Game Theory
Lecture #13 – Mixed Nash and Correlated Equilibria

Focus of Lecture:

- Mixed Nash Equilibria
- Connection Between Security Strategies and Nash Equilibria
- Correlated Equilibria

1 Introduction

Last lecture focused on investigating the implication of mixed strategies in decision-making. While a Nash equilibrium need not exist in general, last lecture demonstrated that by considering mixed strategies and vNM preferences then a Nash equilibrium is always guaranteed to exist. We will begin this lecture by revisiting mixed strategies in zero-sum games, drawing a connection between security strategies and mixed Nash equilibria. Next, we will turn our attention to studying the efficacy of Nash equilibria from a societal perspective. What are the mechanisms available to an engineer to improve upon the efficiency of Nash equilibria? The second part of this lecture will focus on the use a coordinating mechanisms for this purpose. The use of such coordinating mechanisms requires new equilibrium concepts for analyzing strategic behavior in such contexts.

2 Strategic Form Games with Mixed Strategies

Recall the framework of strategic form games with vNM preferences as introduced in the last lecture. The specific components are as follows:

- **Decision-makers:** There are a collection of decision-makers, i.e., $N = \{1, 2, 3, \ldots, |N|\}$.
- **Choice Sets and Mixed Strategies:** Each decision-maker $i \in N$ is associated with a given choice set $\mathcal{A}_i$. Each decision-maker $i \in N$ is now able to employ a probabilistic mixed strategy $p_i \in \Delta(\mathcal{A}_i)$.
- **vNM Utility Functions:** Each decision-maker $i \in N$ is associated with a given utility function $U_i : \mathcal{A} \to \mathbb{R}$ that defines her preference over the joint actions $\mathcal{A}$. The utility of each agent $i \in N$ given a mixed strategy profile $p = (p_1, \ldots, p_n)$ where $p_i \in \Delta(\mathcal{A}_i)$ for each agent $i \in N$ is defined as

$$U_i(p_1, \ldots, p_n) = \sum_{a \in \mathcal{A}} U_i(a) \times p_1^{a_1} \times \cdots \times p_n^{a_n}. \quad (1)$$

A central part of the last lecture was on extending the definition of best response sets and Nash equilibria to the case of mixed strategies, which is summarized as follows.

Definition 2.1 (Best Response) *The best response of player i to the collective strategy of the other players $\alpha_{-i} \in \Delta(\mathcal{A}_{-i})$ is of the form*

$$B_i(\alpha_{-i}) = \{\alpha_i : U_i(\alpha_i, \alpha_{-i}) \geq U_i(\alpha_i', \alpha_{-i}) \text{ for all } \alpha_i' \in \Delta(\mathcal{A}_i)\} \tag{2}$$

Once again, note that the best response function is actually a set.

Definition 2.2 (Mixed Strategy Nash equilibrium) *A mixed strategy profile $\alpha^* = (\alpha_1^*, ..., \alpha_n^*)$ is a mixed strategy Nash equilibrium if for every player $i \in N$*

$$\alpha_i^* \in B_i(\alpha_{-i}^*)$$

The main takeaway for the last lecture was John Nash's famous result that proved that every game has a Nash equilibrium when considering mixed strategies.

Theorem 2.1 (Nash, 1950) *Every strategic form game with vNM preferences in which each player has finitely many actions has a mixed strategy Nash equilibrium.*

This lecture will begin by re-visiting the framework of zero-sum games where we will identify an important connection between security strategies and Nash equilibria. Next, we will re-visit our Hawk Dove example of the previous lecture with a specific interest in analyzing the efficiency of Nash equilibria. In particular, we will seek to identify whether or not there are alternative forms of collective behaviors that are more desirable from a societal perspective. If so, how can we facilitate the emergence of such behavior?

3 Zero-Sum Games – Revisited

Recall the framework of zero-sum games where we have a payoff matrix M with rows $\mathcal{I}$, columns $\mathcal{J}$, and game matrix elements m_{ij} which corresponds to the entry $M(i,j)$. Given this notation, the computation of the security levels of the two players over pure strategies is of the form

$$\underline{v} = \max_{i \in \mathcal{I}} \min_{j \in \mathcal{J}} m_{ij}, \tag{3}$$

$$\overline{v} = \min_{j \in \mathcal{J}} \max_{i \in \mathcal{I}} m_{ij}, \tag{4}$$

Recall that $\overline{v} \geq \underline{v}$ and there are game instances where $\overline{v} = \underline{v}$, i.e., the game has a value. The computation of security levels over mixed strategies takes on a similar form where we are now guaranteed that $\overline{v} = \underline{v}$.

