Signaling Games

Farhad Ghassemi

Abstract - ***We give an overview of signaling games and their relevant solution concept, perfect Bayesian equilibrium. We introduce an example of signaling games and analyze it.***

1 Introduction

In the general framework of *incomplete information* or *Bayesian games*, it is usually assumed that information is equally distributed among players; i.e. there exists a commonly known probability distribution of the unknown parameter(s) of the game. However, very often in the real life, we are confronted with games in which players have *asymmetric information* about the unknown parameter of the game; i.e. they have different probability distributions of the unknown parameter. As an example, consider a game in which the unknown parameter of the game can be measured by the players but with different degrees of accuracy. Those players that have access to more accurate methods of measurement are definitely in an advantageous position.

In extreme cases of asymmetric information games, one player has complete information about the unknown parameter of the game while others only know it by a probability distribution. In these games, the information is completely one-sided. The informed player, for instance, may be the only player in the game who can have different types and while he knows his type, other do not (e.g. a prospect job applicant knows if he has high or low skills for a job but the employer does not) or the informed player may know something about the state of the world that others do not (e.g. a car dealer knows the quality of the cars he sells but buyers do not).

Because of such a total asymmetry of information in one-sided information games, one naturally expects that in many circumstances, the uninformed player may not be even willing to participate in the game. It is therefore common that the informed player sends a signal to the uniformed player to help him decide his action (e.g the job applicant sends a college certificate with a high or low level of credibility to the employer, the car dealer announces different warranty plans for his cars). This signal can be considered as the action of the informed player and because "actions speaks louder than words", the uninformed player now has something to base his action on. This is the general structure of *signaling games*.

It should be noted that even after receiving the signal from the informed player, the uninformed player still has many good reasons not to believe in full credibility of such signals. In fact, the whole study of signaling games evolves around the idea of what signals should be sent and how optimally one should react to these signals. It is in this context that *deception* can be formally defined. Deception strategies are those ones which lead an opponent into a disadvantageous position by a deliberate misrepresentation of the truth.

In Section 2, we define a signaling game and an appropriate solution concept called perfect Bayesian equilibrium. The formal presentation of signaling games in this section is mainly adopted from Chapter 8 of [1]. An interested reader can also refer to Chapter 8 of our textbook [2] or Chapter 24 of [3] for some insightful discussions. In Section 3, we introduce an example of signaling games and analyze it. The game introduced in this section is taken from [4]. However, we generalize some of the results given there and study the game as a signaling game whereas in the original paper, only Nash equilibria are found. Finally, in Section 4, we give some concluding remarks and directions for future research.

It should be also pointed out that in this brief introduction, we only focus on single-stage signaling games where players only act once. Although the literature is somewhat scattered, multi-stage, repeated and stochastic signaling and one-sided information games have also been studied. Of special notes are [5] and [6].

2 Definition and Solution Concept

Figure 1 illustrates the structure of a one-stage signaling game. Player 1 has private information about his type θ in Θ and chooses action (signal) s in S. Player 2 observes s and chooses b in B. Before the game begins, it is common knowledge that player 2 has prior beliefs $p(\cdot)$ about player 1's type. After observing s, player 2 updates his beliefs about θ according to Bayes' rule and base his choice of b on the posterior distribution $\mu(\cdot|s)$ over Θ. A strategy for player 1 prescribes a probability distribution $\sigma_1(\cdot|\theta)$ over actions s for each type θ. A strategy for player 2 prescribes a probability distribution $\sigma_2(\cdot|s)$ over actions b for each action s. The expected payoff for player 1 with type θ and strategy $\sigma_1(\cdot|\theta)$ when player 2 plays $\sigma_2(\cdot|s)$ is

$$u_1(\sigma_1, \sigma_2, \theta) = \sum_s \sum_b \sigma_1(s|\theta)\sigma_1(b|s)u_1(s, b, \theta), \tag{1}$$

and the expected payoff for player 2 conditional on s when he uses strategy $\sigma_2(\cdot|s)$ and posterior belief $\mu(\cdot|s)$ can be computed as follows

$$\begin{aligned} u_2(s, \sigma_2, \mu) &= \sum_\theta \mu(\theta|s)u_2(s, \sigma_2(\cdot|s), \theta) \\ &= \sum_\theta \sum_b \mu(\theta|s)\sigma_2(b|s)u_2(s, b, \theta). \end{aligned} \tag{2}$$

