Game Theory: Bayesian Nash Equilibrium

Adaptive and Cooperative Algorithms (ECE 457A)

ECE, MME, and MSCI Departments,
University of Waterloo, ON, Canada

Course Instructor: Benyamin Ghojogh

Fall 2023

Sequential-Move Games

Sequential-Move Games & Common Knowledge

- In **ordered (sequential-move) games**, the players play the game in order.
- Recall that a ordered (sequential-move) game can be represented by an extensive form or a game tree.
- **Common knowledge** refers to the assumption that:
  - each player knows what the game tree looks like.
  - each player knows that other players know what the game tree looks like.
  - each player knows that other players know that other players know what the game tree looks like.
  - and so on.
- Of course, this assumption is useful when all players play the game **rationally**.
- This reminds us of the episode "The One Where Everybody Finds Out" in the "Friends" series where there is the following conversation: "they don't know that we know they know we know!".

Certain, Symmetric, and Complete Games

- Recall that **nature** can be one of the players which makes random movements with known probabilities.
- **Certain game**: nature does not move after any player but it either does not participate in the game or starts the game.
- **Symmetric game**: no player has information different from other players when it moves or at the end nodes of the game tree.
- **Complete game**: nature does not move first or its initial move is observed by every player so that all players know what situation the game is in.

Bayesian Nash Equilibrium

Bayesian Nash Equilibrium

- **Bayesian Nash equilibrium**, proposed in 1967 [1, 2], considers some probabilistic beliefs for every player.
- Therefore, it is useful if there is some **randomness** in the game, such as when **nature** plays a role in the game.
- Moreover, note that Bayesian Nash equilibrium is usually used for **ordered (sequential-move) games**.
- Every player assumes with some probabilistic beliefs that the other players will have some strategy in playing the game.
- Then, while the players play the game one by one during the game, the players update their beliefs using Bayes' rule.
- In the Bayesian Nash equilibrium, we propose an equilibrium and use it to calculate the beliefs; then, we check whether the strategies are the best response for the generated beliefs.
- By Bayes' rule, we see we are probably in which path of the tree.

Bayesian Nash Equilibrium

- In summary, calculating the Bayesian Nash equilibrium has three steps:
  1. We propose a strategy profile. We usually propose strategies which makes sense according to the payoffs. Note that a rational strategy is usually found as the game is a rational game where some patterns will have more payoffs.
  2. We see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  3. We check whether given those beliefs together with the strategies of the other players, each player is choosing a best response for itself.
- The Bayes' rule:

$$\mathbb{P}(A|B) = \frac{\mathbb{P}(B|A)\,\mathbb{P}(A)}{\mathbb{P}(B)} = \frac{\mathbb{P}(B|A)\,\mathbb{P}(A)}{\sum_A \mathbb{P}(B|A)\,\mathbb{P}(A)}, \tag{1}$$

where $\mathbb{P}(A|B)$ and $\mathbb{P}(B|A)$ are called the posterior and likelihood, respectively, and $\mathbb{P}(A)$ and $\mathbb{P}(B)$ are the priors of A and B, respectively.

Bayesian Nash Equilibrium: Example 1

Bayesian Nash Equilibrium: Example 1

- Consider an example game with the following extensive form where the nature N, then player p, and then player q play in order.
- The game has two players p and q and a nature. The actions of player p are p_1 and p_2 while the actions of player q are q_1 and q_2. The actions of nature are A, B, and C with probabilities 0.7, 0.1, and 0.2, respectively.

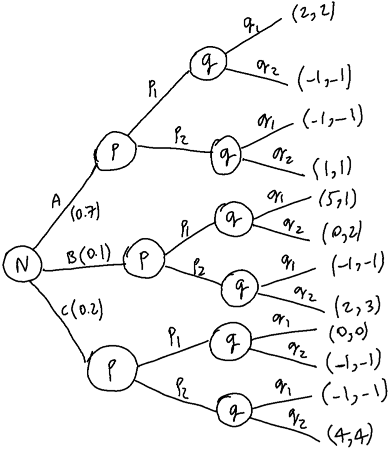

Bayesian Nash Equilibrium: Example 1

- Step 1: we propose a strategy profile.
  - The player p observes the action of nature and then performs its action. The player q does not observe the action of nature but can **guess** it according to the action of the **rational** player p.
  - The prior beliefs of the player q is that nature moves A, B, or C with probabilities 0.7, 0.1, and 0.2, respectively.
  - The player q assumes based on the actions of player p (this assumption makes sense because of the payoffs of the player p at the end nodes of the tree):
    - ★ If the player p chooses action p_1, nature must have chosen A or B probably. So, $\mathbb{P}(p_1|A) = \mathbb{P}(p_1|B) = 1$ and $\mathbb{P}(p_1|C) = 0$.
    - ★ If the player p chooses action p_2, nature must have chosen C probably. So, $\mathbb{P}(p_2|A) = \mathbb{P}(p_2|B) = 0$ and $\mathbb{P}(p_2|C) = 0$.
  - If player p chooses p_1, then player q chooses action q_1. If player p chooses p_2, then player q chooses action q_2. This strategy makes sense because of the payoffs of the player q at the end nodes of the tree.

