Game Theory 2025

Ulle Endriss
Institute for Logic, Language and Computation
University of Amsterdam

Plan for Today

Pure and mixed Nash equilibria are examples for *solution concepts:* formal models to predict what might be the outcome of a game.

Today we are going to see some more such solution concepts:

- *equilibrium in dominant strategies:* do what's definitely good
- *elimination of dominated strategies:* don't do what's definitely bad
- *correlated equilibrium:* follow some external recommendation

For each of them, we are going to see some *intuitive motivation*, then a *formal definition*, and then an example for a relevant *technical result*.

Most of this (and more) is also covered in Chapter 3 of the *Essentials*.

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008. Chapter 3.

Dominant Strategies

Have we maybe missed the most obvious solution concept? . . .

You should play the action $a_i^\star \in A_i$ that gives you a better payoff than any other action a_i', whatever the others do (such as playing $\boldsymbol{s}_{-i}$):

$$u_i(a_i^\star, \boldsymbol{s}_{-i}) > u_i(a_i', \boldsymbol{s}_{-i}) \text{ for all } a_i' \in A_i \setminus \{a_i^\star\} \text{ and all } \boldsymbol{s}_{-i} \in \boldsymbol{S}_{-i}$$

Action $a_i^\star$ is called a *strictly dominant strategy* for player i.

Profile $\boldsymbol{a}^\star \in \boldsymbol{A}$ is called an *equilibrium in strictly dominant strategies* if, for every player $i \in N$, action $a_i^\star$ is a strictly dominant strategy.

Downside: This does not always exist (in fact, it usually does not!).

Example: Prisoner's Dilemma Again

Here it is once more:

	C	D
C	-10, -10	-25, 0
D	0, -25	-20, -20

Exercise: *Is there an equilibrium in strictly dominant strategies?*

Discussion: Conflict between rationality and efficiency now even worse.

Dominant Strategies and Nash Equilibria

Exercise: *Show that every equilibrium in strictly dominant strategies is also a pure Nash equilibrium.*

Elimination of Dominated Strategies

Action a_i is *strictly dominated* by a strategy $s_i^\star$ if, for all $\boldsymbol{s}_{-i} \in \boldsymbol{S}_{-i}$:

$$u_i(s_i^\star, \boldsymbol{s}_{-i}) > u_i(a_i, \boldsymbol{s}_{-i})$$

Then, if we assume i is *rational*, action a_i can be *eliminated*.

This induces a solution concept:

> *all mixed-strategy profiles of the reduced game that survive iterated elimination of strictly dominated strategies (IESDS)*

Simple example:

	L	R
T	4, 4	1, 6
B	6, 1	2, 2

	R
T	1, 6
B	2, 2

	R
B	2, 2

Order Independence of IESDS

Suppose $A_i \cap A_j = \emptyset$. Then we can think of the *reduced game* G^t after t eliminations simply as the subset of $A_1 \cup \cdots \cup A_n$ that survived.

IESDS says: players will actually play G^∞. *Is this well defined?* Yes!

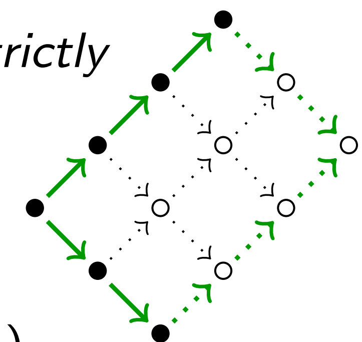

Theorem 1 (Gilboa et al., 1990) *Any order of eliminating strictly dominated strategies leads to the same reduced game.*

Proof: Write $G \twoheadrightarrow G'$ if game G can be reduced to G' by eliminating *one* action. Need to show that trans. closure $\twoheadrightarrow^*$ is *Church-Rosser*. Done if can show that $\twoheadrightarrow$ is C-R (induction!).

Enough to show: if $G \overset{a_i}{\twoheadrightarrow} G'$ and $G \overset{b_j}{\twoheadrightarrow} G''$, then $G' \overset{b_j}{\twoheadrightarrow} G'''$ for some G'''.

