GAME THEORY #6

Reminder

(*) Special case of dynamic games with complete information:
Sequential games: Players move alternatingly, they know everything that happened before (e.g., chess)

(*) Solution concept: Backward induction
(1) Compute best behavior in final stage
(2) Given optimal behavior in final stage, compute best behavior in second-to-final stage.
(3) Repeat until you get to initial stage
The resulting outcome is a Nash equilibrium.

(*) Example: Cooperation and punishment

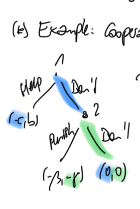

⇝ Solution: $s^{(1)}$ = Don't help
$s^{(2)}$ = Don't punish

Note: If we only required the solution to be a Nash equilibrium, other outcomes would be possible, e.g.

$\left.\begin{array}{l} s^{(1)} = \text{Help} \\ s^{(2)} = \text{Punish} \end{array}\right\}$ NE if $\beta > c$

⇝ Backward induction is an "equilibrium refinement"
It rules out certain Nash equilibria that appear unreasonable, because they require "irrational behavior off the equilibrium path"

§ 3.2 Multi-stage games with observed actions

Example 3.9 (Matching Pennies with outside option)
Consider the following game:

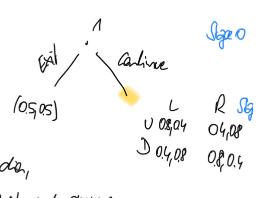

The game has a similar structure as before. However, we can't apply backward induction, because in the final stage there are 2 players who make decisions.

Two questions: (1) How can we formalize games like this?
(2) How to solve them?

Remark 3.10 (Multi-stage game with observed actions)

(*) Informally: (*) Game with T stages $(0,1,\dots,T)$
(*) In stage t, players know what happened in stage $0,\dots,t-1$
(*) In each stage, all players move simultaneously
[However, some players may only have a trivial choice "Do nothing"]

(*) Formally: (-) Players $\mathcal{N} = \{1,\dots,n\}$
(-) Actions may now depend on previous decisions
Action set of player i at stage t takes the form $A^{(i)}(h_t)$ where h_t is the history available at time t.

More specifically:
Stage 0: Available history is $h_0 = \emptyset$
Player's feasible actions are $A^{(i)}(h_0)$
The outcome is $a_0 = (a_0^{(1)},\dots,a_0^{(n)})$

Stage 1: Available history is $h_1 = a_0$
Player's feasible actions are $A^{(i)}(h_1)$
The outcome is $a_1 = (a_1^{(1)},\dots,a_1^{(n)})$

Stage t: Available history is $h_t = (a_0, a_1, \dots, a_{t-1})$
etc.

(-) Payoffs: For a given stage t let H_t be the set of all possible histories up to time t. For a game with T+1 stages, payoffs are a map
$\pi: H_{T+1} \to \mathbb{R}^n$

(-) Behavioral strategies: [A contingent plan how to play in each stage]
Sequence of maps $\{\sigma_t^{(i)}\}_{t=0}^{T}$ with each $\sigma_t^{(i)}: H_t \to \Delta A^{(i)}(h_t)$
[$\Delta A^{(i)}(h_t)$ ... set of all probability distributions over i's available actions after history h_t.]

Examples 3.11

(1) Static games with complete information [Multi-stage game with T = 0]
Players: $\mathcal{N} = \{1,\dots,n\}$
Actions: $A^{(i)}(h_0) = A^{(i)}$
Payoffs: $\pi: \mathcal{A} \to \mathbb{R}^n$
↑ This is H_1

(2) Matching Pennies with a safe outside option

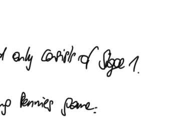

Players: $\mathcal{N} = \{1,2\}$
Actions for $t=0$: $A_0^{(1)} = \{\text{Exit, Continue}\}$
$A_0^{(2)} = \{\text{Do nothing}\}$

Actions for $t=1$
If $h_1 = (\text{Continue, Do nothing})$
$A^{(1)}(h_1) = \{\text{Up, Down}\}$
$A^{(2)}(h_1) = \{\text{Left, Right}\}$

If $h_1 = (\text{Exit, Do nothing})$
$A^{(1)}(h_1) = \{\text{Do nothing}\}$
$A^{(2)}(h_1) = \{\text{Do nothing}\}$

⇝ Everything we have done so far is a subset of multi-stage games with observed actions.