The question that we want to explore in this section is whether or not these security strategies constitutes a Nash equilibrium. We state the following proposition (with regards to pure strategies) to make this connection clear.

Proposition 3.1 *Consider a zero-sum game with payoff matrix M, rows $\mathcal{I}$, and columns $\mathcal{J}$. The following statements are true:*

(i) If the choice profile (i^*, j^*) is a Nash equilibrium, then the game value and $\overline{v} = \underline{v} = m_{i^*j^*}$.

(ii) If the game has a value, i.e., $\overline{v} = \underline{v}$, then the maximizing and minimizing security strategies, i^* and j^*, are a Nash equilibrium.

This proposition demonstrates the equivalence between security strategies and Nash equilibria in cases where the game has a value. The proof of this proposition is given as follows:

Proof 3.1 *We will begin with the proof of (i). Since (i^*, j^*) is a Nash equilibrium, then we know that i^* is a best response to j^* and vice versa. Since j^* is a best response to i^*, we know that for all $j \in \mathcal{J}$*

$$m_{i^*j^*} \leq m_{i^*j}$$

i.e., $m_{i^*j^*}$ is the smallest element in row i^*. Accordingly,

$$m_{i^*j^*} \leq \min_j m_{i^*j} \leq \max_i \min_j m_{ij} = \underline{v}.$$

Since i^* is a best response to j^*, we know that for all $i \in \mathcal{I}$

$$m_{i^*j^*} \geq m_{ij^*}$$

i.e., $m_{i^*j^*}$ is the largest element in column j^*. Accordingly,

$$m_{i^*j^*} \geq \max_i m_{ij^*} \geq \min_j \max_i m_{ij^*} = \overline{v}.$$

Consequently,

$$m_{i^*j^*} \leq \underline{v} \leq \overline{v} \leq m_{i^*j^*},$$

which completes the proof of (i).

We will now move on to the proof of (ii). Let i^*, j^* be security strategies of ROW *and* COL. *Then we know that for all $i \in \mathcal{I}$ and $j \in \mathcal{J}$*

$$\begin{aligned} m_{i^*j} &\geq \underline{v}, \quad \forall j \in \mathcal{J}, &(5)\\ m_{ij^*} &\leq \overline{v}, \quad \forall i \in \mathcal{I}. &(6)\end{aligned}$$

Since $\overline{v} = \underline{v}$ by hypothesis, we know that

$$m_{B(j^*)j^*} \leq \overline{v} = \underline{v} \leq m_{i^*B(i^*)}$$

where $B(\cdot)$ is best response response function for each player. Hence, it follows that $i^* \in B(j^*)$ and $j^* \in B(i^*)$ which completes the proof.

The above proposition demonstrates that if a zero-sum game has a value, then the security strategies constitutes a Nash equilibrium. However, not all zero-sum games have a value and hence the equivalence is incomplete over pure strategies. However, by the Minimax Theorem we know that every zero-sum game does in fact have a value when considering mixed strategies. Hence, security strategies over mixed strategies always represents a mixed Nash equilibrium!

Theorem 3.1 *Consider a zero-sum game with payoff matrix* M*, rows* $\mathcal{I}$*, and columns* $\mathcal{J}$*. A set of mixed strategies* $p^* \in \Delta(\mathcal{I})$ *and* $q^* \in \Delta(\mathcal{J})$ *are security strategies if and only if* (p^*, q^*) *is a mixed Nash equilibrium.*

The proof of this theorem directly follows the proof of Proposition 3.1. Since security strategies always exist, this result was the first result of a well-defined class of games where a (mixed) Nash equilibrium was guaranteed to exists. For historical perspective, it is important to note that the Minimax Theorem predated Nash's Theorem by over 20 years!

4 Hawk versus Dove – Revisited

We will now shift away from analyzing strategic behavior in zero-sum games to investigating the efficacy of Nash equilibria from a societal perspective. To do that we will revisit the Hawk Dove game discussed in last lecture. Recall that the payoff matrix for the Hawk Dove game is as follows

	H	D
H	$0, 0$	$6, 1$
D	$1, 6$	$3, 3$

where the action H, or hawk, represents aggressive behavior and D, or dove, represents passive behavior. Our analysis in the last lecture identified the following three mixed Nash equilibrium.

(i) $(p = 1, q = 0)$ – This mixed Nash equilibrium captures the behavior where ROW always chooses Hawk, and results in an equilibrium payoff of $(6, 1)$.

(ii) $(p = 0, q = 1)$ – This mixed Nash equilibrium captures the behavior where COL always chooses Hawk, and results in an equilibrium payoff of $(1, 6)$.

