

Hannes Strass

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Complexity and Correlated Equilibria

Lecture 3, 27th April 2026 // Algorithmic Game Theory, SS 2026

Previously ...

- A **mixed strategy** is a probability distribution on pure strategies.
- In a **mixed Nash equilibrium**, all players play best responses.
- **Nash's Theorem:** Mixed Nash equilibria always exist (for finite games).
- Nash equilibria for concrete two-player games can be obtained via a translation to a **mixed integer program**:
 - Binary variables model the choices of pure strategies to put in the support;
 - real-valued variables model probabilities, utilities, and regret.

Chicken

Two people, **Eli** and **Fyn**, are racing towards each other in cars. Whoever swerves ("chickens out") loses face. If neither swerves, both get seriously injured.

(Eli, Fyn)	Swerve	RaceOn
Swerve	(2, 2)	(1, 3)
RaceOn	(3, 1)	(0, 0)

Overview

Complexity of Computing Nash Equilibria

Correlated Equilibria

Computing Correlated Equilibria

Complexity of Computing Nash Equilibria

Motivation

- Game theory (normal-form games) can be used to **model** situations (in a simplified way).
- Game-theoretic solution concepts (i.e. equilibria) then make **predictions** about the modelled situation. (E.g., “How are players going to act?”)
- In order to actually obtain such predictions, we must be able to **efficiently compute** solution concepts.
- E.g. pure Nash equilibria can be efficiently computed, but do not always exist.
- On the other hand, **mixed** Nash equilibria always exist.

⇒ Can mixed Nash equilibria be efficiently computed?

Complexity Classes for Decision Problems

Recall

- Complexity class P contains all languages (decision problems) that can be decided by polynomial-time bounded deterministic Turing machines.
- Complexity class NP contains all languages $L \subseteq \Sigma^*$ for which there exists a **polynomial verifier**, that is, a polynomial-time bounded deterministic Turing machine M (and a polynomial p) such that:
 - M accepts only pairs (x, y) of words such that
 - $x \in L$, and
 - the length of y is at most polynomial in the length of x (i.e. $|y| \leq p(|x|)$);
 - for every $x \in L$ there is such a pair (x, y) .
- A (polynomial-time) **(many-one) reduction** from A to B is a (deterministic polynomial-time) computable function f such that $w \in A$ iff $f(w) \in B$.
- A language L is **NP-hard** iff all languages in NP can be reduced to L .
- A language is **NP-complete** iff it is NP-hard and in NP.

Complexity Classes for Search Problems

- “Function” complexity classes contain “function” problems $F \subseteq \Sigma^* \times \Sigma^*$:
 - Input:** A word $x \in \Sigma^*$.
 - Output:** Any one $y \in \Sigma^*$ such that $(x, y) \in F$, if such a y exists; “no”, otherwise.
- Output y can be thought of as **solution to a search problem** instance x .
- Solution y need not be unique for x (relation F need not be functional).
- Complexity class FP contains all search problems F where any y with $(x, y) \in F$ can be computed from x in deterministic polynomial time.
- Class FNP contains all F that are accepted by a polynomial verifier.
- A **reduction** from problem A to problem B is a pair (f, g) of (deterministic polynomial-time) computable functions such that $(x, g(y)) \in A$ iff $(f(x), y) \in B$.

Examples

- Given a propositional formula φ , find a satisfying assignment if one exists.
- Given an undirected graph G and a $k \in \mathbb{N}$, find a k -clique in G if one exists.

Formalising the Search Problem

Proposition

Finding a solution for a mixed integer feasibility problem with binary decision variables is FNP-complete.

Corollary

Finding a Nash equilibrium for a two-player normal-form game is in FNP.

Proof.

Via the MIP formulation from Lecture 2. □

Is finding Nash equilibria of n -player normal-form games also FNP-complete? This is **unlikely**, as every game has at least one Nash equilibrium.

Definition

TFNP is the class of all **total** FNP problems: $\text{TFNP} := \{F \in \text{FNP} \mid \forall x \exists y : (x, y) \in F\}$.

Example: Three-Player Coordination

		1					
		A		B			
(2, 3)	A	(2, 2, 2)	(0, 0, 1)	(2, 3)	A	(1, 0, 0)	(1, 0, 1)
	B	(0, 1, 0)	(0, 1, 1)		B	(1, 1, 0)	(1, 1, 1)

To mix strategies, any two players “team up” to make the third player indifferent: E.g. 2, 3 vs. 1: $U_1(A) = 2 \cdot \pi_2(A) \cdot \pi_3(A) = 1 = U_1(B)$, and similarly for the rest.

