Game Theory
Lecture #10 – Strategic Form Games

Focus of Lecture:

- Strategic Form Games
- Dominant Strategies
- Best Response Sets

1 Introduction

The last lecture focused on investigating strategic decision-making in two-player zero-sum games. A zero-sum game models a strategic environment with two players that have diametrically opposed objectives. We asked the question of what constitutes reasonable strategic behavior in such scenarios and explored the viability of security strategies for this purpose. Recall, that the underpinning of security strategies is that each decision-maker assumes a worst-case model of their opponent, and hence security strategies offer the highest possible worst-case guarantees, which are known as security levels. While security strategies can potentially be highly conservative, we demonstrated that security strategies are not conservative in zero-sum games as both players are guaranteed to have the same security level when considering mixed strategies. This means that if one player is playing a security strategy, then the other player's security strategy is also a best response.

In this lecture we will start to consider strategic environments beyond zero-sum games. We term such environments as strategic form games. Here, we will focus on how to model and analyze such strategic environments.

2 Strategic Form Games

In this section we introduce the formal model for strategic form games. These strategic environments involve multiple (≥ 2) decision-makers, each with their own utility function. The strategic component of such scenarios resides in the enmeshment of the players' utility functions, where the behavior of one player potentially impacts the utility of other players. The specifics of the model are as follows:

- **Decision-makers:** There is a collection of decision-makers, i.e., $N = \{1, 2, 3, \ldots, |N|\}$. We will use the terms decision-makers, players, actors, and agents interchangeably throughout the text.
- **Choice Sets:** Each decision-maker $i \in N$ is associated with a given choice set $\mathcal{A}_i$. We will use the terms choices and actions interchangeably throughout the text.

- **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. Lastly, we will often express a joint choice profile a by (a_i, a_{-i}) where $a_{-i} = (a_1, \ldots, a_{i-1}, a_{i+1}, \ldots, a_n)$ encodes the choice of all decision-makers $\neq i$. The set of joint choices for all agents $\neq i$ is given by $\mathcal{A}_{-i} = \prod_{j \neq i} \mathcal{A}_j$
- **Utility Function:** Each decision-maker $i \in N$ is associated with a given utility function $U_i : \mathcal{A} \to \mathbb{R}$ that defines their preference over the joint actions $\mathcal{A}$. We will use the terms utilities, payoffs, rewards, and objectives interchangeably throughout the text. Note that we now have to contend with $|N|$ different utility functions.

As with zero-sum games, matrix form is a convenient representation for two player strategic form games as well. An example of a strategic form game is given by

		COL		
		L	C	R
	T	v, w	$1, 9$	$2, 8$
ROW	M	$3, 7$	$4, 6$	$5, 5$
	B	$6, 4$	x, y	$7, 3$

with the same interpretation as for zero-sum games. Here, the only distinction resides in the fact that $U_1(a) + U_2(a) \neq 0$ for all $a \in \mathcal{A}$. While simplistic, payoff matrices can model complex strategies and interactions as the following examples highlight.

Example 2.1 (Prisoner's Dilemma) *Consider a scenario where two suspects are being interrogated independently; the prosecutors believe they committed a crime together. The interrogators, who lack significant evidence, are relying on these interrogations to determine the extent of the penalty. Each suspect has two options during these interrogations: Cooperate (C) with their partner and do not confess to the crime or Defect (D) from their partner and confess to committing the crime (including all the details of the other suspect's involvement). If both cooperate and do not confess, then the prosecutors only have enough evidence to sentence each to one year in jail. Alternatively, if both defect and confess then that will provide the prosecutors with enough evidence to sentence each to three years in jail. Lastly, if one cooperates and the other defects, then the suspect who defected is released immediately for cooperating with prosecutors, while the partner receives an extensive four year sentence. The resulting payoff matrix is*

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

Note that the suspects' most desirable collective behavior is cooperation, as (C, C) which yields a one year penalty to each suspect. Is this a reasonable prediction of the emergent collective behavior?

Payoff matrices can also model far richer strategic environment encompassing multiple stages and highly sophisticated strategies.

