Game Theory
Lecture #11 – Nash Equilibrium

Focus of Lecture:

- Strategic Form Games
- Best Response Sets
- Nash Equilibria

1 Introduction

Last lecture focused on investigating strategic decision-making in finite strategic form games. As we turned our focus to such games, we came to the realization that our previous characterizations of strategic behavior in zero-sum games, i.e., security strategies and security levels, were not reasonable for more general game structures. Accordingly, we introduced the notion of dominant strategies and argued that *if* a game has a dominant strategy, then the dominant strategy represents a reasonable description of strategic behavior. However, are work is far from complete since many games do not have a dominant strategy.

This lecture will continue our focus on characterizing reasonable strategic behavior in environments where players do not have a dominant strategy. In doing so, we will introduce the famous solution concept of *Nash equilibrium*, which was introduced by John Nash in 1950. A Nash equilibrium can be viewed as an action profile where the players are all acting as contingent optimizers. Accordingly, we will view Nash equilibria as a viable solution concept of strategic behavior for situations where a dominant strategy does not exist.

2 Strategic Form Games

Recall the framework of strategic form games introduced in the last lecture. The specific components of a strategic form game are as follows:

- **Decision-makers:** There are a collection of decision-makers, i.e., $N = \{1, 2, 3, \ldots, |N|\}$.
- **Choice Sets:** Each decision-maker $i \in N$ is associated with a given choice set $\mathcal{A}_i$.
- **Joint Choice Sets:** The set of joint choices is defined by $\mathcal{A} = \mathcal{A}_1 \times \cdots \times \mathcal{A}_n$. We will denote a joint choice by the tuple $a = (a_1, a_2, \ldots, a_n) \in \mathcal{A}$ where $a_i \in \mathcal{A}_i$ denotes the choice of player i.
- **Utility Function:** 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 $\mathcal{A}$.

2.1 Best Response Sets

The focus of this lecture will center on identifying a solution concept that provides a reasonable description of the emergent collective behavior in strategic form games. Recall that not all games have a dominant strategy, e.g., first price auctions, and hence we need to refine our belief regarding what constitutes reasonable strategic behavior. Given that players are seeking to optimize their utility functions, a crucial component of strategic decision-making in games has to center around the notion of a best response, defined as follows:

Definition 2.1 (Best Response) *The best response of player i to the action of the other players a_{-i} is a function $B_i : \mathcal{A}_{-i} \rightarrow 2^{\mathcal{A}_i}$ where*

$$B_i(a_{-i}) = \{a_i \in \mathcal{A}_i : U_i(a_i, a_{-i}) \geq U_i(a_i', a_{-i}) \text{ for all } a_i' \in \mathcal{A}_i\}. \quad (1)$$

The best response function defined above highlights the optimal choice for a player *conditioned* on the choices of the other players. Hence, the best response is actually a "set". The following example provides the best response sets for a given payoff matrix.

Example 2.1 *Consider the following two player strategic form game with utility functions*

	L	C	R
T	$4,3$	$5,1$	$6,2$
M	$2,1$	$9,4$	$3,6$
B	$3,0$	$9,6$	$2,8$

The best response for the ROW *player is of the form*

$$\begin{aligned} B_{\text{ROW}}(L) &= \{T\} \\ B_{\text{ROW}}(C) &= \{M, B\} \\ B_{\text{ROW}}(R) &= \{T\} \end{aligned}$$

and the best response for the COL *player is of the form*

$$\begin{aligned} B_{\text{COL}}(T) &= \{L\} \\ B_{\text{COL}}(M) &= \{R\} \\ B_{\text{COL}}(B) &= \{R\} \end{aligned}$$

The input of a best response function is the behavior of the other agent, while the output corresponds to the action choices that maximize the player's payoff given this behavior. Focusing on $B_{\text{ROW}}(C)$ above, note that if COL *is playing C, then* ROW *could obtain a payoff of* 5 *for playing T,* 9 *for playing M, and* 9 *for playing B. Hence, either M or B constitutes an optimal choice for this scenario, hence $B_{\text{ROW}}(C) = \{M, B\}$.*

2.2 Nash Equilibrium

Given the description of a strategic form game, which we often denote by the tuple $G = (N, \{\mathcal{A}_i\}_{i\in N}, \{U_i\}_{i\in N})$, we would like to provide a characterization of reasonable strategic behavior. More formally, we would like to identify a collection of joint actions $\mathcal{A}^{\text{sol}} \subseteq \mathcal{A}$ where each joint action $a^{\text{sol}} \in \mathcal{A}^{\text{sol}}$ is a reasonable and plausible description of the collective behavior. Accordingly, what properties should this set of joint actions possess?

