

Correlated Equilibria

Thomas Kesselheim

Last Update: April 23, 2019

In settings where Nash equilibria are hard to compute, it is at least questionable that players will reach them. Fortunately, there are other, weaker equilibrium concepts that generalize Nash equilibria but are easy to compute. We will introduce two of them today.

1 Hierarchy of Equilibrium Concepts

We have already seen pure and mixed Nash equilibria in cost-minimization games, and observed that every pure Nash equilibrium is also a mixed Nash equilibrium.

Definition 6.1. A correlated equilibrium of a cost-minimization game is a probability distribution p on the set of strategy profiles $S = \prod_{i \in \mathcal{N}} S_i$ such that for every $i \in \mathcal{N}$, every strategy $s_i \in S_i$, and every deviation $s'_i \in S_i$ we have

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

Note that the somewhat sloppy notation $\mathbf{E}_{s \sim p} [c_i(s) \mid s_i]$ and $\mathbf{E}_{s \sim p} [c_i(s'_i, s_{-i}) \mid s_i]$ refers to conditional expectations in which we fix s_i to some value. Equivalently, we could write

$$\frac{\sum_{s_{-i}} p(s_i, s_{-i}) c_i(s_i, s_{-i})}{\sum_{s_{-i}} p(s_i, s_{-i})} \leq \frac{\sum_{s_{-i}} p(s_i, s_{-i}) c_i(s'_i, s_{-i})}{\sum_{s_{-i}} p(s_i, s_{-i})} .$$

Importantly, the distribution p in the above definition need not be a product distribution. That is, the strategies of the different players may be correlated. If they are indeed independent, the equilibrium is a mixed Nash equilibrium. A correlated equilibrium protects against conditional deviations of the form “whenever a player played s_i , he now plays s'_i .”

Definition 6.2. A coarse correlated equilibrium of a cost-minimization game is a probability distribution p on the set of strategy profiles $S = \prod_{i \in \mathcal{N}} S_i$ such that for every $i \in \mathcal{N}$ and every deviation $s'_i \in S_i$ we have

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

A coarse correlated equilibrium differs from a correlated equilibrium in that the expectations are not conditional. That is, it protects only against deviations of the form “whatever a player played, how now plays s'_i .”

One can show that every mixed Nash equilibrium is also a correlated equilibrium, and every correlated equilibrium is also a coarse correlated equilibrium. This leaves us with the hierarchy of equilibrium concepts depicted in Figure 1.

2 Examples

Let's consider the following bimatrix cost-minimization game.




Figure 1: Venn diagram of equilibrium concepts.

	A	B	C
A	-1	1	∞
B	1	-1	∞
C	∞	∞	10

There is exactly one pure Nash equilibrium: (C, C) . There is another mixed Nash equilibrium, in which both players play A and B with probability $\frac{1}{2}$ each. This translates to the probability distribution in which $(A, A), (A, B), (B, A), (B, B)$ are played with probability $\frac{1}{4}$ each.

There are many more correlated and coarse correlated equilibria. For example, it is a correlated equilibrium to play $(A, A), (A, B), (B, A), (B, B)$ with probability $\frac{1}{8}$ each and (C, C) with probability $\frac{1}{2}$. This is not a mixed Nash equilibrium because (A, C) does not get any probability mass despite the row player sometimes playing A and the column player sometimes playing C .

The following distribution is a coarse correlated but not a correlated equilibrium: (A, A) and (C, C) are played with probability $\frac{1}{2}$ each.

3 Computing Correlated Equilibria

One interesting feature of correlated equilibria is that they can be computed efficiently.

