Economics of Networks
Repeated Games, Cooperation, and Network Applications

Evan Sadler

Massachusetts Institute of Technology

Agenda

- Game theory review
- Problem of cooperation
- Finitely repeated Prisoner's Dilemma
- Infinitely repeated Prisoner's Dilemma
- Folk theorems
- Prisoner's Dilemma in a network

Reading: Osborne Chapters 14 and 15

Game Theory Review

Elements of a game:

- Players
- Actions (or Strategies)
- Payoffs

Key solution concept: Nash Equilibrium

- Everyone plays a best reply to others' strategies
- Pure vs. Mixed strategies

Normal vs. Extensive form

- Subgame perfection

Prisoner's Dilemma

How to sustain cooperation?

Recall the Prisoner's dilemma, our workhorse model for this lecture:

	Defect	Cooperate
Defect	$(-3,-3)$	$(0,-4)$
Cooperate	(-4,0)	$(-1,-1)$

Recall (D, D) is the unique Nash Equilibrium

- Defecting is a dominant strategy for both players

Repeated Games

Many situations like this where we observe cooperation

- Why?

One idea: players interact repeatedly over time

- Threat of bad future consequences might induce cooperation now

Study a repeated game

- Play the same stage game over and over
- Can express formally as an extensive form game

Discounting

Key new concept: discounting

A dollar tomorrow is worth less than a dollar today

- Opportunity cost of investment (e.g. interest rates)
- Future consumption less valuable, time preference

The standard approach: exponential discounting

- Discount factor $\delta \in [0, 1)$
- Value of payoff t periods from now multiplied by δ^t

Under interest rate interpretation have $\delta = \frac{1}{1+r}$

- In finance, often use the term "net present value"

A Repeated Game, Formally

Start with a normal form game $G = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$ (Stage game)

- Play the game in each of T discrete periods
- Observe outcome of play in all prior periods
- T finite or infinite

Use notation $\mathbf{s} = \{s^t\}_{t=0}^T$ for sequence of action profiles

- $\boldsymbol{\sigma} = \{\sigma^t\}_{t=0}^T$ for mixed strategies

Payoff to player i

$$U_i(\mathbf{s}) = \sum_{t=0}^{T} \delta^t u_i(s_i^t, s_{-i}^t)$$

Denote T-period repeated game with discount factor δ by $G^T(\delta)$

Finitely Repeated Prisoner's Dilemma

What if we play the Prisoner's Dilemma $T < \infty$ times?

	Defect	Cooperate
Defect	$(-3, -3)$	$(0, -4)$
Cooperate	$(-4, 0)$	$(-1, -1)$

First need to decide on solution concept

- Natural choice: subgame perfect Nash equilibrium

Solve via backward induction

- What happends at time T?

Finitely Repeated Prisoner's Dilemma

Defect is a dominant strategy in the last period, so players play (D, D)

Given this, the subgame at $T - 1$ has a dominant strategy: defect

Iterating this argument, we find the unique SPE is to defect in every period

This is a special case of a more general result...

Equilibria of Finitely-Repeated Games

Theorem

Consider the repeated game $G^T(\delta)$ for $T < \infty$. If the stage game G has a unique pure strategy equilibrium σ^*, then G^T has a unique SPE in which σ^* is played every period.

The proof follows the same logic as in the Prisoners' Dilemma example

By backward induction, at time T the unique outcome is σ^*, and taking this as given, we can iterate to construct the unique SPE

Infinitely Repeated Games

Now consider the infinitely repeated game $G^{\infty}(\delta)$

A pure strategy profile $\mathbf{s}$ is now an infinite sequence of action profiles

Payoff to player i is

$$U_i(\mathbf{s}) = \sum_{t=0}^{\infty} \delta^t u_i(s_i^t, s_{-i}^t)$$

Note summation is well defined since $\delta < 1$

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

Trigger Strategies

Trigger strategies are one way to sustain cooperation in an infinitely repeated game

Idea: we have an agreed upon action profile; if you deviate, I will play a "punishment" action

- Infinite repetition ensures we can always punish

Grim trigger strategy: punishment lasts forever after deviation

- Ability to cooperate depends on worst available punishment

Formally, if $\overline{s}$ is the agreed upon profile and $\underline{s}_i$ is the punishment action, the grim trigger strategy is:

$$s_i^t = \begin{cases} \overline{s}_i & \text{if } s^\tau = \overline{s} \text{ for all } \tau < t \\ \underline{s}_i & \text{if } s^\tau \neq \overline{s} \text{ for some } \tau < t \end{cases}$$

Cooperation in the Repeated Prisoner's Dilemma

Recall

	Defect	Cooperate
Defect	$(-3, -3)$	$(0, -4)$
Cooperate	$(-4, 0)$	$(-1, -1)$

Suppose both players adopt the grim trigger strategy

- Cooperate as long as no one has defected

Can this be a subgame perfect Nash equilibrium?