We will start addressing this question by identifying properties that should disqualify a joint action a from being a viable solution. Since players are trying to optimize their utility functions, it seems reasonable that if the joint action a is a viable solution, then any player should not be able to *unilaterally* deviate from a and be better off. More formally, given an action profile $a \in \mathcal{A}$, we will say that a is not a viable solution if there exists a player i with an action a_i' such that

$$U_i(a_i', a_{-i}) > U_i(a_i, a_{-i}),$$

i.e., player i can unilaterally switch from a_i to a_i' and be strictly better off provided all other players continue to play a_{-i}. In terms of best response functions, this condition implies that there exists a player i such that $a_i \notin B_i(a_{-i})$. Ruling out all of these action profiles leads to set of Nash equilibrium, defined as follows:

Definition 2.2 (Nash equilibrium) *An action profile a^* is a Nash equilibrium if for every player $i \in N$,*

$$U_i(a^*) = U_i(a_i^*, a_{-i}^*) \geq U_i(a_i, a_{-i}^*)$$

for every $a_i \in \mathcal{A}_i$.

A Nash equilibrium represents a specific action profile a^* where every agent is a contingent optimizer, i.e., given the choices of other agents, no agent has a unilateral incentive to change strategies. One can also provide an equivalent definition of Nash equilibrium in terms of best response functions as follows:

Definition 2.3 (Nash equilibrium) *An action profile a^* is a Nash equilibrium if for every player $i \in N$,*

$$a_i^* \in B_i(a_{-i}^*)$$

i.e., each player is playing a best response to the actions of the other players.

One way to think about a Nash equilibrium is that it is a "no-regret" point of the game. Suppose a^* is an action profile, and you go to each player one by one *after* a^* is played and ask "do you regret your action, now that you see the actions of the other players?" If no player regrets their action (i.e., no player could have played some other strategy and been better off), then a^* is a Nash equilibrium.

Now that we've formally defined the solution concept of Nash equilibrium, there are several remaining questions that we will seek to address in the coming lectures. First, does a Nash equilibrium constitute a reasonable prediction of societal behavior? Is a Nash equilibrium guaranteed to exist? If so, is there always a unique Nash equilibrium? If there are multiple Nash equilibrium, which Nash equilibrium represents a reasonable description of behavior? We will start with a review of several examples before delving into these questions.

Example 2.2 (Matching Pennies) *Recall the matching pennies game introduced in the zero-sum game lecture with payoff matrix*

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Note that there does not exist a Nash equilibrium for this game as there is always a player that would seek to unilaterally deviate regardless of the specific action profile.

Example 2.3 (Prisoner's Dilemma) *Recall the prisoner's dilemma game introduced in the previous lecture with payoff matrix*

	C	D
C	$-1, -1$	$-4, 0$
D	$0, -4$	$-3, -3$

Note that D strictly dominates C for both players. Accordingly, (D, D) is the unique Nash equilibrium for this game.

Example 2.4 (Social Coordination) *Consider a scenario where a couple is debating about whether to go see Bach or Stravinsky. Either partner would rather be together than apart; however, their preferences over the remaining options differs as one person would rather go to Bach together while the other person would rather go to Stranvinsky together. Accordingly, a payoff matrix modeling these preferences is*

	B	S
B	$2, 1$	$0, 0$
S	$0, 0$	$1, 2$

Observe that there are two Nash equilibria (B, B) and (S, S). Which Nash equilibrium constitutes a reasonable prediction of social behavior?

Example 2.5 (Safety and Social Cooperation) *Consider a scenario where two hunters are going to go hunt together for either Stag or Hare. Stag are much more desirable than Hare, however success in hunting a Stag requires a coordinated effort. Hare, on the other hand, can be successfully hunted in an individual effort. Suppose each hunter drives independently to the hunting ground and is only capable of carrying the hunting accessories for either Stag or Hare, but not both. Without explicit knowledge of what their hunting counterpart will do, each hunter is faced with the following model of strategic interactions*

	Stag	*Hare*
Stag	2, 2	0, 1
Hare	1, 0	1, 1

Observe that there are two Nash equilibria (S, S) *and* (H, H)*. Here,* H *(hare) constitutes the safe choice as a player can get a payoff of* 1 *regardless of what the other player does. On the other hand,* S *(stag) corresponds to the socially cooperative choice where the players can get an improved payoff of* 2 *only if the other player is going from stag as well. Which Nash equilibrium constitutes a reasonable prediction of social behavior?*

