Lecture 7: Correlated Strategies and Correlated Equilibrium

Lecturer: Swaprava Nath *Scribe(s): Claude Scribe*

Disclaimer: *These notes aggregate content from several texts and have not been subjected to the usual scrutiny deserved by formal publications. If you find errors, please bring to the notice of the Instructor.*

7.1 Recap

Last time we established Nash's theorem: every finite game has a mixed-strategy Nash equilibrium (MSNE), but finding one is computationally hard in general (PPAD-complete). Today we look at an entirely different route to equilibrium, one that does not ask each player to randomize independently, but instead brings in a **mediator** who randomizes on the players' behalf.

7.2 Why a Mediator?

Example 7.1 (Delivery robots at a blind junction) *Two autonomous delivery robots, Chandra and Aryabhatta, approach a blind junction with no traffic signal. Each robot independently chooses to* ***Wait*** *or* ***Go***. *If both wait, nothing moves and both get* 0. *If one goes while the other waits, the mover delivers faster and earns more. If both go at once, they collide and both incur a heavy penalty.*

Chandra \ Aryabhatta	*Wait*	*Go*
Wait	$0, 0$	$1, 2$
Go	$2, 1$	$-10, -10$

Delivery Robots at a Blind Junction

The Nash equilibria of this game are: one robot waits while the other goes (two such PSNE), or both randomize with a large probability of waiting (an MSNE). In practice, something simpler and more reliable happens: a junction controller, the robotic equivalent of a traffic light, tells each robot when to wait and when to go.

Why would the robots agree to follow such a controller? Three reasons motivate the entry of a mediating device:

- It offers an alternative explanation of player rationality, players do not need to compute an equilibrium themselves, they only need to trust and follow a recommendation.
- The utility of every player can improve relative to any Nash equilibrium of the original game, as we will see shortly.

- It is often computationally easier to work with than MSNE.

The controller, or **mediator**, does two things: it randomizes over *strategy profiles* (not over individual strategies in isolation), and it then suggests the corresponding action to each player privately. If the suggested strategies are **enforceable**, in the sense that no one wants to deviate once everyone else is expected to comply, this is called a correlated equilibrium.

7.3 Correlated Strategy and Equilibrium

Definition 7.2 (Correlated strategy) *A* ***correlated strategy*** *is a mapping* $\pi : S_1 \times S_2 \times \cdots \times S_n \to [0,1]$ *such that* $\sum_{s \in S} \pi(s) = 1$.

In Example 7.1, a natural correlated strategy for the junction controller is $\pi(\text{Wait,Wait}) = 0$, $\pi(\text{Wait,Go}) = \pi(\text{Go,Wait}) = \frac{1}{2}$, and $\pi(\text{Go,Go}) = 0$: the controller flips a coin and tells exactly one robot to go.

Definition 7.3 (Correlated equilibrium) *A correlated strategy* π *is a* ***correlated equilibrium (CE)*** *if for every player* $i \in N$ *and every pair of actions* $s_i, s_i' \in S_i$,

$$\sum_{s_{-i} \in S_{-i}} \pi(s_i, s_{-i}) \cdot u_i(s_i, s_{-i}) \geqslant \sum_{s_{-i} \in S_{-i}} \pi(s_i, s_{-i}) \cdot u_i(s_i', s_{-i}).$$

In words: conditional on being told to play s_i, player i's expected utility from complying is at least as large as from deviating to any other s_i', given that everyone else complies. The mediator suggests actions after running its randomization device π, and π itself is common knowledge; only the realized recommendation is private.

7.3.1 Worked Examples

Example 7.4 (Chess or Football) *Two friends, Rohan and Kabir, must coordinate on Chess (C) or Football (F).*

		Kabir	
		C	*F*
Rohan	*C*	2, 1	0, 0
	F	0, 0	1, 2

Chess or Football Game

This game has MSNE $\left(\left(\frac{2}{3}, \frac{1}{3}\right), \left(\frac{1}{3}, \frac{2}{3}\right)\right)$, *with expected utility* $\frac{2}{3}$ *to each player. Consider instead* $\pi(C,C) = \pi(F,F) = \frac{1}{2}$: *conditional on being told C, Rohan's payoff from complying is* 2 *versus* 0 *from deviating to F, so complying is optimal; the same check works for every player and every recommendation. This* π *is indeed a CE, and its expected utility is* $\frac{3}{2}$ *to each player, strictly better than the MSNE's* $\frac{2}{3}$.