- Will show it is as long as $\delta > \frac{1}{3}$

Cooperation in the Repeated Prisoner's Dilemma

Step 1: Cooperation is a best response to cooperation

- Suppose at current history there have been no defections

Payoff from cooperation:

$$-\left(1+\delta+\delta^2+...\right)=\frac{-1}{1-\delta}$$

Payoff from defection:

$$0-3\left(\delta+\delta^2+\delta^3+...\right)=\frac{-3\delta}{1-\delta}$$

Cooperation is better if $3\delta > 1$ or $\delta > \frac{1}{3}$

Cooperation in the Repeated Prisoner's Dilemma

Step 2: Defection is a best response to defection

- Suppose someone has defected
- Expect other player to always defect going forward

Defection is unique best response, so grim trigger is subgame perfect

Note: always cooperating is a best response to the grim trigger strategy, but equilibrium requires both players to threaten punishment for defection

- If my opponent always cooperates, I should defect

Multiplicity

Cooperation is an equilibrium, but in general there are many, many subgame perfect equilibria

Another possibility: switching off

- Have one player cooperate and one player defect each period
- Switch roles each period
- If someone deviates from the plan, defect forever (punishment)

In fact, there is a continuum of equilibria

- Very different from case with finite T

Repetition can Support Worse Outcomes

Consider

	A	B	C
A	$(2,2)$	$(2,1)$	$(0,0)$
B	$(1,2)$	$(1,1)$	$(-1,0)$
C	$(0,0)$	$(0,-1)$	$(-1,-1)$

(A, A) is a dominant strategy equilibrium

If $\delta > \frac{1}{2}$, there is a SPE in which (B, B) is played every period

- How can the players support this?

Folk Theorems

In general, can sustain cooperation in essentially any infinitely-repeated game with a sufficiently high discount factor

Results of this type often referred to as "folk theorems"

- Widely believed true before formally proved

Stage game $G = (N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N})$, repeated game $G^\infty(\delta)$

Feasible payoffs

$$V = Conv\,\{\mathbf{v} \in \mathbb{R}^n \,|\, \exists\, \mathbf{s} \in S\; s.t.\, (1-\delta)U(\mathbf{s}) = \mathbf{v}\}$$

Convexity obtained through randomization, normalization by $1 - \delta$

Minimax Payoffs

Minimax payoff of player i: worst payoff opponents can guarantee for i:

$$\underline{v}_i = \min_{s_{-i}} \left\{ \max_{s_i} u_i(s_i, s_{-i}) \right\}$$

Player i can never receive less in any period

Write m^i_{-i} for a profile of others' strategies that forces i to obtain $\underline{v}_i$

Example

	L	R
U	$(-2,-2)$	$(1,-2)$
M	$(1,-1)$	$(-2,2)$
D	$(0,1)$	$(0,1)$

We compute $\underline{v}_1$; write q for probability player 2 plays L

Player 1 earns:

- $1-3q$ from U
- $-2+3q$ from M
- 0 from L

Therefore

$$\underline{v}_1 = \min_{0 \leq q \leq 1} \max\{1-3q, -2+3q, 0\} = 0$$

Folk Theorems

Theorem (Nash Folk Theorem)

If $\mathbf{v}$ is feasible and $v_i > \underline{v}_i$ for all i, then there exists some $\underline{\delta} < 1$ such that for all $\delta > \underline{\delta}$, there is a Nash equilibrium of $G^\infty(\delta)$ with payoffs $(1 - \delta)\mathbf{v}$.

To simplify the argument suppose there is a pure strategy profile $\mathbf{s}$ that delivers the value vector $\mathbf{v}$

Consider the grim trigger strategy for each player i:

- Play s_i as long as no one deviates
- If j deviates, play m_i^j forever

Proof Continued

Can i gain from deviating in period t?

Write $\overline{v}_i$ for i's maximum one period payoff, deviation payoff is bounded by

$$v_i + \delta v_i + ...\delta^{t-1} v_i + \delta^t \overline{v}_i + \delta^{t+1} \underline{v}_i + \delta^{t+2} \underline{v}_i + ...$$

Equilibrium strategy is optimal if

$$\frac{v_i}{1-\delta} \geq \frac{1-\delta^t}{1-\delta} v_i + \delta^t \overline{v}_i + \frac{\delta^{t+1}}{1-\delta} \underline{v}_i$$

which is equivalent to

$$v_i \geq (1-\delta)\overline{v}_i + \delta \underline{v}_i$$

The profile is an equilibrium if $\delta > \underline{\delta} = \max_i \frac{\overline{v}_i - v_i}{\overline{v}_i - \underline{v}_i}$

Some Issues

Can obtain essentially any payoff as a Nash equilibrium with patient players, but punishments can be very costly

- Might not be credible (lack subgame perfection)