Example 2.6 (Social Norms and Conventions) *The last example we will consider pertains to social norms and conventions. In particular, suppose various members of society are debating adopting one of two conventions, the standard and well adopted convention or a new alternative convention that is superior. Here, the benefit associated with choosing a particular convention is intimately tied to the number of societal members choose that convention due to consistency. Consider the layout of a computer keyboard where the primary convention adopted by society is the QWERTY keyboard. Interestingly, some say the Dvorak keyboard is far superior in design to the QWERTY keyboard, however it is much less common. Accordingly, one can model these strategic interactions by a payoff matrix of the form*

	Alt	*Std*
Alt	3, 3	0, 0
Std	0, 0	1, 1

Observe that there are two Nash equilibria (Alt, Alt) *and* (Std, Std)*. Which Nash equilibrium constitutes a reasonable prediction of social behavior?*

2.3 Why Nash equilibrium?

The previous section introduce the solution concept of Nash equilibrium as a reasonable description of behavior in strategic scenarios. The term "equilibrium" suggests that a Nash equilibrium is a rest point, or stable point, of a given dynamical process. Accordingly, consider the following dynamic process known as the Cournot Adjustment Process, which describes a process by which the players are continually seeking to optimize their utility functions.

Definition 2.4 (Cournot Adjustment Process) *Let* $a(0) \in \mathcal{A}$ *denote the joint action profile at time* $t = 0$*. At each time* $t \in \{1, 2, 3, \dots\}$*, the Cournot Adjustment Process chooses the action profile* $a(t) = (a_1(t), \dots, a_n(t)) \in \mathcal{A}$ *according to the rule where for each player* $i \in N$

$$a_i(t) \in B_i(a_{-i}(t-1)).$$

Alternatively, each player selects a best response to the action of the other players at the previous timestep.

Example 2.7 (Matching Pennies) *Recall the matching pennies game highlighted above with payoff matrix*

	H	T
H	$1,-1$	$-1,1$
T	$-1,1$	$1,-1$

If $a(0) = (H,H)$*, then the resulting sequence of joint action profiles chosen according to the Cournot Adjustment Process satisfies* $a(2) = (H,T)$, $a(3) = (T,T)$, $a(4) = (T,H)$, $a(5) = (H,H)$*, and so on.*

Clearly, the Cournot Adjustment Process need not converge to a Nash equilibrium as highlighted above for the matching pennies game. Further, this statement holds true even if a Nash equilibrium exists as one can show from the Stag Hunt game described above, i.e., let $a(0) = (Stag, Hare)$. However, note that a Nash equilibrium is an equilibrium of the Cournot Adjustment Process. That is, if we begin the Cournot Adjustment Process at a Nash equilibrium, it will stay there forever.

2.4 Characterizing a Nash Equilibrium

All of the previous examples of games were quite small, so computing a Nash equilibrium required just an exhaustive search over the joint actions. In general, we will be tasked with characterizing Nash equilibria for broader scenarios where an exhaustive search is not a feasible or desirable approach. Here, we will review two examples where we directly employ the two definitions of Nash equilibria to characterize the Nash equilibrium of the given game.

Example 2.8 (Routing Problem) *Consider a routing problem where there are* $|N|$ *players seeking to traverse over the following two-link network*

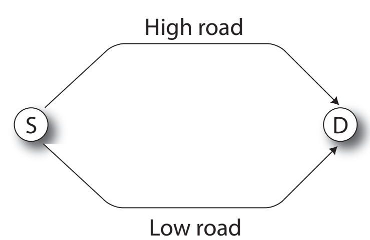

Furthermore, suppose the congestion or latency functions on the high and low road are of the form where for any $k \geq 0$

$$\begin{aligned} c_H(k) &= \bar{c}_H + k \\ c_L(k) &= \bar{c}_L + k \end{aligned}$$

where $\bar{c}_L, \bar{c}_H \geq 0$ *are given constants. Does a Nash equilibrium exist for this routing problem? If so, how do you characterize it?*

