Game theory I

Stefano Caria

Michaelmas Term 2026

Game theory

- Discrete games and the emergence of cooperation
- Hidden action and sharecropping
- Hidden information and the signalling model of education

(I have included some bonus material on continuous games — e.g. Bertrand competition and Cournot competition – on Canvas).

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

Results from field experiments

> *A game is a description of strategic interaction that includes the constraints on the actions that the players can take and the players' interests, but does not specify the actions that the players do take. A solution is a systematic description of the outcomes that may emerge in a family of games. Game theory suggests reasonable solutions for classes of games and examines their properties.*

Osborne and Rubinstein (1994)

Example: cooperation games

"Cooperation is said to occur when two or more individuals engage in joint actions that result in mutual benefits. Examples include the mutually beneficial exchange of goods, the payment of taxes to finance public goods, team production, common pool resource management, collusion among firms, voting for income redistribution to others, participating in collective actions such as demonstrations, and adhering to socially beneficial norms."

Bowles and Gintis (2006)

Stylized cooperation problems

1. Public good game
2. Prisoner dilemma (a special case of the public good game)
3. Common property resource management
4. Trust game

When individuals have a personal incentive to take actions that decrease aggregate welfare, we have a **social dilemma**

Cooperation problems in developing economies

1. Information about new technologies is a valuable **public good** for firms that make limited R&D investment.

2. Rural communities often derive substantial income from **common property resources** (fishing, forest services, etc..).

3. With weak courts, **trust** becomes essential to be able to trade.

What information do we need to represent a game? The normal form representation

Example: Prisoner Dilemma

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 3
	DEFECT	3 0	1 1

Who are the players?

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 3
	DEFECT	3 0	1 1

What actions can they take, and when?

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 3
	DEFECT	3 0	1 1

Who payoff does each player get for all combinations of player actions??

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 3
	DEFECT	3 0	1 1

For each player, what is the set of best responses to all combinations of other player actions?

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 $\underline{3}$
	DEFECT	$\underline{3}$ 0	$\underline{1}$ $\underline{1}$

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

Results from field experiments

Definition (NASH EQUILIBRIUM)
A Nash equilibrium is an intersection of best-response functions. That is, an outcome is a Nash equilibrium if and only if no player can deviate to another action to improve its payoff, holding fixed all of the other players' actions.

(DEFECT; DEFECT) is the NE of the Prisoner Dilemma

		Player 2	
		COOPERATE	DEFECT
Player 1	COOPERATE	2 2	0 $\underline{3}$
	DEFECT	$\underline{3}$ 0	$\underline{1}$ $\underline{1}$

Definition (STRICT DOMINATION)
For player i, an action a_i strictly dominates an action b_i if player i receives a strictly higher payoff from a_i than from b_i, irrespective of the other players' actions. In that situation, we would say that action b_i is strictly dominated.

When both player have a dominant strategy, we call the resulting equilibrium an equilibrium in dominant strategies.

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

Results from field experiments

Suppose now that the PD is repeated an infinite number of times

		Player 2	
		COLLUDE	COMPETE
Player 1	COLLUDE	2 2	0 3
	COMPETE	3 0	1 1

The game that is repeated is called the 'stage game'.

At the start of each repetition, the outcomes of previous plays of the stage games are common knowledge.

i's payoff from the stage game in period t (given actions a_i and a_j) is $u_t(a_i, a_j)$

Player wants to maximize the discounted sum across stage games

$$U = \sum_{t=0}^{\infty} \delta^t \cdot u_t \tag{1}$$

We will consider possible solutions where players play a trigger strategy:

$$a_i^t = \begin{cases} \texttt{Collude} & \text{if outcome of all stage games so far } \texttt{Collude}, \texttt{Collude} \textit{ or } t = 0 \\ \texttt{Compete} & \text{otherwise} \end{cases}$$

Key features of sums across geometric sequences

$$\sum_{t=m}^{n} \delta^t = \frac{\delta^m - \delta^n}{1 - \delta} \tag{2}$$

$$\sum_{t=m}^{\infty} \delta^t = \frac{\delta^m}{1 - \delta} \tag{3}$$

$$\sum_{t=0}^{\infty} \delta^t = \frac{1}{1 - \delta} \tag{4}$$

Discounted average of a geometric sequence: multiply by $(1 - \delta)$ (a constant!)

