Bayesian games with continuous type spaces: The "Study groups" game

Felix Munoz-Garcia

Strategy and Game Theory - Washington State University

Example: Study Groups

- **Tadelis' textbook: section 12.2.2**
- Two students are working together on a project. They can either put in effort ($e_i = 1$) or shirk ($e_i = 0$). If they put in effort, they pay a cost $c < 1$, while shirking has no cost. If either one or both of the students put in the effort than the project is a success, but if both shirk, then it is a failure.
  - We've all been there before.
- Each student varies in how much they care about their success. This is shown by their type, $\theta_i \in [0, 1]$. This type is independently and randomly chosen by nature at the start of the game from a uniform distribution.
  - Recall that a uniform distribution puts equal chance on any of the outcomes between 0 and 1 happening.

Example: Study Groups

- If the project is a success, then each student receives θ_i^2
  - Hence, if the student put in effort, his payoff is $\theta_i^2 - c$. If he shirked, then his payoff is θ_i^2.
- It is common knowledge that the types are distributed independently and uniformly on $[0, 1]$ and that the cost of effort is c.

Example: Study Groups

- This is a Bayesian game with *continuous* type spaces and *discrete* sets of actions.
- Each player needs to determine whether to contribute effort based on their own type, what they believe the type of the other player is, and the cost of contributing effort.
  - We can define this as a strategy $s_i(\theta_i)$ that maps some $\theta_i \in [0, 1]$ onto a corresponding effort $e_i \in \{0, 1\}$. Hence, $s_i(\theta_i)$ will return either a 0 (shirk) or 1 (contribute) depending on what value of θ_i is chosen as player 1's type.
  - Why aren't we mapping θ_j on to this function? Player i cannot observe player j's type.

Example: Study Groups

- Let p be the probability that player j contributes effort to the project. We can then define player i's expected payoff from shirking as

$$\underbrace{p}_{\substack{\text{Player } j \\ \text{contributes}}} \theta_i^2 + \underbrace{(1-p)}_{\substack{\text{Player } j \\ \text{shirks}}} 0 = p\theta_i^2$$

- Therefore, we know that the best response of player i will be to choose effort if his payoff from contributing effort is at least as good as his expected payoff from shirking, or

$$\theta_i^2 - c \geq p\theta_i^2$$

solving for θ_i,

$$\theta_i \geq \sqrt{\frac{c}{1-p}}$$

Example: Study Groups

- From this inequality, notice that the right-hand side is just a constant.
  - This implies that there is some threshold value of θ_i, $\hat{\theta}_i$, for which player 1 will want to contribute effort if $\theta_i \geq \hat{\theta}_i$, while he will not contribute effort if $\theta_i < \hat{\theta}_i$.
  - This is an application of the **threshold rule**.

Example: Study Groups

- This rule is actually quite intuitive:
  - If player i believes that player j will shirk for sure (i.e., $p = 0$), he will only respond contributing if $\theta_i \geq \sqrt{c}$.
  - Since $c < 1$, it is still possible that player i would want to contibute effort and finish the project when his rival shirks.
- However, if player i believes that player j will contribute effort with some positive probability (i.e., $p > 0$), it could cause the value of cutoff $\sqrt{\frac{c}{1-p}}$ to become greater than 1.
  - If that happens, player i would never want to contribute since we know that $\theta_i \in [0, 1]$.
  - Player i would rather free ride at this point (maybe go play some video games).

Example: Study Groups

- So we are now looking for a Bayesian Nash equilibrium in which each student has a threshold type $\hat{\theta}_i \in [0, 1]$ such that

$$s_i(\theta_i) = \begin{cases} 0 & \text{if } \theta_i < \hat{\theta}_i \text{ (shirk)} \\ 1 & \text{if } \theta_i \geq \hat{\theta}_i \text{ (contribute)} \end{cases}$$

- From this observation, we can now derive the best reponse function for player i given some threshold value for $\hat{\theta}_j$.
  - We know that player j will contribute if $\theta_j \geq \hat{\theta}_j$, and from our uniform distribution, we can figure out an exact value for p. $\longrightarrow$

Example: Study Groups

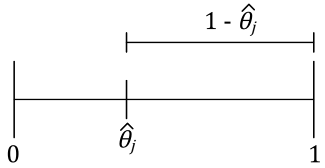

- Putting all of the outcomes from the uniform distribution on a line from 0 to 1, we know that there are $1 - \hat{\theta}_j$ values for θ_j that are above or equal to $\hat{\theta}_j$.
  - This can be interpreted as the probability that $\theta_j \geq \hat{\theta}_j$ (i.e., $p = 1 - \hat{\theta}_j$).