Bayesian Nash Equilibrium: Example 1

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - Updating the beliefs of player q about the nature if the player p chooses action p_1:

$$\mathbb{P}(A|p_1) = \frac{\mathbb{P}(p_1|A)\,\mathbb{P}(A)}{\mathbb{P}(p_1|A)\,\mathbb{P}(A) + \mathbb{P}(p_1|B)\,\mathbb{P}(B) + \mathbb{P}(p_1|C)\,\mathbb{P}(C)}$$
$$= \frac{1 \times 0.7}{(1 \times 0.7) + (1 \times 0.1) + (0 \times 0.2)} = 0.875,$$
$$\mathbb{P}(B|p_1) = \frac{1 \times 0.1}{(1 \times 0.7) + (1 \times 0.1) + (0 \times 0.2)} = 0.125,$$
$$\mathbb{P}(C|p_1) = \frac{0 \times 0.2}{(1 \times 0.7) + (1 \times 0.1) + (0 \times 0.2)} = 0.$$

  - Updating the beliefs of player q about the nature if the player p chooses action p_2:

$$\mathbb{P}(A|p_2) = \frac{\mathbb{P}(p_2|A)\,\mathbb{P}(A)}{\mathbb{P}(p_2|A)\,\mathbb{P}(A) + \mathbb{P}(p_2|B)\,\mathbb{P}(B) + \mathbb{P}(p_2|C)\,\mathbb{P}(C)}$$
$$= \frac{0 \times 0.7}{(0 \times 0.7) + (0 \times 0.1) + (1 \times 0.2)} = 0,$$
$$\mathbb{P}(B|p_2) = \frac{0 \times 0.1}{(0 \times 0.7) + (0 \times 0.1) + (1 \times 0.2)} = 0,$$
$$\mathbb{P}(C|p_2) = \frac{1 \times 0.2}{(0 \times 0.7) + (0 \times 0.1) + (1 \times 0.2)} = 1.$$

Bayesian Nash Equilibrium: Example 1

- Step 3: we check whether given those beliefs together with the strategies of the other players, each player is choosing a best response for itself.
  - The above calculations show that if player p chooses action p_1, the player q believes that nature has chosen A or B with probabilities 0.875 and 0.125, respectively.
  - If player p chooses action p_2, the player q believes that nature has definitely chosen C with probability 1.
  - As a result, according to the above analysis and the payoffs in the tree, player q chooses action q_1 if player p chooses action p_1. Moreover, player q chooses action q_2 if player p chooses action p_2.
  - Note that if nature has chosen B and player p has actually chosen action p_1, it would have more payoff for player q to choose action q_2; however, as it is more probable that nature must have chosen A given the action p_1 of the player p, it makes sense for the player q to choose action q_1 after the action p_1 of the player p.
  - In summary, the strategy of players in the Bayesian Nash equilibrium is as follows:
    - ★ If player p chooses action p_1 (so the nature must have chosen A with higher probability or B with lower probability), then player q chooses action q_1.
    - ★ If player p chooses action p_2 (so the nature must have chosen C with probability 1), then player q chooses action q_2.

Bayesian Nash Equilibrium: Example 1

- Note that here, the proposed strategy was deterministic. We can also propose a stochastic strategy, also called **mixed strategy** (we will see it later).
- For example, we could have proposed the following strategy: the player p chooses action p_1 with probability 0.5 in state A of nature, with probability 0.4 in state A of nature, and with probability 0.1 in state C of nature. In this case, the conditional probabilities would become:

$$\begin{aligned}
\mathbb{P}(A|p_1) &= \frac{\mathbb{P}(p_1|A)\,\mathbb{P}(A)}{\mathbb{P}(p_1|A)\,\mathbb{P}(A) + \mathbb{P}(p_1|B)\,\mathbb{P}(B) + \mathbb{P}(p_1|C)\,\mathbb{P}(C)} \\
&= \frac{0.5 \times 0.7}{(0.5 \times 0.7) + (0.4 \times 0.1) + (0.1 \times 0.2)} = 0.853, \\
\mathbb{P}(B|p_1) &= \frac{0.4 \times 0.1}{(0.5 \times 0.7) + (0.4 \times 0.1) + (0.1 \times 0.2)} = 0.0975, \\
\mathbb{P}(C|p_1) &= \frac{0.1 \times 0.2}{(0.5 \times 0.7) + (0.4 \times 0.1) + (0.1 \times 0.2)} = 0.0487.
\end{aligned}$$

- Likewise, it is possible to have mixed strategies for the action p_2.

