennies

Convention

The column player's utility can be suppressed in the utility matrix (because it is redundant). If the index of utility function is suppressed, the row player's utility is meant.

Symmetric Nash equilibria

Suppose a population consists of rational players. They play a symmetric game against each other with random pairing. Everybody knows the probability distribution over strategies at a random encounter. A **symmetric** Nash equilibrium is a possible state of such a population.

Definition (Symmetric Nash equilibrium)

A mixed strategy α for a symmetric two-person game is a **symmetric Nash equilibrium** iff

$$U(\alpha, \alpha) \geq U(\alpha', \alpha)$$

for each mixed strategy α' .

If a strategy is strictly better against itself than any other strategy (strict reading), we have a **strict** symmetric Nash equilibrium.

Definition (Strict symmetric Nash equilibrium)

A mixed strategy α for a symmetric two-person game is a **strict symmetric Nash equilibrium** iff

$$U(\alpha, \alpha) > U(\alpha', \alpha)$$

for each mixed strategy α' .