The properly mixed Nash equilibrium thus has $\pi_1(A) = \pi_2(A) = \pi_3(A) = \frac{1}{\sqrt{2}}$.

Problem

For games with three or more players, equilibrium probabilities can be irrational. How should we represent (numbers involved in) equilibria in a formal language?

Approximate Nash Equilibria

Idea: Allow to ignore potential deviations below a certain payoff threshold ε .

Definition

Let $(P, \mathbf{S}, \mathbf{u})$ be a game in normal form and let $\varepsilon \in \mathbb{R}_{\geq 0}$.

1. A mixed strategy $\pi_i \in \Pi_i$ is an **ε -best response** of player i to the mixed-strategy profile $\boldsymbol{\pi}_{-i} \in \boldsymbol{\Pi}_{-i}$ iff for all mixed strategies $\pi'_i \in \Pi_i$, we have

$$U_i(\pi_1, \dots, \pi_{i-1}, \pi_i, \pi_{i+1}, \dots, \pi_n) \geq U_i(\pi_1, \dots, \pi_{i-1}, \pi'_i, \pi_{i+1}, \dots, \pi_n) - \varepsilon$$

2. A profile $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)$ is a **ε -Nash equilibrium (in mixed strategies)** iff mixed strategy π_i is an ε -best response to $\boldsymbol{\pi}_{-i}$ for all players $1 \leq i \leq n$.

- Thus π_i is an ε -best response to $\boldsymbol{\pi}_{-i}$ iff $U_i(\pi'_i, \boldsymbol{\pi}_{-i}) - U_i(\pi_i, \boldsymbol{\pi}_{-i}) \leq \varepsilon$ for all $\pi'_i \in \Pi_i$.
- Threshold ε can be chosen arbitrarily small, so equilibria can be approximated arbitrarily precisely.
- In particular, $\varepsilon = 0$ leads to “standard” mixed Nash equilibria (Lecture 2).
- Value ε becomes part of the input, represented as rational number in binary.

Finding Some Nash Equilibrium

Nash

Input: A strategic game G in normal form and an $\varepsilon \in \mathbb{Q}^+$.

Output: An ε -Nash equilibrium of G .

Proposition

Nash \in TFNP.

Proof.

- **Nash** \in FNP: Given a game and an equilibrium candidate (a mixed strategy profile $\boldsymbol{\pi}$), we can verify in deterministic polynomial time (for every $i \in P$ and $s, t \in S_i$) that $U_i(\boldsymbol{\pi}_{-i}, s) > U_i(\boldsymbol{\pi}_{-i}, t) + \varepsilon$ implies $\pi_i(t) = 0$.
- **Nash** \in TFNP: By Nash's theorem, an equilibrium always exists. □


Is **Nash** also TFNP-hard? There are no known TFNP-complete problems.

Total Function Problems (Polynomial Verifiers)

Based on the **type of combinatorial argument** establishing totality of the function problem, TFNP can be subdivided into further classes:

- PPA: polynomial parity argument (for undirected graphs)
“If a graph has an odd-degree node it has another one.”
- PLS: polynomial local search
“Every directed acyclic graph has a sink.”
- PPP: polynomial pigeonhole principle
“Any function mapping n elements to $n - 1$ elements has a collision.”
- PPAD: polynomial parity argument for directed graphs
“If a directed graph has a node with different in- and out-degree, then it has another one.”

The known relationships between (some of) the classes are as follows:




End of the Line and Class PPAD

End-of-the-Line

Input: A directed graph $G = (V, E)$ with $V = \{0, 1\}^n$ where every node has in-/out-degree at most one; G is represented by two Boolean circuits (n inputs/outputs) P and S such that $(u, v) \in E$ iff $u = P(v)$ and $S(u) = v$.

Output: A sink or another source ($0^n = P(0^n)$ is always a source).

"Example"





Definition

PPAD is the set of all problems in FNP that can be reduced to **End-of-the-Line**.

End-of-the-Line: Example

Consider the following instance of **End-of-the-Line**:

$G = (V, E)$ with $V = \{0, 1\}^2 = \{00, 01, 10, 11\}$ where for $v \in V$ we denote $v = v_1v_2$.



