

Mixed and Correlated Equilibria

ETH Zürich

Paolo Penna

In this lecture we consider more general equilibrium concepts, namely, *mixed* and (*coarse*) *correlated equilibria*. We extend the definition of price of anarchy to these equilibria and study under which conditions the results on pure Nash equilibria still hold. Our main motivation is

- (a) Pure Nash Equilibria: **may not exist** in some games, and even when they exist they are **hard to compute**. So it is unlikely that players are always be able to converge to one such equilibrium.
- (b) Mixed Nash Equilibria: **always exist**, but they are still **hard to compute**.
- (c) Correlated Equilibria: **always exist** and **easy to compute** (next lecture).

We shall see that the **smooth framework** can be also used for (**coarse**) **correlated equilibria**, and the previous bounds on the price of anarchy extend to these more general equilibria.

1 Pure, Mixed, and Correlated Equilibria

We are going to study extensions of the pure Nash equilibria introduced in the previous lecture. Before giving the formal definitions, we start building some intuition by looking at simple games.

Pure Nash Equilibria (PNE): Each player chooses **one strategy** and no player has a reason to deviate.

Matching Pennies

		A	B
A	1	-1	-1
B	-1	1	1

NO PNE (best response cycle)

Coordination Game

		A	B
A	2	1	0
B	0	0	1

PNE: (AA) and (BB)

Mixed Nash Equilibria (MNE): Each player chooses a **probability distribution** over his/her strategies, and no player has a reason to switch to another strategy.

The numbers in brackets are the probabilities that the player chooses the corresponding strategy.

	$A \left(\frac{1}{2}\right)$	$B \left(\frac{1}{2}\right)$
$A \left(\frac{1}{2}\right)$	-1 1	1 -1
$B \left(\frac{1}{2}\right)$	1 -1	-1 1

MNE (flip fair coins)

	$A \left(\frac{1}{3}\right)$	$B \left(\frac{2}{3}\right)$
$A \left(\frac{1}{2}\right)$	1 2	0 0
$B \left(\frac{1}{2}\right)$	0 0	1 1

MNE

Consider the **row player** given the probabilities used by the other player. The row player is indifferent between the two strategies that he/she is choosing randomly:

- Switch to A : utility = $2 \times \frac{1}{3}$
- Switch to B : utility = $1 \times \frac{2}{3}$
- Play mixed strategy '1/2-1/2': utility = $\frac{2}{3}$

More precisely, these are *expected utilities*:

$$u_i(p) := \sum_{s \in S} p(s) \cdot u_i(s) = \mathbf{E}_{s \sim p} [u_i(s)] .$$

In the games above, we can say that

$$u_i(p) \geq u_i(s'_i, p_{-i})$$

for $s'_i \in \{A, B\}$, where (s'_i, p_{-i}) is the probability distribution in which i plays s'_i with probability 1.

Coarse Correlated Equilibria: A **trusted device** chooses **randomly** one state (one strategy per player), and no player has a reason to switch to another strategy:

$$u_i(p) \geq u_i(s'_i, p_{-i})$$

For cost-minimization games,

$$c_i(p) \leq c_i(s'_i, p_{-i})$$

The **trusted device** uses this distribution over the four states (numbers in brackets):

	A	B
A	$\begin{matrix} 1 \\ 2 \end{matrix} (1/3)$	$\begin{matrix} 0 \\ 0 \end{matrix} (0)$
B	$\begin{matrix} 0 \\ 0 \end{matrix} (0)$	$\begin{matrix} 1 \\ 1 \end{matrix} (2/3)$

CCE

Again, if a player decides a priori to play A (or B), his/her expected utility is not going to improve. Given that the other player(s) agree to accept the device choice, there is no reason to not do so.

Definition 1. An ϵ -approximate coarse correlated equilibrium (or ϵ -coarse correlated equilibrium) of a cost-minimization game is a probability distribution p on the set of states S such that for every player i and every deviation $s'_i \in S_i$ we have

$$\mathbf{E}_{s \sim p} [c_i(s)] \leq \mathbf{E}_{s \sim p} [c_i(s'_i, s_{-i})] + \epsilon .$$

The case of $\epsilon = 0$ is called coarse correlated equilibrium.

Note that the distribution p in the above definition need not be a product distribution like in mixed Nash equilibria.


A stronger notion is that of correlated equilibria (below). Intuitively, it requires that a player still does not want to deviate even after receiving a ‘signal’ s_i by the trusted device (like in a traffic light, if we see ‘Red’ we know that the other cars crossing our street received ‘Green’).

Definition 2. An ϵ -approximate correlated equilibrium (or ϵ -correlated equilibrium) of a cost-minimization game is a probability distribution p on the set of states S such that for every player i , every strategy $s_i \in S_i$, and every deviation $s'_i \in S_i$ we have

$$\mathbf{E}_{s \sim p} [c_i(s) \mid s_i] \leq \mathbf{E}_{s \sim p} [c_i(s'_i, s_{-i}) \mid s_i] + \epsilon .$$

The case of $\epsilon = 0$ is called correlated equilibrium.