Example 2.2 (Iterated Prisoner's Dilemma) *For example, suppose the prisoner's dilemma game highlighted above is repeated M consecutive stages and the resulting payoff is the sum of the payoffs received in each stage. Now, the actions of a particular player do not represent a single choice of Cooperate or Defect, but rather comprehensive plan of action depending on the behavior of the other player. Example of such comprehensive strategies include:*

- $\mathtt{GT}$ *"grim trigger": Play C until the opponent plays D, then play D forever afterwards*
- $\mathtt{TfT}$ *"tit for tat": At stage k, repeat the opponent's move at stage $k-1$*
- $\mathtt{Cy}$ *"cycle": Play sequence $\{C, D, C, ...\}$*

We can recast new setup in the standard framework where the "action" set for each player is now of the form

$$\{\mathtt{GT}, \mathtt{TfT}, \mathtt{Cy}\}$$

and the resulting payoff is the cumulative payoff which clearly depends on the strategies of the two opposing players. Accordingly, this requires filling in the new matrix game

	$\mathtt{GT}$	$\mathtt{TfT}$	$\mathtt{Cy}$
$\mathtt{GT}$	?, ?	?, ?	?, ?
$\mathtt{TfT}$	?, ?	?, ?	?, ?
$\mathtt{Cy}$	?, ?	?, ?	?, ?

Hence, payoff matrices can capture highly sophisticated strategies in complex environments.

Example 2.3 (Routing Problem) *Consider a routing problem where $|N|$ players are tasked with traversing over a given network from the source to the destination over one of two possible edges, H and L.*

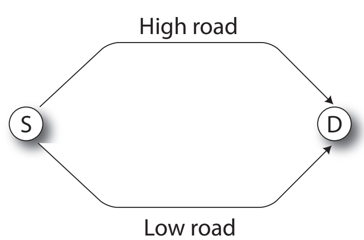

Each edge is associated with a given congestion function (or latency function) that specifies the level of congestion as a function of the level of utilization. The latency of the H (L) edge is given by $c_H : N \to \mathbb{R}$ ($c_L : N \to \mathbb{R}$), where $c_H(k)$ denotes the latency on edge H when there are $k \geq 0$ players on that edge. In the case where $|N| = 2$, this results in the follow cost matrix (as opposed to payoff matrix) of the form

	H	L
H	$c_H(2), c_H(2)$	$c_H(1), c_L(1)$
L	$c_L(1), c_H(1)$	$c_L(2), c_L(2)$

Here, we adopt the convention that players focus on minimizing costs as opposed to maximizing negative utility.

Example 2.4 (First Price Auction) *Consider a situation where there are $N = \{1, 2, \ldots, |N|\}$ individuals participating in an auction for a given good. Each individual $i \in N$ has its own valuation $v_i \geq 0$ for the good and is tasked with making a bid $b_i \geq 0$ for the underlying auction. The auction specifies a protocol that determines the winner and payments associated with a given collection of bids $b = (b_1, \ldots, b_{|N|})$. Informally, in a first price auction the winner is the individual with the highest bid and the winner is charged their bid. Given a bidding profile $b = (b_1, \ldots, b_n)$ with no ties (i.e., $b_i \neq b_j$ for any $i \neq j$), the payoff to each player $i \in N$ is of the form*

$$U_i(b_i, b_{-i}) = \begin{cases} v_i - b_i & \textit{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \textit{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{1}$$

The top condition, i.e., $b_i = \max_{j \in N} b_j$ represents the scenario where player i had the highest bid and won the good and had to pay the price b_i, hence the net benefit to player i is $v_i - b_i$. Note that the utility functions, i.e., preferences, are dependent on the bidding profile $b = (b_1, \ldots, b_{|N|})$ and this bidding profile is what determines the outcome of the auction.

Example 2.5 (Second Price Auction) *Different auctions can vary by the protocol that determines the winner and payments associated with a given collection of bids $b = (b_1, \ldots, b_{|N|})$. In a 2nd-price auction, the winner is the individual with the highest bid; in contrast to the 1st-price auction, here the winner is charged the amount of the 2nd-highest bid. Given a bidding profile $b = (b_1, \ldots, b_n)$ with no ties (i.e., $b_i \neq b_j$ for any $i \neq j$), the payoff to each player $i \in N$ is of the form*

$$U_i(b_i, b_{-i}) = \begin{cases} v_i - \max_{j \neq i} b_j & \textit{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \textit{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{2}$$

3 Strategic Behavior

The only difference between the first price and second price auction is the payment rule. Clearly, this payment rule will impact the strategic behavior of the players and the resulting social behavior. Which auction leads to more efficient societal behavior? How do you even measure the quality of social behavior? To answer these questions we first need to investigate what is a reasonable prediction of social behavior in a given game.