(iii) $(p = 3/4, q = 3/4)$ – This mixed Nash equilibrium captures the behavior where both players choose hawk with probability $3/4$. These strategies induce the following joint distribution of play

	H	D
H	$9/16$	$3/16$
D	$3/16$	$1/16$

and each player's expected payoff is given by 1.5.

Note that all three of these mixed Nash equilibria are undesirable from a societal perspective. In particular, the first two Nash equilibria have a fairness issue (one individual is always given the right away) while the last Nash equilibria suffers from a performance issue as each player's expected payoff is only 1.5. A more desirable form of collective behavior would be

	H	D
H	0	1/2
D	1/2	0

which provides each player an expected payoff of 3.5. This behavior essentially boils down to players taking turns in the (H, D) and (D, H) configurations. Clearly, this form of collective behavior requires a degree of collaboration between the players as independent mixed strategies are unable to result in such a joint distribution of play.

As an engineer, it is important that we identify how can we achieve this socially optimal behavior in situations where players are making decisions in a self-interested fashion. Interestingly, we actually observe data in line with this desired distribution of play when we think about actual driver patterns at a traffic intersection. How is that possible? Are there any systems employed to help self-interested drivers coordinate their behavior for the social benefit? Of course there are as we have traffic lights!

Lets start to think a little deeper about the role of traffic lights. Do traffic lights dictate travel patterns? While most people would say yes, the answer is actually no. The traffic light provides a suggestion for driver behavior, with green suggesting go and red suggesting stop. Whether or not a self-interested driver actually follows these suggestion requires a more thorough investigation of whether or not it is in the driver's best interest to do so.

Our solution concept of mixed Nash equilibria is not sufficient to investigate the role of correlating devices on strategic behavior in games. To that end, we introduce the concept of correlated equilibria which captures this phenomena. Informally, there is a "referee" (e.g., traffic signal) that chooses the distribution of play $z \in \Delta(\mathcal{A})$. Each of the players can either go along with the referee suggestions, at which point the expected utility of each player $i \in N$ would be

$$\sum_{a \in \mathcal{A}} z^a \cdot U_i(a)$$

where z^a is the probability of action profile a in the distribution z. A player could also forgo the referees suggestion and commit to an action $\tilde{a}_i \in \mathcal{A}$ which would result in an expected payoff

$$\sum_{a \in \mathcal{A}} z^a \cdot U_i(\tilde{a}_i, a_{-i}),$$

provided all the other players go along with the referees suggestions. Informally, a (coarse) correlated equilibria is a joint distribution of play $z \in \Delta(\mathcal{A})$ where it is better to follow the recommendation of the referee than to commit to a given action.

Definition 4.1 (Coarse Correlated Equilibrium) *A joint distribution $z \in \Delta(\mathcal{A})$ is a coarse correlated equilibrium if for every player $i \in N$ we have*

$$\sum_{a \in \mathcal{A}} z^a \cdot U_i(a) \geq \sum_{a \in \mathcal{A}} z^a \cdot U_i(\tilde{a}_i, a_{-i}), \quad \forall \tilde{a}_i \in \mathcal{A}_i.$$

Note that this definition of coarse correlated equilibrium implies blind following of the referee. That is, a player completely hands over the decision-making to the referee. In more general scenario, such as traffic lights at an intersection, the referee signals the choice that it wants each agent to take and the agent has a decision about whether or not to go along with the signaled choice. This setup is analyzed by a stricter equilibrium concept, termed correlated equilibrium, that gives players the option to condition whether or not they follow the referee recommendation on the dictated choice. More formally, suppose the referee selects an action profile from the joint distribution $z \in \Delta(\mathcal{A})$. Once the referee selects a joint profile $\tilde{a}$ from z, the referee informs each player $i \in N$ the choice they should play $\tilde{a}_i$, which allows each player to form their conditional belief about the distribution z, i.e.,

$$z^{(a_i,a_{-i})}|_{\tilde{a}_i} = \begin{cases} \frac{z^{(\tilde{a}_i,a_{-i})}}{\sum_{a_{-i} \in \mathcal{A}_{-i}} z^{(\tilde{a}_i,a_{-i})}} & \text{if } a_i = \tilde{a}_i \\ 0 & \text{otherwise} \end{cases}$$

as the player in unaware of the dictate choices $\tilde{a}_{-i}$ the other players. Informally, a correlated equilibria $z \in \Delta(\mathcal{A})$ is any distribution of play where it is always better to follow the recommendation of the referee irrespective of the dictated action.