$G \overset{b_j}{\twoheadrightarrow} G''$ means there is an $s_j^\star$ s.t. $u_j(s_j^\star, \boldsymbol{s}_{-j}) > u_j(b_j, \boldsymbol{s}_{-j})$ for all $\boldsymbol{s}_{-j}$.
This remains true if we restrict attention to $\boldsymbol{s}'_{-j}$ with $a_i \notin$ *support*(s'_i):
$u_j(s_j^\star, \boldsymbol{s}'_{-j}) > u_j(b_j, \boldsymbol{s}'_{-j})$ for all such $\boldsymbol{s}'_{-j}$. So b_j can be eliminated in G'. ✓

I. Gilboa, E. Kalai, and E. Zemel. On the Order of Eliminating Dominated Strategies. *Operations Research Letters*, 9(2):85–89, 1990.

Let's Play: Numbers Game (Again!)

Let's play this game one more time:

> *Every player submits a (rational) number between 0 and 100.*
> *We then compute the average (arithmetic mean) of all the numbers submitted and multiply that number with 2/3.*
> *Whoever got closest to this latter number wins the game.*

The winner gets ₲100. In case of a tie, the winners share the prize.

Analysis

IESDS results in a reduced game where everyone's only action is 0. So, we happen to find the only pure Nash equilibrium this way.

IESDS works on the assumption of *common knowledge of rationality*. In the *Numbers Game*, we have seen:

- Playing 0 usually is not a good strategy in practice, so assuming common knowledge of rationality must be unjustified.
- When we played the second time, the winning number got closer to 0. So by discussing the game, both your own rationality and your confidence in the rationality of others seem to have increased.

Idea: Recommend Good Strategies

Consider the following variant of the game of the *Battle of the Sexes* (previously, we had discussed a variant with different payoffs):

	A	B
A	2, 1	0, 0
B	0, 0	1, 2

Nash equilibria:

- pure AA: utility $= 2 \;\&\; 1$
- pure BB: utility $= 1 \;\&\; 2$
- mixed $((\frac{2}{3}, \frac{1}{3}), (\frac{1}{3}, \frac{2}{3}))$: EU $= \frac{2}{3} \;\&\; \frac{2}{3}$

$\Rightarrow$ *either unfair or low payoffs*

Ask Rowena and Colin to toss a fair coin and to pick A in case of *heads* and B otherwise. They don't have to, but if they do:

$$\begin{aligned} \text{expected utility} \;&=\; \tfrac{1}{2} \cdot 2 + \tfrac{1}{2} \cdot 1 \;=\; \tfrac{3}{2} \;\text{ for Rowena} \\ &\;\; \tfrac{1}{2} \cdot 1 + \tfrac{1}{2} \cdot 2 \;=\; \tfrac{3}{2} \;\text{ for Colin} \end{aligned}$$

Observe: Nobody has an incentive to deviate from their plan!

Correlated Equilibria

A *random public event* occurs. Each player i receives *private signal* x_i. Modelled as random variables $\boldsymbol{x} = (x_1, \ldots, x_n)$ on $D_1 \times \cdots \times D_n = \boldsymbol{D}$ with *joint* probability distribution π (so the x_i can be correlated).

Player i uses function $\sigma_i : D_i \to A_i$ to translate signals to actions.

A *correlated equilibrium* is a tuple $\langle \boldsymbol{x}, \pi, \boldsymbol{\sigma} \rangle$, with $\boldsymbol{\sigma} = (\sigma_1, \ldots, \sigma_n)$, such that, for all $i \in N$ and all alternative choices $\sigma_i' : D_i \to A_i$, we get:

$$\sum_{\boldsymbol{d} \in \boldsymbol{D}} \pi(\boldsymbol{d}) \cdot u_i(\sigma_1(d_1), \ldots, \sigma_n(d_n)) \geqslant \sum_{\boldsymbol{d} \in \boldsymbol{D}} \pi(\boldsymbol{d}) \cdot u_i(\sigma_1(d_1), \ldots, \sigma_{i-1}(d_{i-1}), \sigma_i'(d_i), \sigma_{i+1}(d_{i+1}), \ldots, \sigma_n(d_n))$$

Interpretation: Player i *controls* whether to play σ_i or σ_i', but has to choose *before* nature draws $\boldsymbol{d} \in \boldsymbol{D}$ from π. She *knows* $\boldsymbol{\sigma}_{-i}$ and π.