Nash equilibrium

Suppose Player 1 plays trigger strategies. Will Player 2 find it optimal to play a trigger strategy too?

We need to check two cases

1. Somebody has played `Compete` at least once before.
2. The outcome of the stage game has alwats been `Collude`, `Collude` or $t = 0$

Nash equilibrium

Suppose Player 1 plays trigger strategies. Will Player 2 find it optimal to play a trigger strategy too?

We need to check two cases

1. Somebody has played `Compete` at least once before.
2. The outcome of the stage game has alwats been `Collude`, `Collude` or $t = 0$

Case where somebody has played `Compete`

- In this case, Player 1 will play `Compete` for the rest of the game.
- Given this, Player 2 will want to play `Compete` as well.

Case 2

- In this case, Player 1 will play `Collude` in the next stage game.
- Now we need to compare the payoff that Player 2 will get if they play `Collude` (consistent with the trigger strategy, which will result in sustained cooperation) or if they deviate to `Compete`

Payoff if collusion is sustained over time

		Player 2	
		COLLUDE	COMPETE
Player 1	COLLUDE	2 2	0 3
	COMPETE	3 0	1 1

Average payoff for player 2:

$$= (1-\delta)\frac{2}{1-\delta}$$
$$= 2$$

Player 2's payoff if 2 plays COMPETE

		Player 2	
		COLLUDE	COMPETE
Player 1	COLLUDE	2 2	0 3
	COMPETE	3 0	1 1

Average payoff for player 2:

$$
\begin{aligned}
&= (1-\delta) \cdot \left(3 + 1 \cdot \delta + 1 \cdot \delta^2 + \ldots\right) \\
&= (1-\delta) \cdot \left(3 + \frac{\delta}{1-\delta}\right) \\
&= 3 \cdot (1-\delta) + \delta \\
&= 3 - 2\delta
\end{aligned}
$$

Conditions for cooperation

$$2 \geq 3 - 2\delta$$

$$2 \geq \frac{3}{1+\delta}$$

which implies

$$\delta \geq 0.5$$

$\rightarrow$ If agents are sufficiently patient, the trigger strategies can be a NE of the repeated game, inducing sustained cooperation.

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

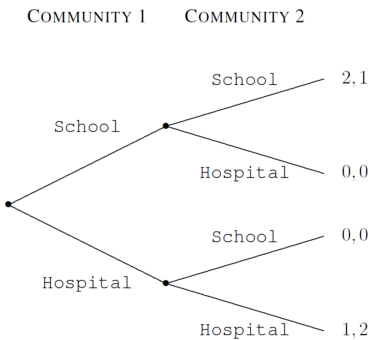
Social preferences and cooperation

Results from field experiments

A coordination problem

		Player 2	
		SCHOOL	HOSPITAL
Player 1	SCHOOL	2 1	0 0
	HOSPITAL	0 0	1 2

Player 1 moves first, player 2 observes their move and then chooses their action: The extensive form representation



Definition (SUBGAME)
For our purposes, a subgame is a game that begins at a node in an extensive form game.

Definition (SUBGAME PERFECT EQUILIBRIUM)
A subgame perfect equilibrium is a combination of strategies that induces a Nash equilibrium in every subgame.

- If community 1 plays `SCHOOL`, community 2 will play `SCHOOL`.
- If community 1 plays `HOSPITAL`, community 2 will play `HOSPITAL`.
- Knowing this, community 1 will play `SCHOOL`.

- The SPE is {`SCHOOL; SCHOOL, HOSPITAL`}

- This method of finding the SPE (starting from the final subgame and going backwards) is called ‘backward induction’

The core idea is that SPE helps us rule out promises and threats that are not credible.

If community 2 would credibly threaten to play `HOSPITAL` even when community 1 plays `SCHOOL`, community 1 would be induced to play `HOSPITAL`.

However, SPE helps us see that this threat would not be credible.

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

- The core model
- Inequality aversion
- Social welfare preferences
- Competitive preferences

Results from field experiments

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

- The core model
- Inequality aversion
- Social welfare preferences
- Competitive preferences

Results from field experiments

- People care deeply about the outcomes of others.

$\rightarrow$ They may be envious of those who have more than them.
$\rightarrow$ They may feel guilty about those who have less than them..

- People also care about the process by which outcome distributions come about.

→ If they feel treated fairly/kindly, they may reciprocate the kindness.