Suppose (n_H, n_L) is a division of traffic that represents a Nash equilibrium. Since this behavior constitutes a Nash equilibrium, then we know that all users on the High road would not choose to unilaterally deviate to the Low road, which mathematically implies that

$$\bar{c}_H + n_H \leq \bar{c}_L + n_L + 1,$$

where the "+1" results from the extra driver that results from the unilateral deviation. Given that there are $|N|$ total users, i.e., $n_H + n_L = |N|$, we can rearrange the above equation to obtain

$$2n_H \leq |N| + \bar{c}_L - \bar{c}_H + 1. \quad (2)$$

Since this behavior constitutes a Nash equilibrium, then we also know that all users on the Low road would not choose to unilaterally deviate to the bottom route, i.e.,

$$\bar{c}_L + n_L \leq \bar{c}_H + n_H + 1,$$

which can be rearranged to obtain

$$2n_H \geq |N| + \bar{c}_L - \bar{c}_H - 1. \quad (3)$$

Accordingly, an action profile a is a Nash equilibrium if and only if (2) and (3) are both satisfied. If $|N| = 100$, $\bar{c}_H = 20$, and $\bar{c}_L = 6$, then a Nash equilibrium must satisfy

$$85 \leq 2n_H \leq 87 \quad \Rightarrow \quad n_H = 43,$$

hence a Nash equilibrium is unique in terms of the distribution of drivers on the network. Alternatively, if $|N| = 100$, $\bar{c}_H = 20$, and $\bar{c}_L = 5$, then a Nash equilibrium must satisfy

$$84 \leq 2n_H \leq 86 \quad \Rightarrow \quad 42 \leq n_H \leq 43.$$

Accordingly, a Nash equilibrium could have either $n_H = 42$ or $n_H = 43$. Note that regardless of the parameters, a Nash equilibrium is characterized by a scenario where both roads have almost the same congestion.

The above example demonstrates how one can directly employ the definition of Nash equilibrium pertaining to the players' cost function to characterize Nash equilibrium. The following example will highlight the value of using best response functions to characterize Nash equilibria.

Example 2.9 *Consider a modified version of the routing problem discussed above where a unit mass of flow needs to be routed across the following network*

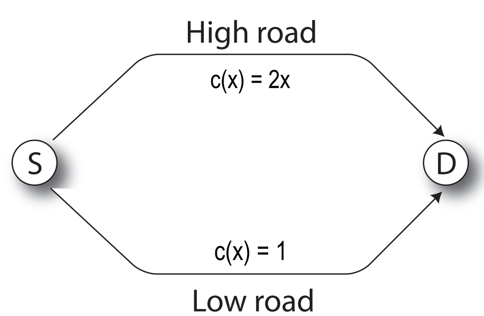

Here, there are two players that each control $1/2$ *units traffic that can divided arbitrarily over the High and Low road. Let* $1/2 \geq f_i^H \geq 0$ *denote the amount of traffic player* i *sends on the High road, which directly implies that* $0.5 - f_i^H$ *is the amount of traffic that the player sends on the Low road. Given a particular joint choice* (f_1^H, f_2^H)*, the cost to player* i *is the total latency experienced by* i*'s traffic, i.e.,*

$$J_i(f_1^H, f_2^H) = f_i^H c_H(f_1^H + f_2^H) + (0.5 - f_i^H) c_L(1 - f_1^H - f_2^H),$$

where we use the notation $J_i(\cdot)$ *for cost as opposed to* $U_i(\cdot)$ *for benefit.*

Here, we seek to identify the routing decisions (f_1^H, f_2^H) *that represent a Nash equilibrium. To do that, we focus on characterizing the best response functions of each player. With regards to player* 1*, this best response function takes on the form*

$$B_1(x) = \underset{0 \leq f_1^H \leq 0.5}{\arg\min} \quad f_1^H \cdot 2(x + f_1^H) + (0.5 - f_1^H),$$

which gives us

$$B_1(x) = \frac{1}{4} - \frac{x}{2}.$$

Since player 1 *and player* 2 *are symmetric we also have*

$$B_2(y) = \frac{1}{4} - \frac{y}{2}.$$

Given these best response functions, we know that a routing profile (f_1^H, f_2^H) *is a Nash equilibrium if and only if*

$$f_1^H = B_1(f_2^H),$$
$$f_2^H = B_2(f_1^H).$$

Plotting each of these functions as in

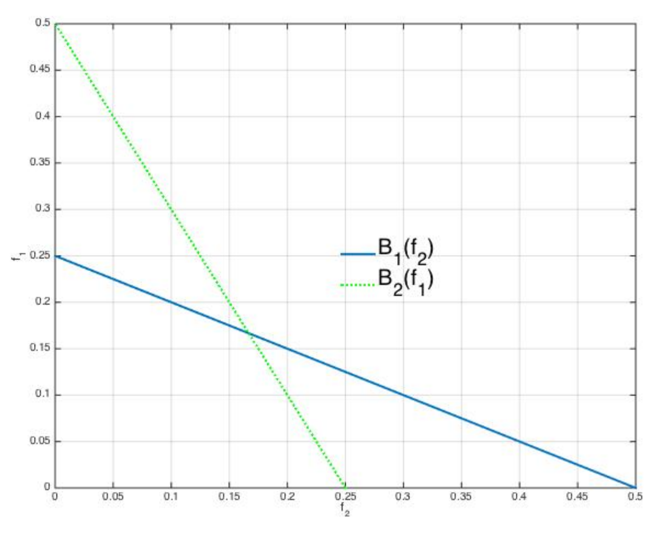

reveals that the Nash equilibrium is a mutual best response where $f_1^H = f_2^H = 1/6$*.*