Example 3.12 (Solving the Matching Pennies with a safe outside option)

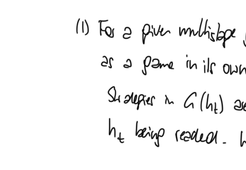

Idea: First solve the simpler game that only consists of Stage 1.
⇝ Nash equilibrium of this Matching Pennies game:
$\hat{\sigma}^{(1)}(\text{Continue}) = (\frac{1}{2}, \frac{1}{2})$
$\hat{\sigma}^{2}(\text{Continue}) = (\frac{1}{2}, \frac{1}{2})$ ⇝ Payoffs (0.6, 0.6)

⇝ Given this, Player 1 can choose backward induction

Equilibrium: Player 1: $\hat{\sigma}_0^{(1)}(\emptyset)$ = Continue
$\hat{\sigma}_1^{(1)}(\text{Continue}) = (\frac{1}{2}, \frac{1}{2})$

Player 2: $\hat{\sigma}_0^{(2)}(\emptyset)$ = Do nothing
$\hat{\sigma}_1^{(2)}(\text{Continue}) = (\frac{1}{2}, \frac{1}{2})$

⇝ Now, let us generalize this idea of looking at simpler parts of the game tree.

Definition 3.13 (Subgame perfection)

(1) For a given multistage game, one can view the game that arises after some given history h_t as a game in its own right. We call it a subgame $G(h_t)$
Strategies in $G(h_t)$ are all strategies of the original game that are compatible with the history h_t being reached. We denote these strategies by $\sigma^{(i)}|_{h_t}$

(2) A strategy profile $\hat{\sigma} = (\hat{\sigma}^{(1)},\dots,\hat{\sigma}^{(n)})$ is a subgame perfect equilibrium (SPE) if for all histories h_t the strategy restrictions form a Nash equilibrium of $G(h_t)$.

Remark 3.14 (On subgame perfection)

(1) Informally, subgame perfection means: Even if the game does not start in the initial node but in some arbitrary node of the game tree, the player's remaining actions still need to be consistent with Nash equilibrium.
(2) Because the whole game is a subgame of itself ⇝ Every SPE is Nash.
(3) A SPE always exists, but it does not need to be unique
↳ Exercises
(4) In sequential games the solution obtained by backward induction is the unique SPE.
(5) Status within game theory: SPE is the standard solution concept for dynamic games with complete information
These games are important ⇝ Nobel prize for Reinhard Selten.

Example 3.15 (Investment decisions under Cournot competition)

Consider two firms that produce an identical good. They need to decide which amount $x^{(i)} \in [0,\infty)$ to produce. They can sell the good at a price $p = 10 - x^{(1)} - x^{(2)}$
The baseline cost to produce one unit of the good is 1 €.
Now firm 1 is considering implementing a new production technology.
This requires an investment of 2 €, but it reduces the per unit cost by 50%

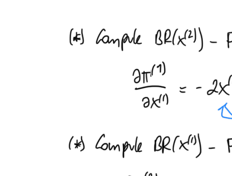

3 Subgames
$G(h_0)$ ... Entire game
$G(h_1 = (\text{Invest}))$ ... Cournot competition after investment
$G(h_1 = (\text{Don't}))$ ... Regular Cournot competition

Regular Cournot competition
We already know (Exercises): $\hat{x}^{(1)} = \hat{x}^{(2)} = 3$ → $\hat{\pi}^{(1)} = \hat{\pi}^{(2)} = 9$

Cournot competition after investment

$\pi^{(1)}(x^{(1)},x^{(2)}) = (10 - x^{(1)} - x^{(2)}) \cdot x^{(1)} - 0.5x^{(1)} - 2 = -{x^{(1)}}^2 + 9.5x^{(1)} - x^{(1)}x^{(2)} - 2$
$\pi^{(2)}(x^{(1)},x^{(2)}) = (10 - x^{(1)} - x^{(2)}) \cdot x^{(2)} - x^{(2)} = -{x^{(2)}}^2 + 9x^{(2)} - x^{(1)}x^{(2)}$

(*) Compute $BR(x^{(2)})$ – For given $x^{(2)}$, what should player 1 do?
$$\frac{\partial \pi^{(1)}}{\partial x^{(1)}} = -2x^{(1)} + 9.5 - x^{(2)} \overset{!}{=} 0 \qquad (a)$$

(*) Compute $BR(x^{(1)})$ – For given $x^{(1)}$, what should player 2 do?
$$\frac{\partial \pi^{(2)}}{\partial x^{(2)}} = -2x^{(2)} + 9 - x^{(1)} \overset{!}{=} 0 \qquad (b)$$
↳ $x^{(1)} = 9 - 2x^{(2)}$

(a)+(b) $\quad -2(9 - 2x^{(2)}) + 9.5 - x^{(2)} = 0$
$\Leftrightarrow \quad 3x^{(2)} = 17/2$
$\Leftrightarrow \quad \hat{x}^{(2)} = 17/6 < 3$
$\hat{x}^{(1)} = 9 - 2 \cdot 17/6 = 10/3 > 3$

⇝ Firm 1 does not only have smaller per unit costs, but it will also get a larger share of the market!