What is a solution?

Complexity of Computing Nash Equilibria

Theorem (Daskalakis, Goldberg, Papadimitriou, 2006)

Nash is PPAD-complete.

Proof (Sketch).

Makes use of the problem **Brouwer**: Given a continuous function C on a compact, convex set (as an arithmetic circuit), compute a (ε) -fixpoint of C .

- Membership: Reduction of **Nash** to **Brouwer** to **End-of-the-Line**.
 - **Nash** to **Brouwer**: Like in the proof of Nash's theorem.
 - **Brouwer** to **End-of-the-Line**: "Digitalise" C using ε and use Sperner's lemma.
- Hardness: Reduction of **End-of-the-Line** to **Brouwer** to **Nash**.
 - **End-of-the-Line** to **Brouwer**: Embed nodes into \mathbb{R}^3 ; use perturbations of C to represent edges; colour regions of \mathbb{R}^3 to get source/sink-fixpoint correspondence.
 - **Brouwer** to **Nash**: Show that games "can do arithmetic" and simulate C . \square

Complexity of Nash Equilibria: Discussion

- Nash's theorem tells us that mixed (Nash) equilibria always exist.
- This pertains to the usefulness of this solution concept:
It will **always** make a prediction about a game.
- The computational complexity results diminish this:
While a (mixed Nash) equilibrium will exist, it is **hard to obtain**.

↪ We will consider another solution concept that is easier to obtain.

Correlated Equilibria

Correlated Equilibria: Motivation

Traffic Lights

Two cars both want to cross an intersection. If a car stops, it does not get to the other side. If only one car goes, it gets to the other side. If both cars go, there is an accident.

(Car1, Car2)	Stop	Go
Stop	(0, 0)	(0, 1)
Go	(1, 0)	(-100, -100)

- The pure Nash equilibria are (Stop, Go) and (Go, Stop):
In both equilibria, one car never gets to move.
- Another mixed Nash equilibrium is $\left(\left(\frac{100}{101}, \frac{1}{101} \right), \left(\frac{100}{101}, \frac{1}{101} \right) \right)$:
Both cars mostly stop and there is a positive probability of accidents.
- A more desirable outcome would be: $\left\{ (\text{Stop}, \text{Go}) \mapsto \frac{1}{2}, (\text{Go}, \text{Stop}) \mapsto \frac{1}{2} \right\}$:
However, mixed Nash equilibria cannot realise this. **Traffic lights can!**

Correlated Equilibrium: Intuition

- An external device chooses a strategy profile $\mathbf{s} \in \mathcal{S}$ **randomly**.
- The distribution $\psi: \mathcal{S} \rightarrow [0, 1]$ for this is fixed and known to all players.
- For a chosen $(s_1, \dots, s_n) \in \mathcal{S}$, each player $i \in P$ gets **private** advice $s_i \in S_i$.
- Knowing $\{\mathbf{s} \in \mathcal{S} \mid \psi(\mathbf{s}) > 0\}$, player i may be able to infer advice of others.
- A distribution ψ being a correlated equilibrium now means:
Given their knowledge of ψ , no player has an incentive to deviate from the signal's advice (under the assumption that no other player deviates).

Example

In the traffic lights game, assume $\psi = \left\{ (\text{Stop}, \text{Go}) \mapsto \frac{1}{2}, (\text{Go}, \text{Stop}) \mapsto \frac{1}{2}, \dots \right\}$:

- If **Car1** receives signal **Stop**, then it knows **Car2** must have received **Go**.
- Thus its best choice is to **Stop**.
- Symmetrically for **Car1** receiving signal **Go**, and **Car2**.

Correlated Equilibrium: Definition

Definition [Aumann, 1974]

Let $(P, \mathbf{S}, \mathbf{u})$ be a game in normal form with $P = \{1, \dots, n\}$.

- For a strategy profile $\mathbf{s} \in \mathcal{S}$ with $\mathbf{s} = (s_1, \dots, s_n)$ and $i \in P$, we denote $(\mathbf{s})_i := s_i$.
- A probability distribution ψ on $\mathcal{S} = S_1 \times \dots \times S_n$ is a **correlated equilibrium** iff for every $i \in P$, $s_j \in S_j$, and $s_k \in S_i$, we have

$$\sum_{\mathbf{s} \in \mathcal{S}, (\mathbf{s})_i = s_j} \left(\psi(\mathbf{s}) \cdot \left(u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s}) \right) \right) \leq 0$$

Roughly: Following the signal's advice incurs no (positive) regret.