→ If they feel treated unfairly/unkindly, they may reciprocate the unkindness.

The model of Charness and Rabin (2002)

Two players A and B have payoffs π_A and π_B.

$$U_B(\pi_A, \pi_B) = (\alpha a + \delta d + \theta q)\pi_A + (1 - \alpha a - \delta d - \theta q)\pi_B \tag{5}$$

where

- $a = 1$ if $\pi_B > \pi_A$, and $a = 0$ otherwise,
- $d = 1$ if $\pi_B < \pi_A$, and $d = 0$ otherwise,
- $q = -1$ if A has misbehaved and $q = 0$ otherwise.

a captures advantageous inequality, d disadvantegeous inequality.

The model of Charness and Rabin (2002)

$$U_B(\pi_A, \pi_B) = (\alpha a + \delta d + \theta q)\pi_A + (1 - \alpha a - \delta d - \theta q)\pi_B \tag{6}$$

- B has social welfare preferences if $\alpha > \delta > 0$
- B is inequality averse if $\alpha > 0 > \delta$
- B has competitive preferences if $0 > \alpha > \delta$

- B is reciprocal if $\theta > 0$

(Note $1 > \alpha$)

Roadmap

- Basic definitions
- Nash equilibrium
- Repeated games
- Subgame perfect equilibrium
- Social preferences and cooperation
  - The core model
  - Inequality aversion
  - Social welfare preferences
  - Competitive preferences
- Results from field experiments

The model of Charness and Rabin (2002)

$$U_B(\pi_A, \pi_B) = (\alpha a + \delta d + \theta q)\pi_A + (1 - \alpha a - \delta d - \theta q)\pi_B$$

- B is inequality averse if $\alpha > 0 > \delta$

The public good (PG) game

- There are $n > 2$ players
- Each player has an endowment y
- Players simultaneously choose how much to contribute to the public good: $g_i \in [0, y]$
- Contributing to the public good is costly!
- The game can be one-shot, or repeated

The payoff function

$$\pi_i(g_1, ..., g_n) = y - g_i + a\sum_{j=1}^{n} g_j \tag{7}$$

$$\frac{1}{n} < a < 1$$

- You can think of the payoff function as having two parts:
  - A private account: $y - g_i$
  - A public account: $a\sum_{j=1}^{n} g_j$.

- $a < 1$ determines a *private incentive* to keep the endowment.
  - If player i contributes 1 unit to the PG, she lowers the value of her private account by 1 and raises the value of her public account by a;
  - $a < 1$: player i's overall payoff goes down.

The payoff function

$$\pi_i(g_1, ..., g_n) = y - g_i + a\sum_{j=1}^{n} g_j \tag{8}$$

$$\frac{1}{n} < a < 1$$

- $\frac{1}{n} < a$ makes PG contributions *socially beneficial*.
  - Consider a PG contribution of 1 unit.
  - The contributor loses 1 in his private account, while n people get a benefit of a in their public account.
  - $na > 1$ and the sum of all players' payoffs increases.

How do selfish players play the cooperation game?

How will selfish players play the PG game?

One-shot game. As $a < 1$, any strategy that specifies a positive level of g_i is dominated by $g_i = 0$.

Finitely repeated game. In the last subgame any $g_i = k > 0$ is dominated by $g_i = 0$. Strategies involving positive PG provision are not part of a subgame perfect equilibrium.

How do real players play the cooperation game?

Finding 1: positive contributions are very common

Positive contributions are the norm, even for the one-shot game!

Zelmer (2013) reviews 27 different experiments and reports that players on average contribute about 37% of their endowment.

Further, we observe substantial levels of variation in cooperation:

1. across groups;
2. for the *same* individual, in the same group, in *different* contexts.

Equilibrium in the one-shot game with inequality-averse players

- Consider a situation where all players are contributing y.
- Player i knows that she would maximise her own payoff π_i by reducing her contribution to 0.
- However, this would also generate inequality ($\pi_i - \pi_j > 0$) and she would suffer a psychic cost (α) for this inequality.

$\rightarrow$ If α is large enough, the optimal action will be to stick to a contribution of y: cooperation can be sustained!

