UPenn NETS 412: Algorithmic Game Theory	February 13, 2018
Lecture 8	
Lecturer: Bo Waggoner	*Scribe: Aaron Roth, Bo Waggoner*

Correlated Equilibria

Consider the following two player traffic light game that will be familiar to those of you who can drive:

	STOP	GO
STOP	0,0	0,1
GO	1,0	-100,-100

This game has two pure strategy Nash equilibria: (GO,STOP), and (STOP,GO) – but these are clearly not ideal because there is one player who never gets any utility.

There is also another mixed strategy Nash equilibrium: Suppose player 1 plays $(p, 1-p)$. If the equilibrium is to be fully mixed, player 2 must be indifferent between his two actions – i.e.:

$$0 = p - 100(1-p) \Leftrightarrow 101p = 100 \Leftrightarrow p = 100/101$$

So in the mixed strategy Nash equilibrium, both players play STOP with probability $p = 100/101$, and play GO with probability $(1-p) = 1/101$. This is even worse! Now both players get payoff 0 in expectation (rather than just one of them), and risk a horrific negative utility. The four possible action profiles have roughly the following probabilities under this equilibrium:

	STOP	GO
STOP	98%	<1%
GO	<1%	≈ 0.01%

A far better outcome would be the following, which is fair, has social welfare 1 (this means the expected sum of the players' utilities is 1), and doesn't risk death:

	STOP	GO
STOP	0%	50%
GO	50%	0%

But there is a problem: *there is no set of mixed strategies that creates this distribution over action profiles.* Therefore, fundamentally, this can never result from Nash equilibrium play.

The reason however is not that this play is not rational – it is! The issue is that we have defined Nash equilibria as profiles of mixed strategies, that require that players randomize independently, without any communication. In contrast, the above outcome requires that players somehow correlate their actions.

Drivers of course do this all the time – the correlating device is a traffic light. The traffic light suggests to each player whether to STOP or GO, and (at least when roads are busy), conditioned on the advice it gives you following its advice is a best response for everyone involved.

Now, our goal is to generalize this into an equilibrium concept. Our definitions of equilibria always look like this: "all players prefer following the equilibrium strategies instead of deviating, given that everyone else is doing so as well." But in this case, what does it mean to deviate? Well, when the player sees the suggestion of the traffic light, she can choose to follow it, or to switch to some other action. In other words, she can deviate to some "swap" function $f : A_i \to A_i$, where when the recommended action is a_i, she instead plays $f(a_i)$.

Definition 1 *A* correlated equilibrium *is a distribution $\mathcal{D}$ over action profiles A such that for every player i, and every function $f : A_i \to A_i$:*

$$\mathop{\mathbb{E}}_{a\sim\mathcal{D}}[u_i(a)] \geq \mathop{\mathbb{E}}_{a\sim\mathcal{D}}[u_i(f(a_i), a_{-i})]$$

In words, a correlated equilibrium is a distribution over action profiles a such that after a profile a is drawn, playing a_i is a best response for player i conditioned on seeing a_i, given that everyone else will play according to a. For example, in the traffic light game, conditioned on seeing STOP, a player knows that his opponents see GO, and hence STOP is indeed a best response. Similarly, conditioned on seeing GO, he knows that his opponents see STOP, and so GO is a best response.

So correlated equilibrium $\mathcal{D}$ models the following timing:

1. A signalling device draws an action profile a from $\mathcal{D}$ and each player i is told their "assigned action" a_i.

2. Each player i chooses the action they actually play. If $\mathcal{D}$ is a correlated equilibrium, then i can compute that, conditioned on being assigned a_i, playing a_i maximizes her expected utility assuming others follow their assigned actions.

Nash equilibria are also correlated equilibria – they are just the special case in which each player's actions are drawn from an independent distribution, and hence conditioning on a_i provides no additional information about a_{-i}. But as we saw above, the set of correlated equilibria is richer than the set of Nash equilibria.

(Notice that a pure-strategy Nash equilibrium a is always a correlated equilibrium, just a boring one: the distribution puts probability 1 on a, and each player always plays according to a. Similarly, a mixed-strategy Nash equilibrium p is also a correlated equilibrium, though this is more subtle: action profile a is drawn with probability $p(a) = p_1(a_1) \cdots p_n(a_n)$. Here it is important that each player sees only her assigned action, not the assigned actions of everyone else. She can Bayesian update on her action to a posterior belief about all others' actions.)

We can define an even larger set still:

Definition 2 *A* coarse correlated equilibrium *is a distribution $\mathcal{D}$ over action profiles A such that for every player i, and every action a'_i:*

$$\mathop{\mathbb{E}}_{a \sim \mathcal{D}}[u_i(a)] \geq \mathop{\mathbb{E}}_{a \sim \mathcal{D}}[u_i(a'_i, a_{-i})]$$

The difference is that, in a coarse correlated equilibrium, we only consider deviations of the form "deviate from following the suggestion, to always playing some fixed strategy a'_i regardless of the suggestion." In other words, following the recommendation need only be a best response in expectation *before* seeing the recommendation a_i. This makes sense if you have to commit to following your suggested action or not up front, and don't have the opportunity to deviate after seeing it. In other words, the timing is:

1. Each player chooses to either commit to a fixed action a'_i, or else to follow the recommendation.

2. A signalling device draws $a \sim \mathcal{D}$. If i chose to follow the recommendation, she plays the recommended a_i; else, she plays the a'_i that she picked.

A coarse correlated equilibrium can for example occasionally suggest that players play obviously stupid actions. Consider the following game, and distribution over action profiles:

	A	B	C
A	(1,1)	(-1,-1)	(0,0)
B	(-1,-1)	(1,1)	(0,0)
C	(0,0)	(0,0)	(-1.1,-1.1)

	A	B	C
A	1/3		
B		1/3	
C			1/3

The payoff for each player for playing according to this distribution is:

$$(1/3) \cdot 1 + (1/3) \cdot 1 - (1/3) \cdot 1.1 = 0.3$$

In contrast the payoff a player would get by playing the fixed action A or B while his opponent randomized would be:

$$(1/3) \cdot 1 - (1/3) \cdot 1 + (1/3) \cdot 0 = 0$$

and the payoff for playing C would be strictly less than zero. Hence, the given distribution is a coarse correlated equilibrium *even though* conditioned on being told to play C, it is not a best response. This means that the given distribution is a coarse correlated equilibrium, *but not* a correlated equilibrium, proving that coarse correlated equilibria are a strictly larger set of distributions.

To recap, we have so far considered several solution concepts: Dominant strategy equilibria (DSE), Pure strategy Nash equilibria (PSNE), mixed strategy Nash equilibria (MSNE), correlated equilibria (CE), and coarse correlated equilibria (CCE), and we know the following strict containments:

$$DSE \subset PSNE \subset MSNE \subset CE \subset CCE$$

where starting at Mixed Nash equilibria, the solution concept is guaranteed to exist (but may still be hard to find). Next lecture, we want to show that starting at correlated equilibria, not only is the solution concept guaranteed to exist, but we can always efficiently compute one.