**CS 6840 Algorithmic Game Theory** February 28, 2020

# Lecture 15: Hierarchy of Equilibrium Concepts

*Instructor: Eva Tardos* *Scribes: Matt Eichhorn, Shawn Ong*

In today's lecture, we will discuss different types of equilibrium in a game. We define the Nash condition or analogue for each of these equilibria, and briefly discuss computational issues. We adopt the following conventions:

- Each game has $k$ players.
- Each player $i \in [k]$ has a set of possible strategies $S_i$ (we assume each $S_i$ is finite).
- A strategy vector $\mathbf{s} = (s_1, \dots, s_k)$, with each $s_i \in S_i$ is an assignment of each player to one of their strategies.
- We use the shorthand $(s_i', \mathbf{s}_{-i})$ to refer to the strategy vector where all players use their strategy from $\mathbf{s}$ except player $i$, who uses strategy $s_i'$.
- $c_i : S \to \mathbb{R}$ is the cost incurred by player $i$ when players follow a particular strategy vector.

## Pure Strategy Nash Equilibrium

A strategy vector $\mathbf{s} = (s_1, \dots, s_k)$ is a pure strategy Nash Equilibrium (pure Nash) if

$$c_i(\mathbf{s}) \le c_i(s_i', \mathbf{s}_{-i})$$

for all $i$, and for all $s_i' \in S_i$. Intuitively, no player is able to decrease their cost through unilateral action (choosing another of their strategies while everybody else remains the same).

We have previously studied pure strategy Nash Equilibria, especially in the context of congestion games, where they are guaranteed to exist. As a reminder, a congestion game admits a potential function $\Phi$ with the property that the change in a player's cost by switching strategies is exactly the change in $\Phi$. Therefore, a pure Nash corresponds to a local minimum of $\Phi$, since no possibility local improvement (unilateral action by a player) ensures no player can unilaterally decrease their cost. Since our game has finitely-many players, each with finitely-many strategies, $\Phi$ can take on only finitely-many values, and therefore has a global minimum, and thus at least one local minimum (so a pure Nash).

We consider the question of how hard it is to find a Nash. Note that this is not a standard decision problem (i.e., given a game, does a pure Nash exist?) because, at least for congestion games, the answer is immediate (yes). Therefore, we need to work harder in order to cast this to a more interesting problem. Some questions we can ask:

> What is the Nash Equilibrium with minimal social cost?
>
> Is there a Nash Equilibrium where one player adopts a certain strategy?

These problems turn out to be NP-Complete.

# Mixed Strategy Nash Equilibrium

Many games, such as the 3-player instance of the Hotelling game from a few lectures ago, do not have pure strategy Nash Equilibria, so we must consider a more general type of equilibrium, the mixed strategy Nash Equilibrium (mixed Nash). Here, instead of selecting a single strategy $s_i \in S_i$, player $i$ selects a probability distribution $\sigma_i$ over $S_i$. We denote the mixed Nash by

$$\sigma = \prod_i \sigma_i,$$

where the $\prod$ indicates that each player's probability distribution $\sigma_i$ is independent of the others. In order for $\sigma$ to be a mixed Nash, we must have

$$\mathop{\mathbb{E}}_{\mathbf{s}\sim\sigma}[c_i(\mathbf{s})] \leq \mathop{\mathbb{E}}_{\mathbf{s}\sim\sigma}[c_i(s_i', \mathbf{s}_{-i})] \tag{1}$$

for all $i$ and all $s_i' \in S_i$.

Note that this condition makes the assumption that players are *expectation neutral*; that is, they are indifferent between strategies with the same expected cost. In reality, this is a pretty dumb assumption (although it makes the math work out nicely). As an example, people would most likely not be willing to play a game with a $\frac{1}{1{,}000{,}000}$ of losing \$9,999.99 and a $\frac{999{,}999}{1{,}000{,}000}$ chance of gaining \$0.01, even though the expected payout is \$0.

It may seem surprising that we are allowing players to choose mixed strategies, yet the condition (1) only compares costs to those incurred if $i$ moves to a different *pure* strategy. The following claim shows that the condition presented above is enough to capture the notion that a player cannot unilaterally deviate to a mixed strategy to improve their cost.

**Claim:** Suppose that a player cannot decrease their cost by switching (deterministically) to a (pure) strategy $s_i'$, as in (1). Then, they also cannot decrease their cost by switching to a mixed strategy.

**Proof.** Intuitively, the expected cost of a mixed strategy is an average of the costs of the pure strategies in its support, weighted by its probability distribution; but an average cannot be less than its smallest argument.