Definition 4.2 (Correlated Equilibrium) *A joint distribution $z \in \Delta(\mathcal{A})$ is a correlated equilibrium if for every player $i \in N$*

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_i} \cdot U_i(a_i, a_{-i}) \geq \sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_i} \cdot U_i(\tilde{a}_i, a_{-i}), \quad \forall a_i, \tilde{a}_i \in \mathcal{A}_i.$$

Note that in order for $z^a|_{a_i}$ to be defined, $z^{\tilde{a}} > 0$ for some action profile of the form $\tilde{a} = (a_i, \tilde{a}_{-i})$. If this is not the case, we will let $z^a|_{a_i} = 0$ for all $a \in \mathcal{A}$.

The definition of correlated equilibrium essentially states that a player's expected utility when the referee recommends actions a_i is optimized when the player actually play a_i as opposed to any other choice $\tilde{a}_i$. While this definition looks complicated, it merely extends the definition of coarse correlated equilibria through this extra conditioning on the joint distributions and expectations.

The following example will shed some light onto these solution concepts.

Example 4.1 *Recall the Hawk Dove example given above. Here we will show that the joint distribution $z \in \Delta(\mathcal{A})$ of the form*

	H	D
H	0	1/2
D	1/2	0

is in fact both a coarse correlated equilibrium and a correlated equilibrium. To show that it is a coarse correlated equilibrium, note that

$$\sum_{a \in \mathcal{A}} z^a \cdot U_1(a) = \sum_{a \in \mathcal{A}} z^a \cdot U_2(a) = 1/2 \cdot 6 + 1/2 \cdot 1 = 3.5$$

Focusing on player 1, if player 1 is committed to either H or D and player 2 follows the referee's recommendation, then the player's expected payoff would be

$$\begin{aligned}
\sum_{a \in \mathcal{A}} z^a \cdot U_1(H, a_{-i}) &= 1/2 \cdot U_1(H, H) + 1/2 \cdot U_1(H, D) = 3 \\
\sum_{a \in \mathcal{A}} z^a \cdot U_1(D, a_{-i}) &= 1/2 \cdot U_1(D, H) + 1/2 \cdot U_1(D, D) = 2
\end{aligned}$$

The analysis of player 2 would provide the same distribution. Hence, this distribution is a coarse correlated equilibrium.

Now lets focus on showing that this distribution is also a correlated equilibrium. Suppose player 1 is told to play H. Given this information, the conditional belief of the distribution $z|_{a_1=H}$ is

	H	D
H	0	1
D	0	0

and the expected utility would be

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_1=H} \cdot U_1(H, a_{-i}) = U_1(H, D) = 6.$$

The expected utility for disobeying any choosing D instead would be

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_1=H} \cdot U_1(D, a_{-i}) = U_1(D, D) = 3.$$

Hence, if player 1 is informed to play H, then it is in the player's best interest to obey. Now, suppose player 1 is told to play D. Given this information, the conditional belief of the distribution $z|_{a_1=H}$ is

	H	D
H	0	0
D	1	0

and the expected utility would be

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_1=D} \cdot U_1(D, a_{-i}) = U_1(D, H) = 1.$$

The expected utility for disobeying and choosing H instead would be

$$\sum_{a_{-i} \in \mathcal{A}_{-i}} z^a|_{a_1=D} \cdot U_1(H, a_{-i}) = U_1(H, H) = 0.$$

Hence, if player 1 is informed to play D, then it is in the player's best interest to obey. Hence, z is also a correlated equilibrium.

5 Conclusion

This lecture covered two main topics. First, we revisited the framework of zero-sum games and demonstrated that if the game has a value, then the security strategies form a Nash equilibrium. Since every zero-sum game has a value over mixed strategies, this means that every zero-sum game has a (mixed) Nash equilibrium. This was the first class of games for which a Nash equilibrium was proven to always exist and predated Nash's famous result by over 20 years.

Our second topic focused on the efficiency of Nash equilibria. Here, we demonstrated that Nash equilibria need not be efficient from a societal perspective. More interestingly, we demonstrating how correlating devices could potentially be employed to try and stabilize self-interested behavior that is more socially desirable. Accordingly, we introduced a solution concept that covers such behaviors termed correlated equilibrium. The take away here is two-fold: (i) self-interested behavior can be highly inefficient from a system-level perspective; (ii) one way to influence societal behavior is to provide a correlating signal. However, understanding the potential impact of such a design choices requires an in depth analysis of the interplay between strategic behavior and these correlating devices which can be modeled through the solution concept of correlated equilibria.

6 Exercises

1. Consider the following two player game. Characterize the set of Nash equilibria, correlated equilibria, and coarse correlated equilibria.

	H	D
H	$0, 0$	$6, 1$
D	$1, 6$	$4, 4$

   This should boil down to constraints on the components of the joint distribution z.

2. Derive the relationship between Nash equilibria, correlated equilibria, and coarse correlated equilibria, illustrating the results using a Venn diagram.