	L	R
U	$(6, 6)$	$(0, -100)$
D	$(7, 1)$	$(0, -100)$

Unique NE is (D, L), minimax payoffs are $\underline{v}_1 = 0$ and $\underline{v}_2 = 1$

Can get (U, L) as NE, but punishing player 1 for deviations is very costly

Subgame Perfect Folk Theorem

Theorem

Let σ^* *be a static equilibrium of the stage game with payoffs* $\mathbf{e}$. *For any feasible payoff* $\mathbf{v} > \mathbf{e}$, *there exists* $\underline{\delta} < 1$ *such that for all* $\delta > \underline{\delta}$, *there exists a subgame perfect Nash equilibrium of* $G^\infty(\delta)$ *with payoffs* $\mathbf{v}$

Proof: Same idea as before using σ^* as the grim trigger punishment

Community Enforcement

Intuition: cooperation is easier to sustain if we can enlist others to punish defectors

Example based on Ali and Miller (2016), "Ostracism and Forgiveness"

Suppose we have three players: Ann, Bob, and Carol

- Time is continuous
- Each pair has an interaction at Poisson arrival times with intensity λ

At each interaction, play a version of the work/shirk game

- Simultaneous choose effort levels $a_i \geq 0$
- Effort costs a^2, benefit to other player $a^2 + a$
- Discount future at interest rate r

Community Enforcement

In any given interaction, there is a clear myopic motive to shirk

- Own effort only benefits the other player

Bilateral enforcement

- Players observe the outcome of their own interactions
- Can sustain positive a with a given partner through future threats

Community enforcement

- At each interaction, players can reveal what happened in interactions with others
- If Ann defects on Bob, Bob can tell Carol the next time he sees her
- Then, both Bob and Carol can punish Ann

Bilateral Enforcement

Baseline: each partnership behaves independently

- How much effort and Ann and Bob sustain on their own?

Grim trigger strategy: both exert a as long as other does so, exert 0 forever after a deviation

Incentive constraint:

$$a + a^2 \leq a + \int_0^\infty e^{-rt} \lambda a \mathrm{d}t$$

One time gain from shirking is less than equilibrium payoff

Binding at level $\underline{a} = \frac{\lambda}{r}$

Permanent Ostracism with Mechanical Communication

Suppose players automatically reveal details of all prior interactions to each partner

Modified grim trigger

- All players exert a as long as no one deviates
- If any player deviates, the victim will report to third player
- Victim and third player permanently exert 0 with guilty player
- Victim and third player cooperate at level $\underline{a}$ going forward

Incentive constraint:

$$a + a^2 \leq a + 2\int_0^{\infty} e^{-rt}\lambda a \mathrm{d}t$$

Guilty player cannot conceal deviation, stronger punishment supports higher equilibrium effort $2\underline{a}$

Permanent Ostracism with Strategic Communication

What if individuals choose which interactions to reveal?

Ann considers shirking on Bob, anticipates he will tell Carol

- Ann can still shirk on Carol if she meets Carol before Bob does

Incentive constraint in a permanent ostracism strategy profile:

$$a + a^2 + \int_0^\infty e^{-rt} e^{-2\lambda t} \lambda (a + a^2) \mathrm{d}t \leq a + 2 \int_0^\infty e^{-rt} \lambda a \mathrm{d}t$$

Payoff from shirking on Carol discounted by $e^{-2\lambda t}$, probability that no one else has met Carol by time t

Constraint binds at $\left(\frac{r+4\lambda}{r+3\lambda}\right) \underline{a}$

Strategic Communication Continued

But will Bob tell Carol about Ann's defection?

No! At $a = \left(\frac{r+4\lambda}{r+3\lambda}\right)$, Bob prefers to conceal Ann's guilt

Key insight: telling Carol forces Bob and Carol to revert to cooperation level $\underline{a}$

- Bob loses from partnership with Carol because they can't sustain the same level of cooperation anymore
- Instead, Bob can profit from his private information and defect on Carol himself (no threat of third-party punishment)

Incentive constraint:

$$a + a^2 \leq \underline{a} + \int_0^\infty e^{-rt} \lambda \underline{a} \mathrm{d}t = \underline{a} + \underline{a}^2$$

Bob reports on Ann only if equilibrium effort a is less than $\underline{a}$

A General Result

Suppose there are n players who each interact in pairs, engage in strategic communication

Theorem

In every permanent ostracism equilibrium, each player's expected equilibrium payoff never exceeds that of bilateral enforcement ($\underline{a}$)

Proof: See Ali and Miller (2016)

Temporary Ostracism

Forgiveness facilitates communication and cooperation

- If Bob knows Ann will eventually be forgiven, he looks forward to working with her
- Concealing information from Carol postpones this prospect
- Communication among innocent players may be incentive compatible

Theorem (Ali and Miller)

If $r < 2\lambda(n-3)$, then there exists a temporary ostracism equilibrium that yields payoffs strictly higher than permanent ostracism.