Example: Study Groups

- Substituting back into our inequality from before:

$$\theta_i \geq \sqrt{\frac{c}{1-p}} = \sqrt{\frac{c}{1-(1-\hat{\theta}_j)}} = \sqrt{\frac{c}{\hat{\theta}_j}}$$

- **What if $\hat{\theta}_j > c$?** Then, the right-side of the inequality will be less than 1, i.e., $\sqrt{\frac{c}{\hat{\theta}_j}} < 1$
  - We can then define the cutoff value for player i to contribute as $\hat{\theta}_i = \sqrt{\frac{c}{\hat{\theta}_j}}$.
- **What if $\hat{\theta}_j < c$?** Then, the right-side of the inequality will be greater than 1, i.e., $\sqrt{\frac{c}{\hat{\theta}_j}} > 1$,
  - And since $\hat{\theta}_i$ is upper bounded at 1, we will have $\hat{\theta}_i = 1$.

Example: Study Groups

- Summarizing, player i's best response is

$$BR_i(\hat{\theta}_j) = \begin{cases} \sqrt{\frac{c}{\hat{\theta}_j}} \text{ if } \hat{\theta}_j \geq c \\ 1 \text{ if } \hat{\theta}_j < c \end{cases}$$

Example: Study Groups

- We can depict this BRF of player 1 as follows:

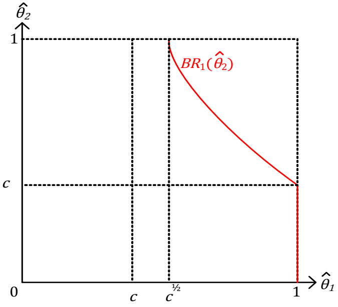

Example: Study Groups

- We can depict this BRF of player 2 as follows:

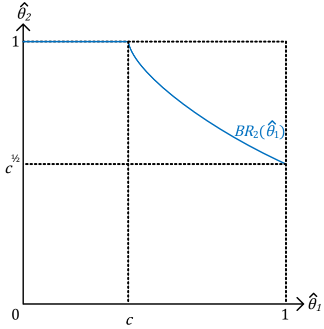

Example: Study Groups

- Implying that the Bayesian Nash Equilibrium (BNE) occurs at the point where both BRFs cross each other.

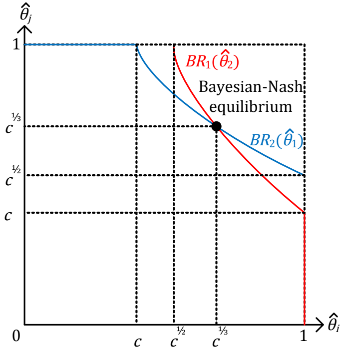

Example: Study Groups

- In order to find the crossing point between both BRFs, we can plug $\hat{\theta}_j = \sqrt{\frac{c}{\hat{\theta}_i}}$ into $\hat{\theta}_i = \sqrt{\frac{c}{\hat{\theta}_j}}$, that is

$$\hat{\theta}_i = \sqrt{\frac{c}{\sqrt{\frac{c}{\hat{\theta}_i}}}} = \frac{c^{1/2}}{\frac{c^{1/4}}{\hat{\theta}_i^{1/4}}} = \frac{c^{1/2}\hat{\theta}_i^{1/4}}{c^{1/4}}$$

- Rearranging,

$$\frac{\hat{\theta}_i}{\hat{\theta}_i^{1/4}} = \frac{c^{1/2}}{c^{1/4}} \implies \hat{\theta}^{3/4} = c^{1/4}$$

and solving for $\hat{\theta}_i$ yields

$$\hat{\theta}_i = \hat{\theta}_j = c^{\frac{1}{3}}$$

Example: Study Groups

- This threshold rule $\hat{\theta}_i = \hat{\theta}_j = c^{\frac{1}{3}}$ is implemented by the following BNE strategy for every player i who, after observing his private type θ_i, chooses the following effort pattern

$$s_i^*(\theta_i) = \begin{cases} 0 \ \ \text{(i.e., shirk)} & \text{if } \theta_i < c^{1/3} \\ 1 \ \ \text{(i.e., effort)} & \text{if } \theta_i \geq c^{1/3} \end{cases}$$

- Thus implying that the student puts effort if and only if his type θ_i is sufficiently high, i.e., $\theta_i \geq c^{1/3}$.