Finding 2: The unravelling of cooperation

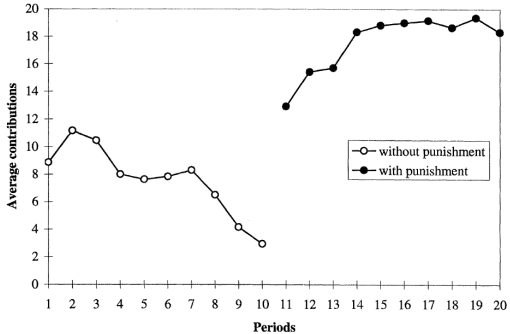


FIGURE 3B. AVERAGE CONTRIBUTIONS OVER TIME IN THE PARTNER-TREATMENT (SESSION 5)

from Fehr and Gachter (2000)

The public good game with punishment

Consider a modified public good game in 2 periods:

- In period 1, individuals choose contributions g_i
- In period 2, individuals observe period 1 contributions and can reduce the payoff other players at a cost $c < 1$.

$p_i = (p_{i1}, ..., p_{in})$ captures the units of punishment given to others. After the two periods, payoffs are given by:

$$\pi_i(g_1, ..., g_n; p_1, ..., p_n) = y - g_i + a\sum_{j=1}^{n} g_j - \sum_{j=1}^{n} p_{ji} - c\sum_{j=1}^{n} p_{ij}$$

Equilibrium when all players are selfish

- This game adds the possibility of costly punishment: you can pay a cost $c < 1$ to decrease the payoff of another player by 1.
- A selfish player would like to contribute 0 and would never punish.
- An inequality-averse player may choose to punish free-riders to improve equality in payoffs.
- To avoid this punishment, selfish players may find it optimal to contribute to the PG.

Higher contributions when punishment is allowed

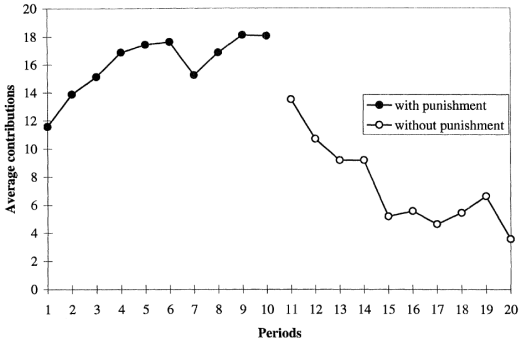
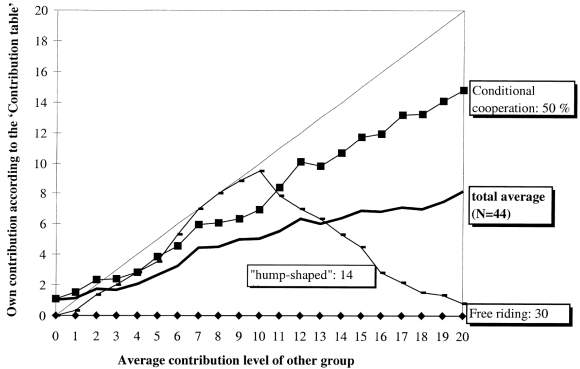


FIGURE 3A. AVERAGE CONTRIBUTIONS OVER TIME IN THE PARTNER-TREATMENT (SESSION 4)

from Fehr and Gachter (2000)

Finding 3: Conditional cooperation



from Fischbacher et al (2001)

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

- The core model
- Inequality aversion
- Social welfare preferences
- Competitive preferences

Results from field experiments

Are people always prepared to hurt others to reduce inequality?

- Consider a ‘dictator game’ with two players — A and B — as in Charness and Rabin (2002):
  - B has to choose between the following allocations:
    1. A: 400; B: 400
    2. A: 750; B: 400
  - 69 percent of player Bs chose allocation 2, which increases disadvantageous inequality!

Do people sacrifice personal welfare for the welfare of others?

- Charness and Rabin (2002) also play the following dictator game:
  - B has to choose between the following allocations:
    1. A: 400; B: 400
    2. A: 750; B: 375
  - $\approx$ 50 percent of player Bs chose allocation 2: they increase disadvantageous inequality at a personal cost!

The model of Charness and Rabin (2002)

$$U_B(\pi_A, \pi_B) = (\alpha a + \delta d + \theta q)\pi_A + (1 - \alpha a - \delta d - \theta q)\pi_B$$

- B has social welfare preferences if $\alpha > \delta > 0$

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

- The core model
- Inequality aversion
- Social welfare preferences
- Competitive preferences

Results from field experiments

- (Some) people are willing to reduce the other player's payoff, at a personal cost, as long as this increases advantageous inequality.