Bayesian Nash Equilibrium: Example 2

Bayesian Nash Equilibrium: Example 2

- Consider an example game with the following extensive form where the nature N, then player p, then player q, and then player r play in order.
- The game has three players p, q, r, and a nature. The actions of player p are p_1 and p_2, the actions of player q are q_1 and q_2, and the actions of player r are r_1 and r_2. The actions of nature are A and B with probabilities 0.5 and 0.5, respectively.

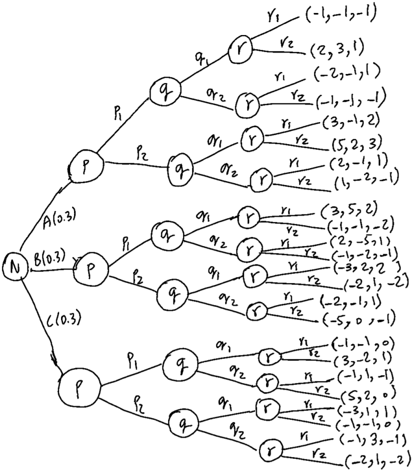

Bayesian Nash Equilibrium: Example 2

- Step 1: we propose a strategy profile.
  - The player p observes the action of nature and then performs its action. The player q does not observe the action of nature but can **guess** it according to the action of the **rational** player p. The player r does not observe the action of nature but can **guess** it according to the action of the **rational** players p and q.
  - The prior beliefs of the player q is that nature moves A, B, or C with probabilities 0.3, 0.3, and 0.3, respectively.
  - The player q assumes based on the actions of player p (this assumption makes sense because of the payoffs of the player p at the end nodes of the tree):
    - ★ If the player p chooses action p_1, nature must have chosen B or C probably. So, $\mathbb{P}(p_1|A) = 0$ and $\mathbb{P}(p_1|B) = \mathbb{P}(p_1|C) = 1$.
    - ★ If the player p chooses action p_2, nature must have chosen A probably. So, $\mathbb{P}(p_2|A) = 1$ and $\mathbb{P}(p_2|B) = \mathbb{P}(p_2|C) = 0$.

Bayesian Nash Equilibrium: Example 2

- Step 1: we propose a strategy profile.
  - The player r assumes based on the action of player q after the action of player p (this assumption makes sense because of the payoffs of the players p and q at the end nodes of the tree):
    - ★ If players p and q do actions p_1 and q_1, respectively, nature must have chosen B probably. So, $\mathbb{P}(q_1|A, p_1) = 0$, $\mathbb{P}(q_1|B, p_1) = 1$, and $\mathbb{P}(q_1|C, p_1) = 0$.
    - ★ If players p and q do actions p_1 and q_2, respectively, nature must have chosen C probably. So, $\mathbb{P}(q_2|A, p_1) = 0$, $\mathbb{P}(q_2|B, p_1) = 0$, and $\mathbb{P}(q_2|C, p_1) = 1$.
    - ★ If players p and q do actions p_2 and q_1, respectively, nature must have chosen A probably. So, $\mathbb{P}(q_1|A, p_2) = 1$, $\mathbb{P}(q_1|B, p_2) = 0$, and $\mathbb{P}(q_1|C, p_2) = 0$.
    - ★ If players p and q do actions p_2 and q_2, respectively, nature must have chosen A probably. So, $\mathbb{P}(q_2|A, p_2) = 1$, $\mathbb{P}(q_2|B, p_2) = 0$, and $\mathbb{P}(q_2|C, p_2) = 0$.
  - If player p chooses p_1, then player q chooses action q_1, then player r chooses action r_1. If player p chooses p_2, then player q chooses action q_1, then player r chooses action r_2. This strategy makes sense because of the payoffs of all players at the end nodes of the tree.