Equilibrium payoffs: $\pi^{(1)}(\hat{x}^{(1)},\hat{x}^{(2)}) = (10 - \frac{17}{6} - \frac{10}{3}) \cdot \frac{10}{3} - 0.5 \cdot \frac{10}{3} - 2 = 82/9 \approx 9.11$
$\pi^{(2)}(\hat{x}^{(1)},\hat{x}^{(2)}) = (10 - \frac{17}{6} - \frac{10}{3}) \cdot \frac{17}{6} - 1 \cdot \frac{17}{6} = \frac{289}{36} \approx 8.03$

Entire game

Comment: In this game, you could have easily arrived at a wrong solution:
"Old" equilibrium $x^{(1)} = 3$
Reduction in per unit-cost 0.5 €
Overall saving 1.5 €
Cost of investment 2 €
} Would have ignored the market consequences

⇝ Subgame perfection allows economists to better analyze non-trivial market decisions

§ 3.3 Repeated games

Remark 3.16 (Setup)

One important special case of a multi-stage game: You play the same game over and over.
Formally: Let $\Gamma = (\mathcal{N}, \mathcal{A}, u)$ be a normal-form game ("stage game"), where $\mathcal{A} = A^{(1)} \times \dots$
Consider the multi-stage game that arises if players face this stage game Γ in each stage $0,\dots,T$
⇓
"round"

(*) Players: $\mathcal{N}$
(*) Actions: $A^{(i)}(h_t) = A^{(i)} \quad \forall h_t$
(*) Payoffs: $\pi: H_{T+1} \to \mathbb{R}^n$
Multiple possibilities. Suppose the history is such that players' actions in round t are $a_t = (a_t^{(1)},\dots,a_t^{(n)})$

Payoff of the finitely repeated game (for T+1 rounds)
$\pi^{(i)} = \frac{1}{T+1} \sum_{t=0}^{T} u^{(i)}(a_t)$ Average payoff per round

Payoff in the infinitely repeated game: ↗ Exercises
(1) "Limit of averages": $\pi^{(i)} = \liminf_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} u^{(i)}(a_t)$
(2) Discounted payoffs: $\pi^{(i)} = (1-\delta) \sum_{t=0}^{T} \delta^t u^{(i)}(a_t)$
Normalization factor. $\qquad \delta < 1$ discount rate
If you get constant payoff $u^{(i)}(a_t) = x$ in every round
⇒ $\pi = x$

Major question: Does the repeated game have equilibria that are systematically different from the equilibria of the stage game?

Examples 3.17 (Finitely repeated games)

1) Repeated prisoner's dilemma

	C	D
C	3,3	0,4
D	4,0	1,1

} Stage game

Stage game has a unique NE (D,D) ⇝ Payoff (1,1)

Now, suppose the stage game is played 2 times.
Can there be a SPE with an average payoff > 1?

No, because in last stage, we always defect independent of previous history. But given play in first round has no effect on behavior in second round, players would like to defect there, too.

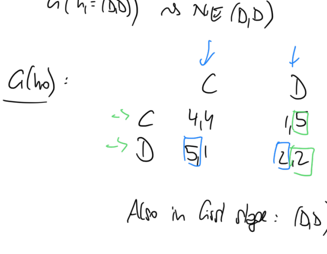

5 possible subgames
$G(h_0)$
$G(h_1 = (C,C))$ ⇝ NE (D,D)
$G(h_1 = (C,D))$ ⇝ NE (D,D)
$G(h_1 = (D,C))$ ⇝ NE (D,D)
$G(h_1 = (D,D))$ ⇝ NE (D,D)

$G(h_0)$:

	C	D
C	4,4	1,5
D	5,1	2,2

} Overall payoffs, taking into account subsequent behavior

Also in first stage: (D,D)

2) Prisoner's dilemma with 2 ways of defecting

	C	D_1	D_2
C	3,3	0,4	-12,0
D_1	4,0	1,1	-10,0
D_2	0,-12	0,-10	-5,-5

If this game is played once
⇝ C is not played in equilibrium
⇝ payoff ≤ 1
If this game is played twice,
payoff > 1 is possible
} Exercise