Observation

Every (mixed) Nash equilibrium $\boldsymbol{\pi} = (\pi_1, \dots, \pi_n)$ induces a correlated equilibrium $\psi_{\boldsymbol{\pi}} := \{(s_1, \dots, s_n) \mapsto \pi_1(s_1) \cdot \dots \cdot \pi_n(s_n) \mid (s_1, \dots, s_n) \in \mathcal{S}\}$.

Correlated: Players no longer mix their strategies independently.

Correlated Equilibrium: Example (1)

Battle of the Partners

Two partners, **Cat** and **Dee**, think about how to spend the evening. Each has their personal preference what to do, but overall they want to spend the evening together.

(Cat, Dee)	Cinema	Dancing
Cinema	(10, 7)	(2, 2)
Dancing	(0, 0)	(7, 10)

For the mixed Nash equilibrium $\boldsymbol{\pi} = (\pi_{\text{Cat}}, \pi_{\text{Dee}}) = \left(\left(\frac{2}{3}, \frac{1}{3} \right), \left(\frac{1}{3}, \frac{2}{3} \right) \right)$, we get

$$\psi_{\boldsymbol{\pi}} = \left\{ \begin{array}{l} (\text{Cinema}, \text{Cinema}) \mapsto \frac{2}{9}, (\text{Cinema}, \text{Dancing}) \mapsto \frac{4}{9}, \\ (\text{Dancing}, \text{Cinema}) \mapsto \frac{1}{9}, (\text{Dancing}, \text{Dancing}) \mapsto \frac{2}{9} \end{array} \right\}$$

with $U_{\text{Cat}}(\psi_{\boldsymbol{\pi}}) = U_{\text{Dee}}(\psi_{\boldsymbol{\pi}}) = 4\frac{2}{3}$.

Correlated Equilibria: Example (2)

Assume that both **Cat** and **Dee** have access to the result of one fair coin toss:

- If the coin shows heads, both go to the concert;
- if the coin shows tails, both go to the cinema.

This leads to the following (additional) correlated equilibrium:

$$\psi = \left\{ (\text{Cinema}, \text{Cinema}) \mapsto \frac{1}{2}, (\text{Dancing}, \text{Dancing}) \mapsto \frac{1}{2}, \dots \right\}$$

with associated payoffs $U_{\text{Cat}}(\psi) = U_{\text{Dee}}(\psi) = \frac{1}{2} \cdot 10 + \frac{1}{2} \cdot 7 = 8\frac{1}{2}$.

To verify that ψ is a correlated equilibrium, we (essentially) verify that:

$$\begin{aligned} \psi(\text{Cinema}, \text{Cinema}) \cdot (u_{\text{Cat}}(\text{Dancing}, \text{Cinema}) - u_{\text{Cat}}(\text{Cinema}, \text{Cinema})) &\leq 0 \\ \psi(\text{Dancing}, \text{Dancing}) \cdot (u_{\text{Cat}}(\text{Cinema}, \text{Dancing}) - u_{\text{Cat}}(\text{Dancing}, \text{Dancing})) &\leq 0 \end{aligned}$$

which holds because $\frac{1}{2} \cdot (0 - 10) = -5 \leq 0$ and $\frac{1}{2} \cdot (2 - 7) = -2\frac{1}{2} \leq 0$.

Correlated Equilibria Form a Convex Set

Theorem

Let $G = (P, \mathbf{S}, \mathbf{u})$ be a strategic game in normal form.

For any two correlated equilibria ψ_1 and ψ_2 , and for any $\alpha \in [0, 1]$, we find that $\psi_\alpha := \{\mathbf{s} \mapsto \alpha \cdot \psi_1(\mathbf{s}) + (1 - \alpha) \cdot \psi_2(\mathbf{s}) \mid \mathbf{s} \in \mathcal{S}\}$ is a correlated equilibrium.

Proof.