Bayesian Nash Equilibrium: Example 2

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - Updating the beliefs of player q about the nature if the player p chooses action p_1:

$$\mathbb{P}(A|p_1) = \frac{\mathbb{P}(p_1|A)\,\mathbb{P}(A)}{\mathbb{P}(p_1|A)\,\mathbb{P}(A) + \mathbb{P}(p_1|B)\,\mathbb{P}(B) + \mathbb{P}(p_1|C)\,\mathbb{P}(C)} = \frac{0 \times 0.3}{(0 \times 0.3) + (1 \times 0.3) + (1 \times 0.3)} = 0,$$

$$\mathbb{P}(B|p_1) = \mathbb{P}(C|p_1) = \frac{1 \times 0.3}{(0 \times 0.3) + (1 \times 0.3) + (1 \times 0.3)} = 0.5.$$

  - Updating the beliefs of player q about the nature if the player p chooses action p_2:

$$\mathbb{P}(A|p_2) = \frac{\mathbb{P}(p_2|A)\,\mathbb{P}(A)}{\mathbb{P}(p_2|A)\,\mathbb{P}(A) + \mathbb{P}(p_2|B)\,\mathbb{P}(B) + \mathbb{P}(p_2|C)\,\mathbb{P}(C)} = \frac{1 \times 0.3}{(1 \times 0.3) + (0 \times 0.3) + (0 \times 0.3)} = 1,$$

$$\mathbb{P}(B|p_2) = \mathbb{P}(C|p_2) = \frac{0 \times 0.3}{(1 \times 0.3) + (0 \times 0.3) + (0 \times 0.3)} = 0.$$

Bayesian Nash Equilibrium: Example 2

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - According to the chain rule in probability, we have:

$$\mathbb{P}(A, p_1, q_1) = \mathbb{P}(q_1|A, p_1)\mathbb{P}(A, p_1) = \mathbb{P}(q_1|A, p_1)\mathbb{P}(p_1|A)\mathbb{P}(A).$$

  - Updating the beliefs of player r about the nature and player p if the player q chooses action q_1 (the case where player p has chosen action p_1):

$$\begin{aligned}
\mathbb{P}(A, p_1|q_1) = &\Big\{\mathbb{P}(q_1|A, p_1)\mathbb{P}(p_1|A)\mathbb{P}(A)\Big\}\times \\
&\Big\{\mathbb{P}(q_1|A, p_1)\mathbb{P}(p_1|A)\mathbb{P}(A) + \mathbb{P}(q_1|B, p_1)\mathbb{P}(p_1|B)\mathbb{P}(B) \\
&+ \mathbb{P}(q_1|C, p_1)\mathbb{P}(p_1|C)\mathbb{P}(C)\Big\}^{-1} \\
&= \frac{0 \times 0 \times 0.3}{(0 \times 0 \times 0.3) + (1 \times 1 \times 0.3) + (0 \times 1 \times 0.3)} = 0, \\
\mathbb{P}(B, p_1|q_1) = &\frac{1 \times 1 \times 0.3}{(0 \times 0 \times 0.3) + (1 \times 1 \times 0.3) + (0 \times 1 \times 0.3)} = 1, \\
\mathbb{P}(C, p_1|q_1) = &\frac{0 \times 1 \times 0.3}{(0 \times 0 \times 0.3) + (1 \times 1 \times 0.3) + (0 \times 1 \times 0.3)} = 0.
\end{aligned}$$

Bayesian Nash Equilibrium: Example 2

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - Updating the beliefs of player r about the nature and player p if the player q chooses action q_1 (the case where player p has chosen action p_2):

$$\begin{aligned}
\mathbb{P}(A, p_2|q_1) &= \Big\{\mathbb{P}(q_1|A, p_2)\mathbb{P}(p_2|A)\mathbb{P}(A)\Big\}\times \\
&\quad \Big\{\mathbb{P}(q_1|A, p_2)\mathbb{P}(p_2|A)\mathbb{P}(A) + \mathbb{P}(q_1|B, p_2)\mathbb{P}(p_2|B)\mathbb{P}(B) \\
&\quad + \mathbb{P}(q_1|C, p_2)\mathbb{P}(p_2|C)\mathbb{P}(C)\Big\}^{-1} \\
&\quad = \frac{1 \times 1 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 1, \\
\mathbb{P}(B, p_2|q_1) &= \frac{0 \times 0 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 0, \\
\mathbb{P}(C, p_2|q_1) &= \frac{0 \times 0 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 0.
\end{aligned}$$