Theorem 6.3. *There is an algorithm to compute a correlated equilibrium whose running time is polynomial in the number of strategy profiles.*

Proof. We will devise a linear program whose solutions correspond to correlated equilibria. The number of variables and number of constraints will be polynomial in the number of strategy profiles. We use variables $x(s)$ to denote the probability mass that is put on state s . The equilibrium condition then requires that for all $i \in \mathcal{N}$, and all pairs $s_i, s'_i \in S_i$, we have

$$\frac{\sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i})x(s_i, s_{-i})}{\sum_{s_{-i} \in S_{-i}} x(s_i, s_{-i})} \leq \frac{\sum_{s_{-i} \in S_{-i}} c_i(s'_i, s_{-i})x(s_i, s_{-i})}{\sum_{s_{-i} \in S_{-i}} x(s_i, s_{-i})} .$$

By multiplying both sides by $\sum_{s_{-i} \in S_{-i}} x(s_i, s_{-i})$, we can make these linear constraints.

Computing an equilibrium is actually a search problem, not an optimization problem. However, in this case, we can maximize the overall probability mass as follows.

$$\begin{aligned}
& \max \sum_{s \in S} x(s) \\
& \text{s.t.} \quad \sum_{s_{-i} \in S_{-i}} c_i(s_i, s_{-i}) x(s_i, s_{-i}) \leq \sum_{s_{-i} \in S_{-i}} c_i(s'_i, s_{-i}) x(s_i, s_{-i}) \quad \text{for all } i \in \mathcal{N}, s_i, s'_i \in S_i \\
& \quad \sum_{s \in S} x(s) \leq 1 \\
& \quad x(s) \geq 0 \quad \text{for all } s \in S
\end{aligned}$$

Note that, as there is always a mixed Nash equilibrium, there is a solution of value 1. Any such solution is a correlated equilibrium. So solving the LP we find a correlated equilibrium, which is not necessarily a mixed Nash equilibrium. \square

4 No-Regret Dynamics

We have already seen one piece of evidence why correlated equilibria are a reasonable solution concept. However, it might not be ultimately convincing: How are the players supposed to find the distribution if they each act independently. And even if there is a way to find it, what is the source of the randomization? Who is the arbiter that tells the players what to do?

We will give another, arguably more compelling argument. We will devise a distributed protocol such that if all players follows it they will wind up playing an approximate coarse correlated equilibrium. So, there does not have to be somebody coordinating the actions. Rather it will be the independent behavior of the agents that coordinates itself.

We start with a simple observation. Consider a sequence of strategy profiles $s^{(1)}, \dots, s^{(T)}$ with the property that for all $i \in \mathcal{N}$ and all $s'_i \in S_i$

$$\sum_{t=1}^T c_i(s^{(t)}) \leq \left(\sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)}) \right) + \epsilon T .$$

Then the probability distribution p that chooses one of $s^{(1)}, \dots, s^{(T)}$ uniformly at random is an ϵ -approximate coarse correlated equilibrium, that is,

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

We will call the quantity

$$\max_{s'_i \in S_i} \sum_{t=1}^T c_i(s^{(t)}) - \sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)})$$

the (*external*) *regret* of player i .

Overall, if in a dynamic process each player ensures that her regret is bounded by ϵT , then the players effectively play an ϵ -approximate coarse correlated equilibrium over time. Interestingly, it is possible for a player to ensure that the regret only grows at a rate of $O(\sqrt{T})$. So the longer the process runs, the smaller ϵ becomes.

4.1 Playing against an Adversary

To describe the behavior of a single agent in our dynamics, we will first turn to a different setting. There is only a single player playing against an adversary. Eventually, we will use the adversary to model the behavior and reactions of the other players.

The single player plays T rounds against an adversary, trying to minimize her cost. In each round, the player chooses a probability distribution over N strategies (also termed actions here). After the player has committed to a probability distribution, the adversary picks a cost vector fixing the cost for each of the N strategies.

In round $t = 1, \dots, T$, the following happens:

- The player picks a probability distribution $p^{(t)} = (p_1^{(t)}, \dots, p_N^{(t)})$ over his strategies.
- The adversary picks a cost vector $\ell^{(t)} = (\ell_1^{(t)}, \dots, \ell_N^{(t)})$.
- A strategy $a^{(t)}$ is chosen according to the probability distribution $p^{(t)}$. The player incurs this strategy's cost and gets to know the entire cost vector.

