

GAME THEORY: REPEATED GAMES

Guillaume Pommey

University of Rome Tor Vergata
guillaume.pommey@uniroma2.eu

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Static and dynamic games of complete information

In the first part of the class we have studied:

Static games

- Players, actions, payoffs (normal-form)
- Solution concept: **Nash Equilibrium**

Example: $N = \{1, 2\}$, $A_1 = A_2 = \{\text{Cooperate}, \text{Betray}\}$, payoffs:

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1


Static and dynamic games of complete information

Dynamic games

- Players, actions, payoffs **and** root/decision/terminal nodes, information sets
- Solution concepts: Nash Equilibrium (unsatisfactory), **Subgame Perfect Nash Equilibrium**

Example:

- $N = \{1, 2\}$
- $A_1 = \{L, M, R\}$, $A_2 = \{A, B\}$
- $X_1 = \{x_0\}$, $X_2 = \{x_2, x_3\}$
- $I_1 = \{\{x_0\}\}$, $I_2 = \{\{x_2\}, \{x_3\}\}$
- $r = \{x_0\}$
- $T = \{x_1, x_4, x_5, x_6, x_7\}$



Static and dynamic games of complete information

We used **static games** to represent

- One-shot situations
- Simultaneous choice of actions between players
- Sequential moves but without observability of other players' actions

Examples: Coordination game, prisoner's dilemma, Cournot.

And **dynamic games** to represent

- Sequential moves with observability
- Games that *evolve* depending on the history of actions

Examples: Stackelberg, entry game.

Repeated games

A **repeated game** is some kind of a *combination* of static games and dynamic games

- We take a particular static game...
- and we **repeat it** a given number of times (finite or infinite)
- All past actions are common knowledge among players

The resulting game is obviously a **dynamic game**

- ▷ but **each new period** we play the same static base game

Example: Repeated prisoner's dilemma for $T \geq 2$ periods.

Repeated games

Why is this of any **interest**?

Take the **Prisoner's dilemma problem**:

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1

The **only Nash Equilibrium** of this game is (Betray, Betray)

- Which is clearly a **suboptimal** outcome
- We would obviously prefer (Cooperate, Cooperate)

Repeated games

If the two players were to play this game not only once **but several times**.

Could we achieve a *better* outcome?

With **repeated interactions** we might expect players to use

- **Promises:** If we cooperate today, I will cooperate tomorrow
- **Threats/Revenge:** If you betray me today, I will never cooperate anymore (or for some time)

Repeated games

Repeated interactions have a clear **theoretical interest**

- Allow for some better outcome than the nonrepeated one

But they also have an interest to explain some **social phenomena**

- **Long-term** interactions
- Emergence of **social norms**

Apparently *nonrational* behavior (cooperating in the Prisoner's dilemma) may be the result of **rational** behavior when the game is played **more than once**.

- ▷ **Cooperation** may emerge from a theory of **non-cooperating** players

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Finite and infinite repetitions

A crucial point about repetitions of a game is whether they are:

Finite: The game **ends** after a given number of periods $T < \infty$

Infinite: The game **never ends**, $T = \infty$

We call those games **finitely repeated games** and **infinitely repeated games**, respectively.

Finite and infinite repetitions

Does it make any sense to consider **infinitely** repeated games?

- After all, everything ends at some point (even states)

As often with modeling, we must not take too seriously the **literal meaning** of the model and it is better to interpret it

- Utility theory **does not mean** that people have a utility function in their brain
- Using calculus **does not mean** that all the variables we use are continuous (like prices)
- ...

Finite and infinite repetitions

Here **finiteness** must be understood as the fact that **players perceive the end of the game**.

On the contrary, **infiniteness** represents situations in which players

- are **uncertain whether the game is going to continue** after each period (e.g. employment contract)
- or know the game will end but the **horizon is too far** to be perceived

The choice of modeling must be made accordingly to **what best represents the situation** and not accordingly to the literal meaning of finiteness and infiniteness.

▷ Important as **results differ substantially** with the two approaches

Outline

The theoretical analysis of repeated games is **very dense and rich**

- ▷ We are going to cover **only some** important and interesting topics

Namely,

- Subgame perfect Nash equilibrium in **finitely** and **infinitely** repeated games
- **Folk theorems**
- **Tacit collusion** in Bertrand oligopoly (for Cournot, see the next practice session)

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Finitely repeated games: Definition

We begin with **finitely repeated games**.

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ define a static game of complete information where

- N is the set of players
- $(A_i)_{i \in N}$ is the collection of the sets of players' actions
- $(u_i)_{i \in N}$ are the payoffs

We say that G is the **stage game**.