Figure 1 and the discussion above show that a signaling game can be modeled as an imperfect information extensive form game with Bayesian inference. It is

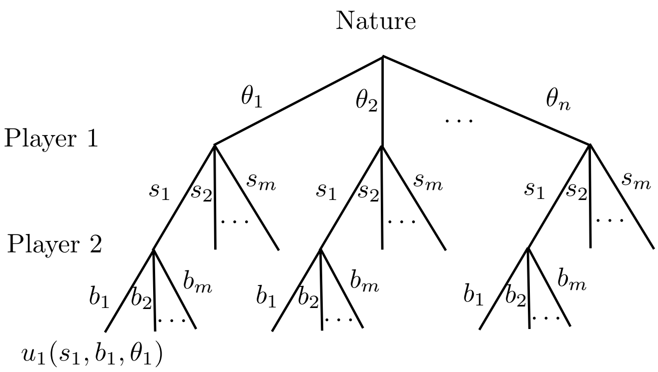

Figure 1: Model of a one-stage signaling game.

therefore natural to expect a solution concept for this game combines the ideas of subgame perfection, Bayes-Nash equilibrium and Bayesian inference.

Definition 1. A perfect Bayesian equilibrium of a signaling game is a strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ and posterior beliefs $\mu(\cdot|s)$ such that

$$\forall \theta,\ \sigma_1^* \in \operatorname*{argmax}_{\sigma_1} u_1(\sigma_1, \sigma_2^*, \theta), \tag{3}$$

$$\forall s, \sigma_2^* \in \operatorname*{argmax}_{\sigma_2} \sum_{\theta} \mu(\theta|s) u_2(s, \sigma_2, \theta) \tag{4}$$

$$\mu(\theta|s) = \frac{p(\theta)\sigma_1^*(s|\theta)}{\sum_{\theta'} p(\theta')\sigma_1^*(s|\theta')} \tag{5}$$

In Definition 1, equation (4) ensures that σ_2^* is the Bayes-Nash equilibrium of the subgames of player 2. Equation (3) ensures that σ^* is the Nash equilibrium of the subgames of each type of player 1. Equation (5) determines how posterior beliefs are obtained. There is one subtlety here and that is we have assumed $\sum_{\theta'} p(\theta')\sigma_1^*(a_1|\theta') > 0$. This assumption holds if player 1 plays actions consistent with his equilibrium. If player 1 plays any action that is not in the support of σ_1^*, it is a common practice to assign an arbitrary posterior distribution for player 2's beliefs of player 1's type.

A signaling game may have different perfect Bayesian equilibria. In a *separating equilibrium*, player 1 sends different signals for each of his types. In other words, he completely reveals his type to player 2. In a *pooling equilibrium*, player 1 sends the same signal for all his types. In this case he does not reveal any new information to player 2 by sending his signals. There can also be *hybrid* or *semi-separating equilibria* in which player 1 randomizes between pooling and separating. It is the structure of the game that determines what types of equilibria exist. In strictly competitive settings, the informed player may want to confuse the uninformed player as much as he can so he may play a pooling strategy. In a less competitive setting, the informed player may want to convey as much information as he can to the uninformed player and therefore he plays a separating strategy.

3 An Example of a Signaling Game

Consider the following game. There are n items and two boxes. Let assume n is an odd number. Player 1 places x number of items in box number 1 and the rest of the items in box number 2. Player 2 picks one box and the items inside it become his. Player 1 takes the items in the other box. The appreciation of players for the number of items they own, $f(i)$, is a strictly increasing function but not necessarily linear. What are the optimal strategies for each player?