3.1 Security Strategies

Do security strategies represent a reasonable prediction of behavior in strategic form games? To investigate this question, consider the following two player game

	L	R
T	$0,0$	$1,1$
B	$1,1$	$0,0$

Here, in pure strategies, the security levels of both players are 0 and any pure strategy constitutes a security strategy. Hence, if both players play pure security strategies, a possible outcome is (T, L) which would yield a payoff of 0 to both players. Mixed security strategies help a little; here, the mixed security strategies are $(1/2, 1/2)$ for each player, and if both play these strategies the payoffs to each player are $1/2$.

However, given this payoff matrix, it seems highly likely that either the pure action profiles (B, L) or (T, R) would emerge as the result of intelligent strategic behavior. For instance, suppose the row player selected action T; the column player's best response is clearly to play R, which would yield a payoff of 1 to each player.

An even more extreme example is the following:

	L	C	R
T	$0,0$	$1,1$	$0,-1$
M	$1,1$	$0,0$	$-1,0$
B	$-1,0$	$0,-1$	$0,0$

Here, the players' pure security strategies are unique; the row and column players' pure security strategies are T and L, respectively, with security levels of $\underline{v} = \overline{v} = 0$. If both play pure security strategies, the outcome is *always* (T, L), whereas just as in the above example, the players would both be far happier with outcomes (T, C) or (M, L). Hence, security strategies are not necessarily indicative of reasonable and rational play in games beyond zero-sum games.

3.2 Dominant and Dominated Strategies

Reasoning about what constitutes a reasonable description of societal behavior can be a challenging task. Accordingly, in this section we start to see if we can reason about what should *not* constitute a reasonable description of behavior.

Example 3.1 (Prisoner's Dilemma) *Recall the payoff matrix for the prisoner's dilemma game discussed above which is of the form*

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

Is (C,C) a reasonable prediction of the emergent collective behavior? Note that if (C,C) is choses, each player has a unilateral incentive to switch from C to D which would increase their payoff from -1 to 0. Further, note that a player always has an incentive to play D regardless of the behavior of the other player. Accordingly, would a player ever select C is this game was played just a single time?

We now introduce the concept of weakly and strictly dominant strategies which formalize the intuition highlighted above for the prisoner's dilemma game.

Definition 3.1 (Strictly Dominant Strategy) *The action a_i' strictly dominates action a_i if $U_i(a_i', a_{-i}) > U_i(a_i, a_{-i})$ for all $a_{-i} \in \mathcal{A}_{-i}$. Alternatively, we will say that a_i is strictly dominated by a_i'.*

Definition 3.2 (Weakly Dominant Strategy) *The action a_i' **weakly dominates** action a_i if $U_i(a_i', a_{-i}) \geq U_i(a_i, a_{-i})$ for all $a_{-i} \in \mathcal{A}_{-i}$ and there exists at least one $a_{-i} \in \mathcal{A}_{-i}$ such that $U_i(a_i', a_{-i}) > U_i(a_i, a_{-i})$. Alternatively, we will say that a_i is weakly dominated by a_i'.*

Given the above definitions, a strategy a_i' (strictly) dominates a strategy a_i if the strategy a_i' performs strictly better than the strategy a_i for all possible choices of the other players a_{-i}. Accordingly, if a strategy is strictly dominated by some other strategy it seems highly unlikely that this strategy would be employed by a strategic decision-maker. While not all games possess a strictly dominant or dominated strategy, the existence of such a strategy could be extremely valuable in reasoning about the likely emergent behavior. In the above prisoner's dilemma game, note that D strictly dominates C.

The takeaway from this section is that one can remove strictly dominated actions from consideration when arguing about reasonable strategic behavior. Iteratively eliminating such strategies can greatly simplify the analysis of a given game as shown in the following example.

Example 3.2 (Successive Iteration of Strictly Dominated Strategies) *Consider the following two player strategic form game with utility functions*

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

Note that ROW *player has no strictly dominated strategies. However, the same does not hold true for* COL *as the choice C is strictly dominated by R. Accordingly, we can remove the choice C from the choice set of* COL *which results in the following reduced payoff matrix*

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

Since C should never be played by COL, *we can now focus on the remaining payoff matrix and see if there are now any further strictly dominated strategies in this reduced payoff matrix. Observer, that* ROW *now has two strictly dominated strategies, as T strictly dominates both M and B. Accordingly, since* COL *will not play C, we can then argue that* ROW *will not play either M or B, which results in the following reduced payoff matrix*

	L	R
T	$4,3$	$6,2$

Continuing in the same fashion, observe that COL *will not player R, which suggests that (T, L) is the most reasonable description of societal behavior.*

Note that the above example pertains to the elimination of *strictly* dominated strategies, which seems reasonable given the above definition. However, it need not be the case that *weakly* dominated strategies should removed from consideration; nonetheless, they still can be valuable in reasoning about the emergent collective behavior.