Definition (nonformal): A finitely repeated game of G is an extensive game with complete information such that players play G at each stage for $T < \infty$ periods. Outcomes and actions of all past periods are observed by all players.

An example: Repeated prisoner's dilemma


Consider the previous **prisoner's dilemma** problem:

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1

And assume it is **played twice**, i.e., $T = 2$

An example: Repeated prisoner's dilemma

First, it is useful to see how this game could be written in **extensive form**



Payoffs

Notice that from our definition of a repeated game as being the succession of plays of the same game we have implicitly assumed that **they receive some payoff at each period**.

- ▷ In previous dynamic games we have conveniently assumed that **payoffs were obtained after the game ends**

We need to define how players perceive this **stream of payoffs**.

- For **finite** repetitions we will simply say that final payoffs are the sum of each period's payoff
- For **infinite** repetitions we will introduce a new tool

Equilibrium concept (finitely repeated games)

The most natural **equilibrium concept** is the **subgame-perfect Nash equilibrium**.

A repeated game is a dynamic game so we want players to be **sequentially rational**.

- ▷ Recall the **failure of Nash equilibrium** in dynamic games (noncredible threats)

In finitely repeated games, we can still rely on **backward induction**

- ▷ not possible anymore in infinitely repeated games

SPNE: Prisoner's dilemma with $T = 2$

Let us consider once again the **prisoner's dilemma** for $T = 2$:

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1

Backward induction:

- Consider the last period of the game, here period 2
- At this point, **whatever happened in period 1 is irrelevant** for period 2's choice of action
- Hence, in period 2, players **choose the Nash Equilibrium actions** of the stage game
- Here it is (Betray, Betray)

SPNE: Prisoner's dilemma with $T = 2$

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1

Backward induction (continued):

- Coming to the first period, both players know that (Betray, Betray) will occur in period 2, **no matter what they do now**
- Hence they **will also play** (Betray, Betray) in the first period

We obtain that the **SPNE** of this twice-repeated prisoner's dilemma is $\{(B, B), (B, B)\}$

- ▷ Notation: (X, Y) is the vector of players' action in a given period and $\underbrace{\{(B, B)\}}_{\text{period 1}}, \underbrace{\{(B, B)\}}_{\text{period 2}}$

SPNE: Prisoner's dilemma with $2 \leq T < \infty$

The resulting SPNE when $T = 2$ is disappointing

- ▷ Still no cooperation

Can we do better with **more** periods $2 \leq T < \infty$?

- ▷ Unfortunately no, the **same backward argument** holds
- Starting at T , players must both choose B , regardless of the history of the game
- In $T - 1$, players know that **they will betray each other** in period T , so the same argument applies: **They must betray each other** now regardless of the history
- In $T - 2, \dots$

SPNE in Finitely repeated games

This sad result unfortunately **holds for a larger class of games**.

Theorem (Osborne and Rubinstein, 1994 - adapted)

*Let $T < \infty$. Consider the T -period repeated game of $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$. If the stage game G has a unique pure-strategy Nash Equilibrium then the T -period repeated game of G has a **unique SPNE**. Furthermore this SPNE is such that players **play the Nash equilibrium at every period**.*

In particular for our **prisoner's dilemma**, (B, B) is the only NE and players play it in every period when the game is **repeated a finite number of periods**.

Cooperation in finitely repeated games

When the stage game has **more than one** Nash Equilibrium, it is possible to find games in which **cooperation occurs as an equilibrium outcome**.

Notably, this analysis give some insights such as

- Players act **cooperatively** when the horizon is *distant enough*
- and **opportunistically** when the horizon is near

However, we **will not discuss** this here as it is more fruitful to investigate cooperative behavior in **infinitely repeated games**.

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Infinitely repeated games

Assume that a stage game G is repeated an **infinite number of periods** $T = \infty$.

The first thing we notice is that we cannot use **backward induction** anymore.

- ▷ There is **no last period**

Hence we cannot simply consider each period separately.

- We will have to define more carefully what a **strategy** means for infinitely repeated games
- But for this we also need to define **how players evaluate payoffs**

Action profiles and discounted utility

Consider the **stage game** $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$.

Let $a_i^t \in A_i$ denote player i 's action at period t and $(a_i^t)_{t=0}^\infty$ denote the **infinite sequence** of player i 's actions.

Similarly, let $a = (a^t)_{t=0}^\infty$ be the infinite sequence of **action profiles** where a^t is the vector of actions of all players at period t .