Formally, let $\sigma$ be a mixed strategy profile satisfying (1), let $p$ be a mixed strategy for player $i$, and let $p_{s_i'}$ denote the probability of sampling $s_i'$ from $p$. Then,

$$\begin{aligned}
\mathop{\mathbb{E}}_{\substack{\mathbf{s}\in\sigma \\ s_i'\sim p}}[c_i(s_i', \mathbf{s}_{-i})] &= \mathop{\mathbb{E}}_{\mathbf{s}\in\sigma}\left[\sum_{s_i'} p_{s_i'} \cdot c_i(s_i', \mathbf{s}_{-i})\right] \\
&= \sum_{s_i'} p_{s_i'} \cdot \mathop{\mathbb{E}}_{s\in\sigma}[c_i(s_i', \mathbf{s}_{-i})] && \text{(linearity of expectation)} \\
&\geq \sum_{s_i'} p_{s_i'} \cdot \mathop{\mathbb{E}}_{\mathbf{s}\sim\sigma}[c_i(\mathbf{s})] && \text{(by (1))} \\
&= \mathop{\mathbb{E}}_{\mathbf{s}\sim\sigma}[c_i(\mathbf{s})]
\end{aligned}$$

Therefore, $p$ does not decrease expected cost. ■

**Theorem 1.** *(Nash's Theorem) Every finite game has a mixed strategy Nash Equilibrium.*

We will prove this later in the course. Again, this makes reasoning about complexity non-obvious (the decision problem of mixed Nash existence is trivial). But like before, asking more nuanced questions, such as finding the Nash with minimal social cost, or finding a Nash with restrictions on a player's allowable strategies, are NP-complete.

# Correlated Equilibrium

The concept of a Correlated Equilibrium (CE) originates in the work of Robert Aumann (1974), for which (among other things) he was ultimately awarded a Nobel prize in economics. Aumann observed that we may be able to reach a more socially-optimal equilibrium by introducing a third-party coordinating agent. As a motivating example, we consider the Chicken Game.

**Example 1.** *(Chicken Game)*

The chicken game is used to model the idiotic road game "chicken", where players drive toward each other, trying to get the other player to "chicken out" and swerve first to avoid collision. Each player has two strategies: $C$, the chicken strategy, will always swerve, and $D$, the daring strategy, will never swerve. The pay-off bimatrix for the Chicken Game is:

| | $C$ | $D$ |
|---|---|---|
| $C$ | $2$<br>$2$ | $3$<br>$0$ |
| $D$ | $0$<br>$3$ | $-10$<br>$-10$ |

Here, the numbers in the lower-left corner of each matrix cell are the payoff to the row player, and the numbers in the upper-right corner are the payoff to the column player. In particular, note that since we are dealing with payoffs and not costs, higher values are now preferable.

We consider the Nash Equilibria of this game. There are two pure strategy Nash equilibria: $(D, C)$ and $(C, D)$. This is because if either player switches strategies (with the other player keeping the same strategy), their payoff strictly decreases, i.e. chickening is a best response if you are sure your opponent is daring and vice versa. Note that both of these pure strategies have social welfare 3. We claim that there is a third Nash in mixed strategies.

From the argument above, we see that if one player chooses a pure strategy, the other player has a unique pure strategy best response. Therefore, the remaining mixed Nash must have both players mixing between both of their strategies with non-zero probability. Therefore, at this mixed Nash, both players must be indifferent between their strategies (otherwise, they would deterministically chose the one with better payout).

Assume that the column player's mixed strategy chooses $C$ with probability $p$ and $D$ with probability $1-p$. Then, the expected payout to the row player if they choose $C$ is

$$2 \cdot p + 0 \cdot (1-p) = 2p,$$

and their expected payout if they choose $D$ is

$$3 \cdot p - 10 \cdot (1-p) = 13p - 10.$$

Equating these, we find that

$$2p = 13p - 10 \implies p = \frac{10}{11},$$

so the column player must choose $C$ with probability $\frac{10}{11}$ and $D$ with probability $\frac{1}{11}$. By symmetry of the Chicken Game, the row player must adopt the same probability distribution. The expected social welfare of this mixed Nash is

$$4 \cdot \frac{100}{121} + 3 \cdot \frac{20}{121} - 20 \cdot \frac{1}{121} = \frac{440}{121} = \frac{40}{11} \approx 3.636.$$

Now, we consider correlated equilibria for this game. The intuition behind correlated equilibrium is that there is a signal that tells each player what to do, and that after receiving this signal, all players prefer following this instruction rather than deviating. Below are two examples of coordinated equilibria:

(i) Choose uniformly between $(C, D)$ and $(D, C)$.

Note that this model the behavior of a traffic light, allowing exactly one car to pass through the intersection ($C$) while the other is stopped ($D$). The expected social welfare 3.