The player's (*external*) (*expected*) *regret* is defined as

$$\mathbf{E} \left[\sum_{t=1}^T \ell_{a^{(t)}}^{(t)} - \min_{j \in [N]} \sum_{t=1}^T \ell_j^{(t)} \right].$$

Let us understand what this definition means. The player and the adversary interact and this way generate the sequences of probability distributions $p^{(t)}$, cost vectors $\ell^{(t)}$, and actions $a^{(t)}$. Now, the player compares her average cost to the cost of a *single action* in *hindsight*. Note that this is possibly a strange point of comparison: If the player had always played i instead of $a^{(t)}$ then the adversary might have reacted differently. This is not taken into consideration here.

Theorem 6.4. *There is an algorithm whose external expected regret is bounded by $O(\sqrt{T \log N})$ if $\ell_i^{(t)} \in [0, 1]$ for all i and t .*

Note that the guarantee means that for larger and larger values of T the regret per step goes to 0. Generally, any such algorithm, whose regret is bounded by some function $o(T)$ is called a *no-regret algorithm*.

We will get to know the algorithm and its analysis in the next lecture. Let us first come back to multiple players and coarse correlated equilibria.

4.2 Back to Multiple Players

In the setting of multiple players, we assume that all players play simultaneously. Their choices will be randomized. So, overall, a random sequence $s^{(1)}, \dots, s^{(T)}$ of strategy profiles will be generated.

To decide what to do in step t each player uses a no-regret algorithm to choose which strategy to play next. There is no adversary but we treat the strategy choices of the other players as if they were adversarial.

In more detail, in every step, we set $\ell_j^{(t)} = c_i(j, s_{-i}^{(t)})$. That is, to define the vector $\ell^{(t)}$ that player i sees in step t , we let the other players choose their strategies. For every possible choice of a strategy for player i , this defines the cost that she will incur: The cost when choosing strategy j will be $c_i(j, s_{-i}^{(t)})$.

The no-regret guarantee holds again whatever the adversary could do in the single-player setting. This is one particular choice. So, the above guarantee holds.

Theorem 6.5. *Let $s^{(1)}, \dots, s^{(T)}$ be a random sequence generated by each player using a no-regret algorithm with regret bound $R^{(T)}$. Then the uniform distribution over $s^{(1)}, \dots, s^{(T)}$ is a $R^{(T)}/T$ -approximate coarse correlated equilibrium.*

Proof. For every player i , the regret bound gives us

$$\mathbf{E} \left[\sum_{t=1}^T c_i(s_i^{(t)}, s_{-i}^{(t)}) - \min_{s'_i \in S_i} \sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)}) \right] \leq R^{(T)} ,$$

or equivalently by linearity of expectation

$$\mathbf{E} \left[\sum_{t=1}^T c_i(s_i^{(t)}, s_{-i}^{(t)}) \right] \leq R^{(T)} + \mathbf{E} \left[\min_{s'_i \in S_i} \sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)}) \right] \leq R^{(T)} + \min_{s'_i \in S_i} \mathbf{E} \left[\sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)}) \right] .$$

Therefore, for the probability distribution over strategy profiles and every $s'_i \in S_i$, we have

$$\mathbf{E}_{s \sim p} [c_i(s)] = \mathbf{E} \left[\frac{1}{T} \sum_{t=1}^T c_i(s^{(t)}) \right] \leq \frac{R^{(T)}}{T} + \mathbf{E} \left[\frac{1}{T} \sum_{t=1}^T c_i(s'_i, s_{-i}^{(t)}) \right] = \frac{R^{(T)}}{T} + \mathbf{E}_{s \sim p} [c_i(s'_i, s_{-i})] .$$

This is exactly the definition of the approximate coarse correlated equilibrium. \square

There are similar results that hold for approximate correlated equilibria. These use a different notion of regret.

Further Reading

- Tim Roughgarden's lecture notes <http://theory.stanford.edu/~tim/f13/1/113.pdf> and lecture video <https://youtu.be/aV16MDoRZoc>