Every mixed Nash equilibrium is also a correlated equilibrium, and every correlated equilibrium is also a coarse correlated equilibrium. This leaves us with the following hierarchy of equilibrium concepts:



Unlike pure Nash equilibria, mixed Nash equilibria always exist:

Theorem 3 (Nash). *Every finite game has a mixed Nash equilibrium.*

We next extend the price of anarchy to these equilibria concepts. Because finding a mixed Nash equilibrium is also computationally hard, we will derive a natural algorithm for computing the most general equilibria (coarse correlated).

2 Price of Anarchy (revisited)

We consider *cost-minimization* games like in the previous lecture. That is, each player i has a cost $c_i(s)$ and the *social cost* of a state s is the sum of all players' costs

$$cost(s) = \sum_i c_i(s).$$

When dealing with mixed and correlated equilibria, it is natural to consider the *expected social cost*:

$$cost(p) := \sum_{s \in S} p(s) cost(s) = \mathbf{E}_{s \sim p}[cost(s)] . \tag{1}$$

The Price of Anarchy compares the **worst equilibrium** with the **optimum**. In particular, we will take the worst equilibrium of a **certain type** and consider its expected cost:

Definition 4 (Price of Anarchy). *For a cost-minimization game, the price of anarchy for Eq is defined as*

$$PoA_{\text{Eq}} = \frac{\max_{p \in \text{Eq}} cost(p)}{\min_{s \in S} cost(s)} ,$$


where $cost(p)$ is the expected social cost (1) and Eq is a set of probability distributions over the set of states S .

Observation 5. *Take $\text{Eq} = \text{PNE}$ and observe that this is the Price of Anarchy for pure Nash equilibria in the previous lecture.*

Observation 6. *The hierarchy of equilibrium concepts says that the price of anarchy can get worst when we consider more general notions of equilibria:*

$$PoA_{\text{PNE}} \leq PoA_{\text{MNE}} \leq PoA_{\text{CE}} \leq PoA_{\text{CCE}} .$$

Exercise 1. *Consider the following simple network congestion game with two players:*



Show that $PoA_{MNE} > PoA_{PNE}$ in this game.

Recall that for congestion games with affine latency functions we have proven

$$PoA_{PNE} = 5/2$$

but we also know that pure Nash equilibria are hard to compute.

What do we do with the bounds from previous lecture?

It turns out that whatever bounds we obtained with the “smooth framework”, automatically extend to *all* equilibria in the hierarchy above. Recall the definition of smooth game from last lecture:

Definition 7. A game is called (λ, μ) -smooth for $\lambda > 0$ and $\mu < 1$ if, for every pair of states $s, s^* \in S$, we have

$$\sum_i c_i(s_i^*, s_{-i}) \leq \lambda \cdot cost(s^*) + \mu \cdot cost(s) .$$

Observe that this condition needs to hold for *all* states $s, s^* \in S$, as opposed to only pure Nash equilibria or only social optima. The following theorem says that the bounds for PNE obtained via this technique extend to all equilibria (in particular to the most general ones):

Theorem 8. In a (λ, μ) -smooth game, the PoA for coarse correlated equilibria (PoA_{CCE}) is at most

$$\frac{\lambda}{1 - \mu} .$$

Proof Idea. The proof for pure Nash equilibria (lecture 2) can be adapted. Let s be a coarse correlated equilibrium and s^* be an optimum solution, which minimizes social cost. Then:

$$cost(p) = \mathbf{E}_{s \sim p}[cost(s)] = \mathbf{E}_{s \sim p} \left[\sum_i c_i(s) \right] \quad (\text{definition of social cost})$$

$$\vdots \quad (\mathbf{Exercise!})$$

$$\leq \lambda \cdot cost(s^*) + \mu \cdot cost(p)$$

and by rearranging the terms we get

$$\frac{cost(s)}{cost(s^*)} \leq \frac{\lambda}{1 - \mu}$$

for any $s \in CCE$ and any social optimum s^* . That is, $PoA_{CCE} \leq \frac{\lambda}{1 - \mu}$. □

For congestion games with affine delay functions, PNE are hard to compute.

The above result says that the price of anarchy for coarse correlated equilibria is still 5/2 for these games. We shall see that **coarse correlated equilibria** are **easy to compute** instead.

Exercise 2. Consider the game in Exercise 1. Can this game be a congestion game with affine delay functions for all values of M ?

Recommended Literature

- Tim Roughgarden's lecture notes, <http://theory.stanford.edu/~tim/f13/f13.pdf> (General reference)
 - Chapter 13 for definitions and hierarchy of equilibrium concepts;
- T. Roughgarden. Intrinsic Robustness of the Price of Anarchy. STOC 2009. (Smoothness Framework and PoA)

A significant part of this notes is from last year's notes by Paul Dütting available here:

- http://www.cadmo.ethz.ch/education/lectures/HS15/agt_HS2015/

Exercises

(during this exercise class - 8.10.2019)

We shall discuss and solve together this exercise.

Exercise 3. Consider a symmetric network congestion game with four players. Suppose the network consists of the source s , the target t , and six parallel edges from s to t each with cost function $c(x) = x$. Consider the distribution σ over states that randomizes uniformly over all states with the following properties:

- There is one edge with two players.
- There are two edges with one player each (so three edges are empty).
- The set of edges with at least one player is either $\{1, 3, 5\}$ or $\{2, 4, 6\}$.

Prove that σ is a coarse correlated equilibrium but not a correlated equilibrium.