Without loss of generality, we can scale $f(i)$ such that $f(0) = 0$ and $f(n) = 1$. Let denote the action of player 1 by θ_i where i is between 0 and n and reflects the number of items, player 1 puts in box 1. Let denote the action of player 2 by b_j where j is 1 or 2 and indicates the box player 2 picks. The matrix of this game can then be shown as

$$\begin{array}{ccccccccc} & \theta_0 & & \theta_1 & & \cdots & & \theta_{n-1} & & \theta_n \\ b_1 & 1 & \leftarrow & f(n-1) & \leftarrow & \cdots & \leftarrow & f(1) & \leftarrow & 0 \\ & \downarrow & & & & & & & & \uparrow \\ b_2 & 0 & \rightarrow & f(1) & \rightarrow & \cdots & \rightarrow & f(n-1) & \rightarrow & 1 \end{array}$$

The matrix cells represent the payoff of player 1. We note that this is a constant-sum game and therefore the payoff of player 1 is 1 minus the payoff player 2 at each entry. It is obvious that the game does not have any Nash equilibrium in pure strategies. In fact, by any counter clock-wise move, one player can increase his payoff as shown in the game matrix. However, since the game is constant-sum, the minmax theorem tells us that the game should have a unique value and any strategy that obtains this value is a Nash equilibrium in mixed strategies. The special structure of the game matrix allows us to quickly find this value. We plot each column expected payoff of player 1 when player 2 randomizes between b_1 and b_2 with probability q assigned to b_1. For each pair of columns i and $n+1-i$, the minimum value of player 1's best response occurs at $q = 1/2$. Therefore the global minimum also occurs at this point and the optimal strategy for player 2 as a minimizer is $(1/2, 1/2)$. Player 1 must select a pair of column actions that maximizes his payoff at $q = 1/2$. Hence he must randomize between columns i^* and $n + 1 - i^*$ for which $f(i^*) + f(n + 1 - i^*) > f(i) + f(n + 1 - i)$ for all i's. This strategy means if, for instance, the items are a collection of stamps and player 1 truly prefers to have all or none of them, then he should only randomize between these two choices.

We now extend this game to a signaling game. In the new game, player 1 still places a certain number of items in box 1 and the rest in box 2. However, he also requires to partially open both boxes and show the content of the boxes to player 2. We assume that the number of items that are revealed to player 2 is always less than half of n. Otherwise the selection is easy for player 2. We also assume if the box is not empty, at least one item is revealed and always unequal numbers of items are revealed. We ask again what the optimal strategies are player 1 and player 2.

To keep it tractable, we assume $n = 3$. In the terminology of Section 2, the types in this game are $\theta_1 = (0, 3)$,$\theta_2 = (1, 2)$, $\theta_3 = (2, 1)$, and $\theta_4 = (3, 0)$, where the first digit indicates the number of items in box 1 and the second digit indicates the number of items in box 2. Signals are $s_1 = (0, 1)$ and $s_2 = (1, 0)$, where the first digit indicates the number of items which are revealed in box 1 and the second digit indicates the number of items that are revealed in box 2. We assume the prior beliefs of player 2 for types of player 1 is (1/4,1/4,1/4,1/4). Finding an equilibrium for this game then involves finding the following probabilities:

$$\begin{aligned}
\sigma_1^*(s_1|\theta_1) = 1 = p_1 &\quad \sigma_1^*(s_2|\theta_1) = 0 = 1 - p_1 \\
\sigma_1^*(s_1|\theta_2) = p_2 &\quad \sigma_1^*(s_2|\theta_2) = 1 - p_2 \\
\sigma_1^*(s_1|\theta_3) = p_3 &\quad \sigma_1^*(s_2|\theta_3) = 1 - p_3 \\
\sigma_1^*(s_1|\theta_4) = 0 = p_4 &\quad \sigma_1^*(s_2|\theta_4) = 1 = 1 - p_4 \\
\sigma_2^*(b_1|s_1) = q_1 &\quad \sigma_2^*(b_2|s_1) = 1 - q_1 \\
\sigma_2^*(b_1|s_2) = q_2 &\quad \sigma_2^*(b_2|s_2) = 1 - q_2.
\end{aligned}$$

We note that when type is θ_1 or θ_4, it is not possible to randomize between signals, therefore there remain 4 parameters (p_2, p_3, q_1, q_2) that must be assigned to have an equilibrium. We next note that the belief update equations are as follows

$$\begin{aligned}
\mu(\theta_1|s_1) = 0 &\quad \mu_1^*(\theta_1|s_2) = \frac{1}{3 - p_2 - p_3} \\
\mu(\theta_2|s_1) = \frac{p_2}{1 + p_2 + p_3} &\quad \mu(\theta_2|s_2) = \frac{1 - p_2}{3 - p_2 - p_3} \\
\mu(\theta_3|s_1) = \frac{p_3}{1 + p_2 + p_3} &\quad \mu(\theta_3|s_2) = \frac{1 - p_3}{3 - p_2 - p_3} \\
\mu(\theta_4|s_1) = \frac{1}{1 + p_2 + p_3} &\quad \mu_1^*(\theta_1|s_2) = 0.
\end{aligned}$$