Bayesian Nash Equilibrium: Example 2

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - Updating the beliefs of player r about the nature and player p if the player q chooses action q_2 (the case where player p has chosen action p_1):

$$
\begin{aligned}
\mathbb{P}(A, p_1|q_2) &= \Big\{\mathbb{P}(q_2|A, p_1)\mathbb{P}(p_1|A)\mathbb{P}(A)\Big\}\times \\
&\Big\{\mathbb{P}(q_2|A, p_1)\mathbb{P}(p_1|A)\mathbb{P}(A) + \mathbb{P}(q_2|B, p_1)\mathbb{P}(p_1|B)\mathbb{P}(B) \\
&+ \mathbb{P}(q_2|C, p_1)\mathbb{P}(p_1|C)\mathbb{P}(C)\Big\}^{-1} \\
&= \frac{0\times 0\times 0.3}{(0\times 0\times 0.3)+(0\times 1\times 0.3)+(1\times 1\times 0.3)} = 0, \\
\mathbb{P}(B, p_1|q_2) &= \frac{0\times 1\times 0.3}{(0\times 0\times 0.3)+(0\times 1\times 0.3)+(1\times 1\times 0.3)} = 0, \\
\mathbb{P}(C, p_1|q_2) &= \frac{1\times 1\times 0.3}{(0\times 0\times 0.3)+(0\times 1\times 0.3)+(1\times 1\times 0.3)} = 1.
\end{aligned}
$$

Bayesian Nash Equilibrium: Example 2

- Step 2: we see what beliefs the strategy profile generates when players update their beliefs in response to each others' moves.
  - Updating the beliefs of player r about the nature and player p if the player q chooses action q_2 (the case where player p has chosen action p_2):

$$
\begin{aligned}
\mathbb{P}(A, p_2|q_2) = &\Big\{\mathbb{P}(q_2|A, p_2)\mathbb{P}(p_2|A)\mathbb{P}(A)\Big\}\times \\
&\Big\{\mathbb{P}(q_2|A, p_2)\mathbb{P}(p_2|A)\mathbb{P}(A) + \mathbb{P}(q_2|B, p_2)\mathbb{P}(p_2|B)\mathbb{P}(B) \\
&+ \mathbb{P}(q_2|C, p_2)\mathbb{P}(p_2|C)\mathbb{P}(C)\Big\}^{-1} \\
&= \frac{1 \times 1 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 1, \\
\mathbb{P}(B, p_1|q_2) = &\frac{0 \times 0 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 0, \\
\mathbb{P}(C, p_1|q_2) = &\frac{0 \times 0 \times 0.3}{(1 \times 1 \times 0.3) + (0 \times 0 \times 0.3) + (0 \times 0 \times 0.3)} = 0.
\end{aligned}
$$

Bayesian Nash Equilibrium: Example 2

- Step 3: we check whether given those beliefs together with the strategies of the other players, each player is choosing a best response for itself.
  - The above calculations show that if player p chooses action p_1, the player q believes that nature has chosen B or C with probabilities 0.5 and 0.5, respectively.
  - If player p chooses action p_2, the player q believes that nature has definitely chosen A with probability 1.
  - If player q chooses action q_1, the player r believes that nature has chosen B and player p has chosen action p_1 definitely with probability 1.
  - If player q chooses action q_1, the player r believes that if nature has chosen A, player p has chosen action p_1 and if nature has chosen B, player p has chosen action p_2.
  - If player q chooses action q_2, the player r believes that if nature has chosen A, player p has chosen action p_2 and if nature has chosen C, player p has chosen action p_1.
  - As a result, according to the above analysis and the payoffs in the tree, the strategy of players in the Bayesian Nash equilibrium is as follows:
    - ★ If player p chooses action p_1 (so the nature must have chosen B or C with equal probability), then player q chooses either action q_1 or action q_2. Subcase (1): if player q chooses action q_1, player r chooses action r_1. Subcase (2): if player q chooses action q_2, player r chooses either action r_1 or action r_2.
    - ★ If player p chooses action p_2 (so the nature must have chosen A with probability 1), then player q chooses action q_1. Then, player r chooses action r_2.

Acknowledgment

- Some slides of this slide deck are inspired by teachings of Prof. Stanko Dimitrov at the University of Waterloo, Department of Management Science and Engineering.
- Some slides of this slide deck are based on the following book: Eric Rasmusen, "Games and Information: An Introduction to Game Theory", 4th Edition, 2007, [3] https://www.rasmusen.org/GI/download.htm

References

[1] J. C. Harsanyi, "Games with incomplete information played by Bayesian players," *Management science*, vol. 14, no. 3, pp. 159–182, 1967.

[2] D. Fudenberg and J. Tirole, "Perfect Bayesian equilibrium and sequential equilibrium," *journal of Economic Theory*, vol. 53, no. 2, pp. 236–260, 1991.

[3] E. Rasmusen, *Games and information: An introduction to game theory*. Wiley-Blackwell, 4 ed., 2007.