Cooperation and communication dynamics
Session 3

Olivier Gossner

PSE

Roadmap

1. Infinitely repeated games
2. Long-run versus short-run players

Summary on discounted games

Given δ, a set E of vector payoffs is self-generating if, for every $x \in E$, there exists $c\colon A \to E$ such that x is a NE payoff of the game with payoffs

$$(1-\delta)g(a) + \delta c(a)$$

Characterization of E'_δ

E'_δ is the largest bounded self-generating set.

Folk Theorem (Fudenberg Maskin 1988)

Assume that $F \bigcap IR$ is full dimensional, then for every x that is feasible and strictly individually rational, there exists δ_0 such that, for every $\delta > \delta_0$, $x \in E'_\delta$.

It is enough to prove that the set of feasible and individually rational payoffs can be approximated by convex self-generating sets E.

On strategies

The strategies keep track of a "target payoff" in $x \in E$ to be achieved.
By the self-generating property, x is a NE payoff of the game with payoffs:

$$(1-\delta)g(a) + \delta c(a)$$

for some $c \colon A \to E$.

After the play of a, $c(a)$ becomes the new target payoff, and so on...

Limitations of the approach

- Strategies lack strategic appeal
- They depend very much on the details of the game, g, δ.

Why "strictly individually rational"?

	L	R
T	$0, -1$	$1, 1$
B	$-1, -1$	$-1, 0$

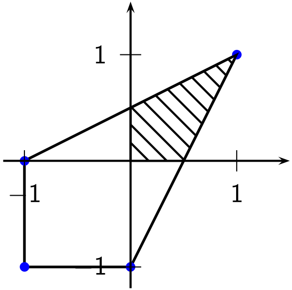

Claim: $(\frac{1}{2}, 0)$ is not a NE payoff of G_δ.

- Only way to generate the payoff is to play (T, L) and (T, R),
- So player 1 must play T in the first stage,
- By playing R forever, player 2 ensures

$$(1-\delta).1 + \delta.0 > 0$$

Construction of SPNE: some difficulties

Easy proof in Nash, difficulties arise because of SPNE

1. Punishing is costly, deviations from punishments must be punished
   Avoid infinite sequences of longer and longer punishments
2. Punishments are in mixed strategies
   Mixed strategies are not observable
3. For strategies to be SPNE, we must check that
   - no deviation is profitable
   - after any history

   There are many such strategies and histories

A tool to check strategies are SPNE

A strategy profile σ is immune to one-shot deviations if, for every h, no player i has a profitable deviation of the form:

- Choose $a^i \neq \sigma^i(h)$ after h
- Follow σ^i afterwards

An SPNE is immune to one-shot deviations

One-shot deviation principle

In a repeated game with continuous payoffs (G_δ, G_n, not G_∞), a strategy profile is a SPNE if and only if it is immune to one-shot deviations.

- Proof in G_n
- Not true for G_∞
- No similar principle for NE

A simplified proof of the FT assuming m.s. are observable

Let x be feasible and strictly individually rational, induced by a cycle of actions $\tilde{a}$. For every i, let x_i feasible and strictly individually rational such that $x^i > x_i^i$, induced by a cycle of actions $\tilde{a}_i$.

- **MP** Play $\tilde{a}$. In case of a deviation of i, go to P(i)
- **MP(i)** Play $\tilde{a}_i$. In case of a deviation of i, go to P(i)
- **P(i)** Play m_i^{-i} for P stages, then return to MP or MP(i). If player j deviates, go to MP(j), otherwise go to MP(i)

We use the OSDP to check that, for P, δ large enough, these strategies form a SPNE.

How do we deal with mixed strategies?

We use a statistical test in order to, at the end of a punishment phase, declare a set of effective punishers. Only effective punishers are rewarded in subsequent play.

A effective punisher is a player j whose action frequency:

- Are close to m_j^i
- Independently of the actions chosen by other players

Properties of the test:

1. Efficiency: If all punishers are effective, the punished player's payoff is at most v_i (up to some ε)
2. Achievability: If a player plays m_j^i repeatedly for P periods, and P is large enough, this punisher is effective with large probability.

Structure of FT strategies

Let $\tilde{a}$ be a cycle of actions with $g(\tilde{a}) = x$. For $J \subset I$, select $\tilde{a}_J$ such that

- $g(\tilde{a}_J) = r_{i+}$ if $i \in J$
- $g(\tilde{a}_J) = r_{i-}$ if $i \notin J$
- $r_{i+} > x_i > r_{i-} > v_i$

MP Play $\tilde{a}$

P(i) Play for P periods. Go to $R(J)$ where J are the effective punishers

R(J) Play $\tilde{a}_J$ for $R \gg P$ periods, then return to MP

Start with MP. If some player i deviates from MP or R(J), start P(i).

Sketch of the proof

Rewards for being a effective punisher are large $\implies$ every punisher passes the review with high probability $\implies$ deviators are effectively punished $\implies$ no incentives to deviate.

Some remark on the FT algorithms

They are incomplete:

- There are histories after which strategies are not defined by the algorithm
- We assume that players play a SPNE of the game played in the punishment phase, and show that all these SPNE have the property that all punishers are effective with large probability

They are robust:

- Independent of δ, provided large enough
- Do not depend on the exact payoff function
- Payoffs could be stochastic, i.e., depend on past actions
- Could relax common knowledge of payoffs

Incompleteness is a necessary condition for robustness.

Roadmap

Motivation

We so far assumed that all players are equally patient.

Many situations such as

- Central bank versus the market
- Cook versus clients (non returning)
- Firm versus customers

are better captured by a patient player facing impatient opponents.

The chain-store game

Consider the following entry game, with $a > 1$, $0 < b < 1$:

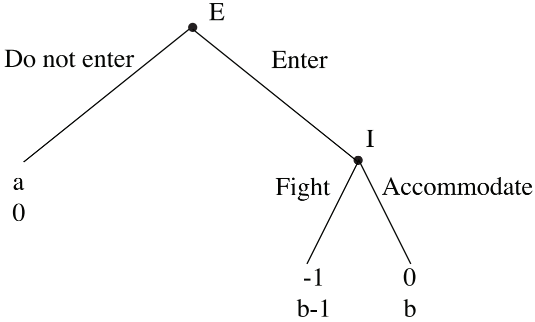

- What are the NE? The SPNE?
- What happens if a long-run incumbent sequentially faces two short-run entrants?
- What happens with a sequence of 100 entrants?

Possibility of a tough incumbent

With (small) probability α, the long-run player is "tough" and a tough player always fights. The game is thus a game of incomplete information.

For $\alpha > b$ no entrant wishes to enter. Now consider $\alpha < b$.

1. There are no SE in which I always accommodates the first entrant.
2. There are no SE in which I always fights the first entrant.
3. Following "Fight", the second entrant is indifferent between E and N, so his belief that the incumbent is tough is b
4. The probability of "Fight" f in the first stage satisfies $\alpha = b(\alpha + f)$
5. The first entrant enters if $\alpha < b^2$, does not enter if $\alpha > b^2$.

Generalization to k entrants, the first entrant does not enter if $\alpha > b^k$.