We next note that the immediate payoff for both players does not rely on the signals sent. In other words, $u_1(s_i, b_1, \theta_k) = f(n + 1 - k)$ and $u_1(s_i, b_2, \theta_k) = f(k - 1)$. Therefore player 1's expected payoff does not directly rely on how he personally randomizes his signals. Two reasonable choices, however, are to pick $p_2 = p_3 = 1$ or $p_2 = p_3 = 0$. These choices give maximum uncertainty on one side of posterior beliefs of player 2. For instance with the first choice, the posterior beliefs of player 2 are $\mu(\cdot|s_1) = (0, 1/3, 1/3, 1/3)$ and $\mu(\cdot|s_2) = (1, 0, 0, 0)$. As for player 2's optimal strategy, he has to maximize his payoff with respect to q_1 and q_2 for each signal and given values of p_2 and p_3. We outline the steps for signal s_1 and $p_2 = p_3 = 1$. Player 2's expected payoff in this case is

$$
\begin{aligned}
u_2(s_1, \sigma_2, \mu) &= \sum_{j=1}^{4} \mu(\theta_j|s_1)[q_1 u_2(s_1, b_1, \theta_j) + (1-q_1) u_2(s_1, b_2, \theta_j) \\
&= \sum_{j=1}^{4} \mu(\theta_j|s_1)[q_1 f(j-1) + (1-q_1) f(n-j+1) \\
&= \sum_{j=1}^{4} \mu(\theta_j|s_1) q_1 (f(j-1) - f(n-j+1)) + \sum_{j=1}^{4} \mu(\theta_j|s_1) f(n-j+1).
\end{aligned} \tag{6}
$$

The last line of equation (6) indicates that maximization with respect to q_1 only concerns with the first sum. If this sum is positive, then $q_1 = 1$, if it is negative, $q_1 = 0$ and if it is zero, q_1 can be chosen arbitrarily. We therefore expand the first sum

$$
\begin{aligned}
\sum_{j=1}^{4} \mu(\theta_j|s_1) q_1 (f(j-1) - f(n-j+1)) = &\frac{p_2}{1+p_2+p_3}[f(1) - f(2)] \\
&+ \frac{p_3}{1+p_2+p_3}[f(2) - f(1)] \\
&+ \frac{1}{1+p_2+p_3}[1-0].
\end{aligned} \tag{7}
$$

And since $p_2 = p_3 = 1$, the above sum is equal to $1/3 > 0$. As a result, q_1 should be one. This means if player 2 sees an item in box 1, he must select that box and otherwise box 2.

4 Conclusion

We provided a brief overview of signaling games. We also investigated the set of strategy solutions for one type of signaling games. In general, because of a cycle, it is harder to find the equilibria of incomplete information extensive form games than the equilibria of complete information extensive form games. One cannot anymore just apply backward induction because the beliefs are updated with the strategies and the strategies are optimal given the beliefs.

The example game provided here can be further investigated in several different directions. The whole set of equilibria can be found and the meaning of each one is explored. The game can be further made complicated by considering other non-uniform priors or assigning a charge to player 2 for partially observing the boxes.

References

[1] D. Fudenberg and J. Tirole, *Game Theory.* MIT Press, 1991.

[2] Y. Shoham and K. Leyton-Brown, *Multi Agent Systems.* Draft Edition, 2005.

[3] P. K. Dutta, *Strategies and Games: Theory and Practice.* MIT Press, 1999.

[4] J. P. Hespanha, Y. Ateskan, and H. Kizilocak, "Deception in non-cooperative games with partial information," in *Proc. of the 2nd DARPA-JFACC Symp. on Advances in Enterprise Control*, 2000.

[5] R. J. Aumann and M. Maschler, *Repeated Games with Incomplete Information.* MIT Press, 1995.

[6] S. Sorin, *A First Course on Zero-Sum Repeated Games.* Springer-Verlag, 2002.