Let $\alpha \in [0, 1]$ and consider any $i \in P, s_j, s_k \in S_i$. We have

$$\begin{aligned} & \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} (\psi_\alpha(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s}))) \\ &= \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} ((\alpha \cdot \psi_1(\mathbf{s}) + (1 - \alpha) \cdot \psi_2(\mathbf{s})) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s}))) \\ &= \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} \left((\alpha \cdot \psi_1(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s}))) + ((1 - \alpha) \cdot \psi_2(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s}))) \right) \\ &= \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} \left(\alpha \cdot \psi_1(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s})) \right) + \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} \left((1 - \alpha) \cdot \psi_2(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s})) \right) \\ &= \alpha \cdot \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} \left(\psi_1(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s})) \right) + (1 - \alpha) \cdot \sum_{\mathbf{s} \in \mathcal{S}, \mathbf{s}_i = s_j} \left(\psi_2(\mathbf{s}) \cdot (u_i(s_k, \mathbf{s}_{-i}) - u_i(\mathbf{s})) \right) \leq 0 \quad \square \end{aligned}$$

Equilibrium Concepts: Overview

$$\text{DSE} \subsetneq \text{PNE} \subsetneq \text{MNE} \subsetneq \text{CE}$$

- DSE** Dominant Strategy Equilibrium: All players play their dominant strategy.
- PNE** Pure Nash Equilibrium: All players play pure-strategy best responses to the other players' pure strategy profile.
- MNE** Mixed Nash Equilibrium: All players play mixed-strategy best responses to the other players' mixed strategy profile.
- CE** Correlated Equilibrium: All players have non-positive regret when playing as signalled.

Computing Correlated Equilibria

Linear Programming (in a Nutshell)

Definition

- A **linear program** is of the form

$$\begin{aligned} & \text{maximise} && \mathbf{c}^T \mathbf{x} \\ & \text{subject to} && \mathbf{Ax} \leq \mathbf{b}, \\ & && \mathbf{x} \geq 0, \\ & && \text{and } \mathbf{x} \in \mathbb{R}^k \end{aligned}$$

where \mathbf{x} is a vector of **decision variables**, and \mathbf{A} , \mathbf{b} , \mathbf{c} are a matrix and two vectors of real values; the expression $\mathbf{c}^T \mathbf{x}$ is the **objective function**.

- If there is no objective function the program is a **feasibility problem**.
- A **solution** is a variable-value assignment that satisfies all constraints.
- A linear program is a special case of a mixed integer program (Lecture 2).
- Linear programming problems can be solved in deterministic polynomial time.

CE: Linear Programming Formulation

For every pure strategy profile $\mathbf{s} \in \mathcal{S}$, we introduce a decision variable $p_{\mathbf{s}}$ with domain $[0, 1]$ expressing the probability of \mathbf{s} being played.

$$\begin{aligned} &\text{find} && p_{\mathbf{s}_1}, \dots, p_{\mathbf{s}_\ell} && \text{for } \mathcal{S} = \{\mathbf{s}_1, \dots, \mathbf{s}_\ell\} \\ &\text{subject to} && p_{\mathbf{s}} \geq 0 && \text{for all } \mathbf{s} \in \mathcal{S} \\ & && \sum_{\mathbf{s} \in \mathcal{S}} p_{\mathbf{s}} = 1 \\ & && \sum_{\substack{\mathbf{s} \in \mathcal{S}, \\ \mathbf{s}_i = s_j}} p_{\mathbf{s}} \cdot u_i(s_k, \mathbf{s}_{-i}) - \sum_{\substack{\mathbf{s} \in \mathcal{S}, \\ \mathbf{s}_i = s_j}} p_{\mathbf{s}} \cdot u_i(s_j, \mathbf{s}_{-i}) \leq 0 && \text{for all } i \in P \text{ and } s_j, s_k \in S_i \\ & && \text{and } p_{\mathbf{s}_1}, \dots, p_{\mathbf{s}_\ell} \in \mathbb{R} \end{aligned}$$

Correlated equilibria are “easier” to compute than Nash equilibria. (Unless $\text{FP} = \text{PPAD}$.)

Conclusion

Summary

- Approximate Nash Equilibria allow instability up to a certain threshold ϵ .
- Computing an ϵ -**Nash** equilibrium of a normal-form game is **PPAD-complete**.
- A **correlated equilibrium** can be seen as providing players with private signals they can use to best-respond to each other's strategies.
- Computing a **correlated** equilibrium of a normal-form game is **in FP**.

Action Points

- Obtain correlated (non-Nash) equilibria for chicken and interpret them.