Example 3.3 (Second Price Auction) *Recall the framework of second price auctions discussed above where there are $N = \{1, 2, \ldots, |N|\}$ individuals participating in an auction for a given good. Recall that given a bidding profile $b = (b_1, \ldots, b_n)$ with no ties (i.e., $b_i \neq b_j$ for any $i \neq j$), the payoff to each player $i \in N$ is of the form*

$$U_i(b_i, b_{-i}) = \begin{cases} v_i - \max_{j \neq i} b_j & \text{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \text{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{3}$$

Now we will show that the bidding strategy where each player bids its true valuation, i.e., $b_i = v_i$, is a weakly dominant strategy. To see this, we will show that for all bidding profiles b_{-i} and bids b_i we have

$$U_i(b_i = v_i, b_{-i}) \geq U_i(b_i, b_{-i}). \tag{4}$$

There are several cases that one needs to consider for this problem relating to the terms b_i, v_i, and b_{-i} to show that $b_i = v_i$ is a weakly dominant strategy. First note that the entire bid profile b_{-i} is not necessarily important as only the maximum bid in this collection is necessary, which we define as $\bar{b} = \max_{j \neq i} b$. Now, lets consider whether having a bid $b_i > v_i$ could ever be beneficial. Given the set of parameters, there are three possible options that one needs to consider: (i) $b_i > v_i > \bar{b}$, (ii) $b_i > \bar{b} > v_i$, and (iii) $\bar{b} > b_i > v_i$. For each of these scenarios, we are tasked with showing that (4) holds. Consider situation (ii) as an illustration where we have

$$\begin{aligned} U_i(b_i = v_i, b_{-i}) &= 0 \\ U_i(b_i, b_{-i}) &= v_i - \bar{b} < 0 \end{aligned}$$

where the first expression is 0 *because individual* i *did not win the object while the second expression is* $v_i - \bar{b}$ *since individual* i *wins the object at a price* $\bar{b} > v_i$. *Continuing on in this fashion shows that individual* i *does not ever have an incentive to submit a bid* $b_i > v_i$. *A similar analysis can be conducted to show that individual* i *also does not ever have an incentive to submit a bid* $b_i < v_i$.

Here, we've seen two examples of games that have a dominant strategy, either strict or weak. The following is an example of a game where a dominant strategy does not exist.

Example 3.4 (First Price Auction) *Recall the framework of first price auctions discussed above where there are* $N = \{1, 2, \ldots, |N|\}$ *individuals participating in an auction for a given good. Recall that given a bidding profile* $b = (b_1, \ldots, b_n)$ *with no ties (i.e.,* $b_i \neq b_j$ *for any* $i \neq j$*), the payoff to each player* $i \in N$ *is of the form*

$$U_i(b_i, b_{-i}) = \begin{cases} v_i - b_i & \textit{if} \quad b_i = \max_{j \in N} b_j \\ 0 & \textit{if} \quad b_i < \max_{j \in N} b_j \end{cases} \tag{5}$$

Here, lets explore whether is $b_i = v_i$ *a weakly or strictly dominant strategy for this scenario. To that end, lets consider whether having a bid* $b_i < v_i$ *could ever be beneficial. Given the set of parameters, there are three possible options that one needs to consider: (i)* $b_i < v_i < \bar{b}$*, (ii)* $b_i < \bar{b} < v_i$*, and (iii)* $\bar{b} < b_i < v_i$. *Focusing on scenario (iii), we have that*

$$\begin{aligned} U_i(b_i = v_i, b_{-i}) &= 0 \\ U_i(b_i, b_{-i}) &= v_i - b_i > 0 \end{aligned}$$

Hence, having $b_i < v_i$ *is strictly better than* $b_i = v_i$. *This implies that* $b_i = v_i$ *is not a dominant strategy. Continuing in a same fashion could show that there are no dominant strategies (either weakly or strictly) in first price auctions.*

3.3 Best Response Sets

Not all games have a dominant strategy as we saw in the case of first price auctions. Hence, we need to refine our belief of what constitutes reasonable and strategic play. 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 3.3 (Best Response) *The best response of player* i *to the action of the other players* a_{-i} *is*

$$B_i(a_{-i}) = \{a_i : U_i(a_i, a_{-i}) \geq U_i(a_i', a_{-i}) \textit{ for all } a_i' \in \mathcal{A}_i\}\,. \tag{6}$$

Note that the best response is actually a "set".