We assume that player i **evaluates payoffs** according to

$$U_i(a) = \sum_{t=0}^{\infty} \delta^t u_i(a_i^t, a_{-i}^t), \quad (1)$$

where $\delta \in [0, 1)$ is the **discount factor**.

Discounting: Interpretation

The discount factor can be interpreted as the fact that players prefer **sooner than later** utility.

It can be seen as a **behavioral** interpretation.

Or, simply as the result of the fact that if you get 1 euro today **you can invest it** and get $1 * (1 + r)$ euros tomorrow, where r is the interest rate.

- ▷ Hence the **present value** of 1 euro tomorrow is $\frac{1}{1+r} =: \delta$
- ▷ Indeed, I need to invest $\frac{1}{1+r}$ if I want to get $\frac{1}{1+r}(1 + r) = 1$ tomorrow

Discounting: Interpretation

The discount factor can also embed the **probability that the game ends**.

Let $p \in [0, 1]$ be the probability that the game ends in each period.

Assume that you expect to receive π in the next period, then you evaluate it as follows in the current period

$$(1 - p) \frac{1}{1 + r} \pi,$$

that is, you evaluate π in terms of present values **and** you take into account the **probability that the game continues**.

In two periods you evaluate it as $\left(\frac{1-p}{1+r}\right)^2 \pi$

▷ Hence defining $\delta := \frac{1-p}{1+r}$ makes sense (and it belongs to $[0, 1]$)

Discounting: An example

Hence, if you expect to receive the following sequence of payoffs

$$(2, 3, 0, 1, 1, \dots),$$

then your discounted utility will be

$$\delta^0 * 2 + \delta^1 * 3 + \delta^2 * 0 + \sum_{t=3}^{\infty} \delta^t * 1$$

Notice that we need that $\delta < 1$, otherwise the infinite sum **diverges**

Infinitely repeated games: Definition

Let $G = \langle N, (A_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be our stage game and assume that all players share the same discount factor $\delta \in [0, 1)$.

Definition (nonformal): A δ -discounted infinitely repeated game of G is an extensive game with complete information such that players play G at each stage for $T = \infty$ periods. Outcomes and actions of all past periods are observed by all players. Moreover, each player's payoff at the **infinite sequence of action profiles** $a = (a^t)_{t=0}^{\infty}$ is evaluated with the δ -discounted utility formula (1).

Note: *There exists other ways to evaluate the infinite sequence of payoffs. Here we focus on discounted utility, the most commonly used criteria.*

Histories

In order to properly define the equilibrium concept, we need to define first **histories**, **strategies**, and **subgames**.

A **period- t history** $h^t = \{a^0, \dots, a^{t-1}\}$ is the collection of all players' past actions from period 0 to period $t - 1$.

▷ We denote by H^t the set of all period- t histories

Example: Consider the infinitely repeated prisoner's dilemma. At period 2, a possible history might be $h^2 = \{(B, B), (C, B)\}$.

- It indicates that both players have chosen to **betray** in period 0
- Player 1 **cooperated** in period 1
- Player 2 **betrayed** in period 1

Strategies

We can now define what is a **strategy** in an infinitely repeated game.

Definition: In an infinitely repeated game, player i 's **pure strategy** is a infinite sequence of actions **for every possible history of the game**. Formally, each player's pure strategy is $s_i = (s_i^t)_{t=0}^{\infty}$ where $s_i^t : H^t \rightarrow A_i$.

In other words, a strategy describes what action each player will choose **in every period** and for **every possible history** at this period.

- ▷ Strategies are **contingent on the current history** of the game

Subgames

What is a **subgame** in an **infinitely repeated game** of G ?

- For instance, *cut* all periods from 0 to t and keep only the repetitions starting at $t + 1$
- The remaining part goes from $t + 1$ to infinity
- Hence the remaining part is **identical to the original infinitely repeated game**

Subgames

Be careful, the subgame that starts at $t + 1$ is indeed **identical to the original infinite repeated game**

- ▷ in terms of possible actions and payoffs,
- ▷ but it may follow from very different **past histories** of play

Hence, what differentiates one subgame from another **is not about what happens next** but about **what happened before**.

At period $t + 1$, there are **as many subgames as there are possible histories** at $t + 1$.

Subgame Perfect Nash Equilibrium

The definition of **SPNE in infinitely repeated games** is very close to the one you have already seen before.

- ▷ the only subtlety is the definition of **subgames** in such infinitely repeated games

Definition: A Nash equilibrium is subgame-perfect if the players' strategies constitute a Nash equilibrium in every subgame.