Returning to Example 7.1, is $\pi(\text{Wait,Go}) = \pi(\text{Wait,Wait}) = \pi(\text{Go,Wait}) = \frac{1}{3}$ a CE? Are there other CEs of this game? Both are worth checking directly against the definition above.

7.4 Computing Correlated Equilibrium

Finding a CE amounts to solving a system of linear inequalities, a **feasibility linear program**:

$$\sum_{s_{-i}\in S_{-i}} \pi(s_i, s_{-i}) \cdot u_i(s_i, s_{-i}) \geqslant \sum_{s_{-i}\in S_{-i}} \pi(s_i, s_{-i}) \cdot u_i(s_i', s_{-i}),\ \forall s_i, s_i' \in S_i,\ \forall i \in N,$$

$$\pi(s) \geqslant 0,\ \forall s \in S, \qquad \sum_{s\in S} \pi(s) = 1.$$

Assuming $|S_i| = m$ for all $i \in N$, the first line contributes $O(n \cdot m^2)$ inequalities, and there are m^n variables (one per strategy profile). Compare this to MSNE: the number of support profiles that must be checked there is $O(2^{mn})$. Taking logarithms of both quantities shows that CE requires exponentially fewer constraints than the support-enumeration approach to MSNE. As a bonus, the same linear program can optimize an objective over the feasible set of correlated strategies, for example maximizing the sum of the players' utilities, something MSNE computation does not offer directly.

7.5 Comparison with Previous Equilibrium Notions

Theorem 7.5 *For every MSNE σ^*, there exists a CE π^*.*

Proof: Let $\pi^*(s_1, \dots, s_n) = \prod_{i=1}^n \sigma_i^*(s_i)$, the product (independent) distribution induced by σ^*. Since σ^* is an MSNE, for every i and every $s_i' \in S_i$,

$$u_i(\sigma_i^*, \sigma_{-i}^*) \geqslant u_i(s_i', \sigma_{-i}^*).$$

By the MSNE characterization theorem, for every $s_i \in \delta(\sigma_i^*)$ (the support of σ_i^*),

$$u_i(\sigma_i^*, \sigma_{-i}^*) = \sum_{s_i\in\delta(\sigma_i^*)} \sigma_i^*(s_i)\, u_i(s_i, \sigma_{-i}^*) = u_i(s_i, \sigma_{-i}^*),$$

so for all $s_i \in \delta(\sigma_i^*)$ and $s_i' \in S_i$,

$$u_i(s_i, \sigma_{-i}^*) \geqslant u_i(s_i', \sigma_{-i}^*) \implies \sigma_i^*(s_i)\, u_i(s_i, \sigma_{-i}^*) \geqslant \sigma_i^*(s_i)\, u_i(s_i', \sigma_{-i}^*).$$

Summing over $s_{-i} \in S_{-i}$ and rewriting $\sigma_i^*(s_i)\sigma_{-i}^*(s_{-i})$ as $\pi^*(s_i, s_{-i})$ gives exactly the CE inequality for π^*. ■

This places CE strictly above MSNE in generality: every equilibrium concept we have seen so far nests inside the next.

$$SDSE \subseteq WDSE \subseteq PSNE \subseteq MSNE \subseteq CE.$$

Example 7.6 (Two AI Labs) *Two AI labs, Saarthi AI and Drishti AI, are racing to release a new model. Each lab can Test Thoroughly (slower, safer) or Ship Fast (quick, risky). If both test, both build trust and profit well. If one ships fast while the other tests, the fast shipper grabs the market. If both ship fast, a public failure hurts both.*

Saarthi AI \ *Drishti AI*	*Test*	*Ship*
Test	5, 5	0, 6
Ship	6, 0	1, 1

Two AI Labs

What is the CE of this game?

7.6 Summary and Preview

We have covered: normal form games; rationality, intelligence, and common knowledge; strategy and action; dominance (strict and weak) and its equilibria SDSE and WDSE; PSNE from unilateral deviation, generalized to MSNE with Nash's existence guarantee; the characterization and computational hardness of MSNE; and now, correlated strategies and correlated equilibrium, introduced through a trusted mediator, together with the fact that CE is computationally easier to find than MSNE and can weakly Pareto-dominate every MSNE. Next time: richer, sequential representations of games, starting with perfect information extensive form games.

References

[Aum74] R. Aumann, "Subjectivity and correlation in randomized strategies," *Journal of Mathematical Economics*, 1(1), 1974, pp. 67–96.

[MSZ13] M. Maschler, E. Solan and S. Zamir, *Game Theory*, Chapter 8, Cambridge University Press, 2013.