The best response function defined above highlights the optimal choice for a player *conditioned* on the choices of the other players. Best response sets will play a pivotal role in analyzing strategic behavior in forthcoming lectures. Here, we will provide some examples to make sure this concept is clear.

Example 3.5 (Second Price Auction) *Recall the framework of second price auctions discussed above. The above analysis shows that $v_i \in B_i(b_{-i})$ for all b_{-i}. Note that there may in fact be other bids in the best response set of player.*

Example 3.6 *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*

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

Note that the input to the best response function is the behavior of the other agents, which 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 constitute an optimal choice for this scenario, hence $B_{\text{ROW}}(C) = \{M, B\}$.*

4 Conclusion

This lecture focused on strategic decision-making in strategic form games. We demonstrated that security strategies are not necessarily reasonable strategies in scenarios beyond zero-sum games. We then turned the question from what constitutes reasonable behavior to what does not constitute reasonable behavior. Along this line, we introduced the notion of dominant strategies and demonstrated how ruling out dominated strategies can be an extremely valuable tool for analyzing strategic behavior. We concluded by looking at best response sets and will learn in the next lecture how these can be used to provide a reasonable model of the emergent social behavior.

5 Exercises

1. Consider a three player game where each player has two moves, i.e., $a_i \in \{0,1\}$, and the payoff to the first player is
$$u_1(a_1, a_2, a_3)$$
Suppose each player independently uses a mixed strategy $(p_i, 1-p_i)$, where p_i is the probability of player i selecting 0. Write an expression for the expected utility for player 1.

2. Consider a two player game where the action sets are $\mathcal{A}_1 = \{T, M, B\}$ and $\mathcal{A}_2 = \{L, C, R\}$ and a payoff matrix of the form

	L	C	R
T	1,2	3,1	1,2
M	1,1	0,0	3,2
C	2,2	1,2	0,1

   (a) Suppose player two is using a strategy $p_2 = (p_2^L = 1/3, p_2^C = 1/3, p_2^R = 1/3) \in \Delta(\mathcal{A}_2)$, i.e., player 2 uses a strategy that selects each action with probability 1/3. What is the expected utility of player 1 when playing each of the three actions T, M, and B conditioned on player two playing the strategy p_2?

   (b) What is the probability mass function over the joint action space $\mathcal{A}$ when player 1 uses the strategy $p_1 = (1/2, 1/4, 1/4)$ and player 2 uses the strategy p_2 above?

   (c) What is the expected utility of player 1 and player 2 when player 1 and player 2 use the strategies p_1 and p_2 highlighted above?

3. Consider a two player game where the action sets are $\mathcal{A}_1 = \{T, M, B\}$ and $\mathcal{A}_2 = \{L, C, R\}$. Consider the following distribution over the joint action set $\mathcal{A}$

	L	C	R
T	1/16	1/24	1/48
M	1/4	1/6	1/12
C	3/16	1/8	1/16

   Can the above joint distribution be realized by mixed strategies $p_1 \in \Delta(\mathcal{A}_1)$ and $p_2 \in \Delta(\mathcal{A}_2)$ where each player selects the action independently in accordance with their strategy? If so, what are the players' strategies?

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. Model this scenario as a strategic form game and provide the payoff matrix.

5. **Grad School Competition:** Two students sign up to prepare an honors thesis with a professor. Each can invest time in his own project: either no time, one week, or two weeks (these are the only three options). The cost of time is 0 for no time, and each week costs 1 unit of payoff. The more time a student puts in the better his work will be, so that if one student puts in more time than the other there will be a clear "leader." If they put in the same amount of time then their thesis projects will have the same quality. The professor, however, will give out only one grade of A. If there is a clear leader, then he will get the A, while if they are equally good then the professor will toss a fair coin to decide who gets the A. The other student will get a B. Since both wish to continue on to graduate school, a grade of A is worth 3 while a grade of B is worth 0. Model this scenario as a strategic form game and provide the payoff matrix.