3 Conclusion

This lecture focused on characterizing strategic behavior in strategic form games. Here, we introduced the solution concept of Nash equilibria, which characterizes a solution where the players are contingent optimizers. We demonstrated examples where a Nash equilibrium did not exists and also where there were multiple Nash equilibria. The following lecture will consider a similar theme were we shift the focus from pure strategies to mixed strategies.

4 Exercises

1. **Tragedy of the Commons:** Suppose 10 families share a plot of land. A goat that grazes on fraction $a \in [0, 1]$ of land produces

$$b = e^{1-\frac{1}{10a}}$$

   bucket of milk. A social planner would like to maximize total milk production. How many goats should be owned among the families to maximize total milk production? Now consider the situation where each family gets to keep only their own milk. Model this situation as a strategic game and identify all Nash equilibria. How does the total milk production change as we transition from a social optimization to a family optimization? Assume throughout that land is divided equally amongst the goats.

2. **Routing:** Consider the routing problem discussed in lecture. In this routing problem, there is 1 unit of divisible traffic that needs to be routed from the start to the destination. There are two possible routing choices either the **High** road or the **Low** road. The cost on the high road is $c_H(x) = x$ where x is the fraction of traffic using the high road. The cost on the low road is $c_L(x) = 1$ for all $x \in [0, 1]$.

   (a) If a social planner controls all traffic, what is the routing profile that minimizes the total cost? The total cost of a routing profile (f_H, f_L) is $f_H c_H(f_H) + f_L c_L(f_L)$ where f_L, f_H are the fractions of traffic on the high and low road.

   (b) Suppose there are two decision makers that each control 1/2 of the traffic. Each decision maker only cares about the total cost of his traffic. Model this situation as a strategic game and analyze the Nash equilibria. How does the total cost compare with the total cost from part (a).

   (c) Suppose there are n decision makers that each control $1/n$ of the traffic. Model this situation as a strategic game and analyze the Nash equilibria. How does the total cost compare with the total cost from part (a). What happens as $n \to \infty$?

3. **Auctions:** An object is to be assigned to a player in the set $\{1, 2, \ldots, n\}$ in exchange for a payment. Player i's valuation of the object is v_i, and $v_1 > v_2 > \cdots > v_n > 0$. The mechanism used to assign the object is a (sealed-bid) auction: the players simultaneously submit bids (nonnegative numbers), and the object is given to the player with the lowest index among those who submit the highest bid, in exchange for a payment. Different auctons differ by the derivation of this payment amount.

(a) **First Price Auction:** In a first price auction the payment that the winner makes is the price that he bids. Formulate a first price auction as a strategic game and analyze its Nash equilibria. In particular, show that in all equilibria player 1 obtains the object.

(b) **Second Price Auction:** In a second price auction the payment that the winner makes is the highest bid among those submitted by the players who did not win (so that if only one player submits the highest bid then the price paid is the second highest bid). Formulate a second price auction as a strategic game. Show that in a second price auction the bid $b_i = v_i$ of any player is a *weakly dominant* action, i.e., player i's payoff when he bids $b_i = v_i$ is at least as high as his payoff when he submits any other bid, regardless of the actions of the other players. Show that nevertheless there are "inefficient" equilibria in which the winner is not player 1.

4. **Dividing money:** Two people have $10 to divide between themselves. They use the following procedure. Each person names a number of dollars (a nonnegative integer), at most equal to 10. If the sum of the amounts that the people names is at most 10, then each person receives the amount of money she named (and the remainder is destroyed). If the sum of the amounts that the people name exceeds 10 and the amounts named are different, then the person who named the smaller amount receives that amount and the other person receives the remaining money. If the sum of the amounts that the people name exceeds 10 and the amounts named are the same, then each person receives 5.

   (a) What is the best response function / curve for either player? Plot them as done in class.

   (b) What are the Nash equilibria of the above game?