Example: Approaching an Intersection

Rowena and Colin both approach an intersection in their cars and each of them has to decide whether to *drive on* or *stop*.

	D	S
D	-10, -10	3, 0
S	0, 3	-2, -2

Nash equilibria:

- pure DS: utility $= 3 \,\&\, 0$
- pure SD: utility $= 0 \,\&\, 3$
- mixed $((\frac{1}{3}, \frac{2}{3}), (\frac{1}{3}, \frac{2}{3}))$: EU $= -\frac{4}{3} \,\&\, -\frac{4}{3}$

$\Rightarrow$ *the only fair NE is pretty bad!*

Could instead use this "randomised device" to get CE:

- $D_i = \{\mathit{red}, \mathit{green}\}$ for both players i
- $\pi(\mathit{red}, \mathit{green}) = \pi(\mathit{green}, \mathit{red}) = \frac{1}{2}$
- recommend to each player to use $\sigma_i : d_i \mapsto \begin{cases} \mathit{drive} & \text{if } d_i = \mathit{green} \\ \mathit{stop} & \text{if } d_i = \mathit{red} \end{cases}$

Expected utility: $\frac{3}{2} \,\&\, \frac{3}{2}$, no incentive to deviate

Correlated Equilibria and Nash Equilibria

Theorem 2 (Aumann, 1974) *For every Nash equilibrium there exists a correlated equilibrium inducing the same distribution over outcomes.*

Proof: Let $\boldsymbol{s} = (s_1, \ldots, s_n)$ be an arbitrary Nash equilibrium.

Define a a tuple $\langle \boldsymbol{x}, \pi, \boldsymbol{\sigma} \rangle$ as follows:

- let the domain of each x_i be $D_i := A_i$
- fix π so that $\pi(\boldsymbol{a}) = \prod_{i \in N} s_i(a_i)$ [the x_i are independent]
- let each $\sigma_i : A_i \to A_i$ be the identity function [i accepts recomm.]

Then $\langle \boldsymbol{x}, \pi, \boldsymbol{\sigma} \rangle$ is the kind of correlated equilibrium we want. ✓

Corollary 3 *Every normal-form game has a correlated equilibrium.*

Proof: Follows from Nash's Theorem. ✓

R.J. Aumann. Subjectivity and Correlation in Randomized Strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

Even More Solution Concepts

There are several other solution concepts in the literature. Examples:

- *Iterated elimination of weakly dominated strategies:* eliminate a_i in case there is a strategy $s_i^\star$ such that $u_i(s_i^\star, \boldsymbol{s}_{-i}) \geqslant u_i(a_i, \boldsymbol{s}_{-i})$ for all $\boldsymbol{s}_{-i} \in \boldsymbol{S}_{-i}$ and this inequality is strict in at least one case.

- *ϵ-Nash equilibrium:* no player can gain more than ϵ in utility by unilaterally deviating from her assigned strategy.

  Exercise: *How does the standard definition of NE relate to this?*

K. Leyton-Brown and Y. Shoham. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008. Chapter 3.

Summary

We have reviewed several *solution concepts* for normal-form games.

- *equilibrium in dominant strategies:* great *if* it exists
- *IESDS:* iterated elimination of strictly dominated strategies
- *correlated equilibrium:* accept external advice

These solution concepts give rise to the following hierarchy:

$$\underbrace{\text{Dom} \subseteq \text{PureNash}}_{\text{might be empty}} \subseteq \underbrace{\text{Nash} \subseteq \text{CorrEq} \subseteq \text{IESDS}}_{\text{always nonempty}}$$

What next? Focus on the special case of zero-sum games.