- For example, some Bs would choose the first allocation in a game like this:
  1. A: 100; B: 200
  2. A: 200; B: 250

The model of Charness and Rabin (2002)

$$U_B(\pi_A, \pi_B) = (\alpha a + \delta d + \theta q)\pi_A + (1 - \alpha a - \delta d - \theta q)\pi_B$$

- B has competitive preferences if $0 > \alpha > \delta$

Roadmap

Basic definitions

Nash equilibrium

Repeated games

Subgame perfect equilibrium

Social preferences and cooperation

Results from field experiments

Rustagi et al. (2010) study communally managed forests in Ethiopia

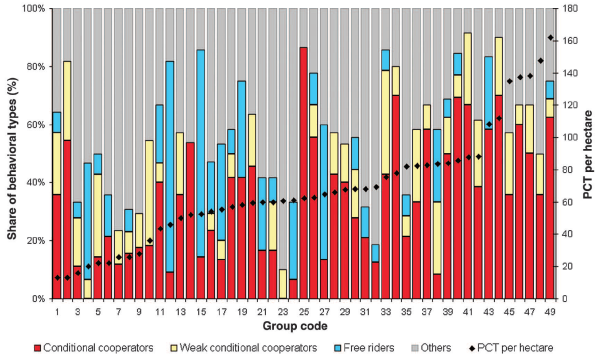


Common property management in Ethiopia

1. They collect information about tree coverage and time spent patrolling the woods (a major form of monitoring)

2. They run public good games and measure the share of ‘conditional cooperators’ in each community

Large variation in share of conditional cooperators across communities



from Rustagi et al (2010)

Share of conditional cooperators predicts forest management outcomes

$$\text{forest outcome}_c = \gamma + \beta * \text{conditional cooperator share}_c + \boldsymbol{\delta X}_c + u_c$$

	Dependent variable: Forest management outcome	
Variables	1	2
Conditional cooperator share		52.085*** (13.117)
Market distance	−22.779*** (4.355)	−20.568*** (4.057)
Time	36.479*** (6.409)	32.185*** (6.004)
Constant	105.325*** (21.719)	81.438*** (19.186)
Additional controls	Yes	Yes
Adjusted r^2	0.58	0.67
Observations	49	49

from Rustagi et al (2010)

Share of conditional cooperators predicts time spent monitoring

$$\text{monitoring}_c = \gamma + \beta * \text{conditional cooperator share}_c + \boldsymbol{\delta X}_c + u_c$$

	Dependent variable: Time spent on monitoring	
Variables	1	2
Conditional cooperator share		23.782*** (8.700)
Market distance	−5.972*** (1.518)	−4.951*** (1.543)
Time	11.428*** (3.814)	9.079** (3.477)
Constant	49.467*** (16.464)	42.092** (15.553)
Additional controls	Yes	Yes
Adjusted r^2	0.31	0.42
Observations	45	45

from Rustagi et al (2010)

How general is that finding that individuals punish anti-social behaviour at personal cost?

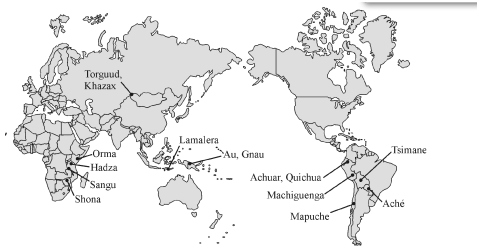


Figure 1. Locations of the 15 small-scale societies.

Henrich et al. (2006) find evidence of widespread punishment in simple games in 15 societies.

Starred Readings

I followed Simon Quinn's presentation of this material. You can check Simon's extensive lecture notes here.

Gibbons, Robert. "A primer in game theory." (1992).
$\rightarrow$ very intuitive exposition of game theory
Charness, G. and Rabin, M., 2002. Understanding social preferences with simple tests. The quarterly journal of economics, 117(3), pp.817-869.

Rustagi, Devesh, Stefanie Engel, and Michael Kosfeld. "Conditional cooperation and costly monitoring explain success in forest commons management." Science (2010): 961-965.

Henrich, Joseph, et al. "Costly punishment across human societies." Science (2006): 1767-1770.