(ii) Choose $(C, C)$ with probability $\frac{10}{12}$, and $(C, D), (D, C)$ each with probability $\frac{1}{12}$.

Note here that the signal does not relay information about what the other player was told. We argue that this is an equilibrium by showing that neither player can expect a better payout if they do not obey the signal (assuming that the other player does obey). If a player has been told to play $D$, they know exactly which strategy pair is being played, which is a pure Nash, so they are not incentivized to disobey. Suppose that the row player has been told to play $C$. Then, their expected payout if they play $C$ is

$$\frac{2 \cdot \Pr(C, C) + 0 \cdot \Pr(C, D)}{\Pr(C, C) + \Pr(C, D)} = \frac{2 \cdot \frac{10}{12}}{\frac{11}{12}} = \frac{20}{11}$$

and their expected payout if they disobey and play $D$ is

$$\frac{3 \cdot \Pr(C, C) - 10 \cdot \Pr(C, D)}{\Pr(C, C) + \Pr(C, D)} = \frac{3 \cdot \frac{10}{12} - 10 \cdot \frac{1}{12}}{\frac{11}{12}} = \frac{20}{11},$$

so there is no incentive to disobey. Note that we could in fact make following the signal strictly better by slightly perturbing the probabilities to be $10/12 - 2\varepsilon$ and $1/12 + \varepsilon$. Observe that the expected social welfare for this CE is

$$4 \cdot \frac{10}{12} + 3 \cdot \frac{2}{12} = \frac{46}{12} \approx 3.833.$$

Thus, the correlated equilibrium allows us to achieve better social welfare in a way agreeable to all players.

With this example complete, we can formally define the Correlated Equilibrium. A CE is a probability distribution $\sigma$ on strategy vectors such that

$$\mathop{\mathbb{E}}_{\mathbf{s} \sim \sigma} [c_i(\mathbf{s}) \mid s_i] \leq \mathop{\mathbb{E}}_{\mathbf{s} \sim \sigma} [c_i(s_i', \mathbf{s}_{-i}) \mid s_i]$$

for all $i$ and all $s_i' \in S_i$. Thus, when told to do $s_i$ by the signal, it is in player $i$'s best interest to actually do it. Note that, contrary to the definition of the mixed Nash, we do not require that $\sigma$ is a product of independent distributions (it will not be if our signal improves social welfare beyond Nash).

As a final example, we consider a modified form of Rock, Paper, Scissors that disincentivizes ties.

**Example 2.** *(Rock, Paper, Scissors 2.0)*

| | $R$ | $P$ | $S$ |
|---|---|---|---|
| $R$ | $-9$, $-9$ | $-1$, $1$ | $1$, $-1$ |
| $P$ | $1$, $-1$ | $-9$, $-9$ | $-1$, $1$ |
| $S$ | $-1$, $1$ | $1$, $-1$ | $-9$, $-9$ |

There are clearly no pure Nash because a player always has a winning response to any fixed strategy of the other player. Therefore, any Nash equilibrium for this game must be mixed. Suppose that the column player adopts a mixed strategy that only mixes between two symbols (say $R$ and $P$). Then, there is one symbol (in our example, $R$) that guarantees a payout at most $-1$ for the row player, and another symbol ($S$) that guarantees a payout more than $-1$, so the row player will also mix with at most two strategies. We are again caught in the same cycle of best responses. Therefore, both players must mix among all 3 strategies.

Suppose that the column player plays $R$ with probability $r$, $P$ with probability $p$, and $S$ with probability $1 - r - p$. Then, the expected payoff to the row player if they play $R$ is

$$-9r - p + (1 - r - p) = 1 - 10r - 2p,$$

if they play $P$ is

$$r - 9p - (1 - r - p) = 2r - 8p - 1,$$

and if they play $S$ is

$$-r + p - 9(1 - r - p) = 8r + 10p - 9.$$

Equating these, we find that $r = p = \frac{1}{3}$. Therefore, the unique Nash occurs when both players mix uniformly among all three pure strategies. In this case, the expected social welfare is

$$-18 \cdot \frac{1}{3} + 0 \cdot \frac{2}{3} = -6.$$

One example of a CE is to avoid the diagonal symmetric strategy pairs and sample uniformly over the other six off-diagonal strategy pairs. To see that this is a CE, note that when a player is told a symbol, they know that the other player was told one of the other two symbols. By disobeying the signal, they run the risk of paying a penalty of 9, which is not possible if they follow the signal. If the other player does follow the signal, their expected utility in this case would be at most $-4$, strictly worse than the 0 expected by following the signal. In this CE, the expected social welfare has increased to 0.