It means that at any period, for every possible history, the infinite sequence of strategies of the **remaining game must be a Nash equilibrium**.

Finding SPNE in infinitely repeated games

Contrary to your previous experience, there is **no systematized procedure to find all SPNE** in a given infinitely repeated game.

- ▷ The set of equilibria can be *very large* and difficult/impossible to fully characterize.

Instead of trying to do so, the usual approach consists in **guessing a candidate** for a SPNE and then **proving** it is indeed a SPNE.

- ▷ We will focus on equilibria that exhibit a **cooperative behavior**

Grim trigger strategies

A popular type of equilibrium strategy profile are the so-called **grim trigger strategies**.

- ▷ Also called **non-forgiving** trigger strategies

The idea is the following:

- Players *agree* (implicitly) on a **cooperative action** and on a **punishment action**
- Then, everyone plays their cooperative action as long as **none of the player has ever played something else than their cooperative actions** in all previous periods
- If for a single time, players observe a **past action different from the cooperative action** they **all play their punishment action forever**

Grim trigger strategies

More formally, let $a^C = (a_i^C)_{t=0}^\infty$ and $a^P = (a_i^P)_{t=0}^\infty$ denote the **cooperative actions** and the **punishment actions** sequences, respectively.

Then, the grim trigger strategy of any player i , at any period t , and for any history h^t writes:

$$s_i^t = \begin{cases} a_i^C & \text{if } a^\tau = a^C \text{ for all } \tau < t \\ a_i^P & \text{if } a^\tau \neq a^C \text{ for some } \tau < t. \end{cases}$$

Note: The adjective “grim” or “non-forgiving” refers to the fact that if someone does not cooperate for even a single time, everyone reverts to their punishment action forever. In fact, there exists more sophisticated trigger strategies in which punishments occur only for a finite number of periods before going back to cooperative behavior.

Grim trigger strat. in the prisoner's dilemma (1)

Consider the **infinitely repeated** previous **prisoner's dilemma**:

1\2	Cooperate	Betray
Cooperate	3, 3	0, 5
Betray	5, 0	1, 1

Can we find a **grim-trigger-strategy SPNE**?

A **good candidate** for

- the cooperative action is: C
- the punishment action is: B

Grim trigger strat. in the prisoner's dilemma (2)

Notice that the notions of “cooperative” and “punishment” are solely an **interpretation** of the game's actions.

- ▷ There is **no game theoretic definition** of what is a cooperative/punishment action

Here we choose C as the cooperative action because we would like to see if some equilibrium in which players plays the *best* outcome is possible.

And we choose B as the punishment action because **it is the Nash equilibrium of the stage game.**

- ▷ See below for details

Grim trigger strat. in the prisoner's dilemma (3)

Formally, the **grim trigger strategy** of player i , at period t writes

$$s_i^t = \begin{cases} C & \text{if } a_1^\tau = a_2^\tau = C \text{ for all } \tau < t \\ B & \text{otherwise,} \end{cases}$$

that is, **continue to cooperate** if cooperation occurred in all previous periods, and **betray forever** if for some period someone did not cooperate.

Once again, so far this strategy is **only a postulate**, not an equilibrium. We now have to investigate **under which conditions** our postulate is correct.

Grim trigger strat. in the prisoner's dilemma (4)

To find whether our postulate can be sustained as an equilibrium we have to

- ▷ compute the **payoffs that players can achieve** under the postulated strategy
- ▷ and compare them with the **best possible deviation**

If we find that the players are **always better-off playing the postulated strategy** than **deviating** at any possible period and after any possible past history

- ▷ we will have shown that the postulated strategy is a SPNE

Grim trigger strat. in the prisoner's dilemma (5)

Let us consider any **subgame** starting at period t .

- ▷ Remember that the only thing that distinguishes a subgame from another is the **different past history** between them

Here we can easily divide the set of subgames starting at t in **two categories**:

- ▷ Those for which the **past history** is such that **everyone has always played C in all past histories**
- ▷ and all the others

Grim trigger strat. in the prisoner's dilemma (6)

Let us start with the subgames such that **someone has chosen to betray in the past**.

Hence, at the beginning of the subgame, player i expects the other to choose B .

▷ One-shot deviation principle.

If player i chooses to **cooperate anyway** then player i will get 0 instead of 1 if they choose to betray as well.

▷ In period t or in any other subsequent period.

Then, it is clear that if someones has betrayed in the past, the **best thing to do is to betray forever as well**.

Grim trigger strat. in the prisoner's dilemma (7)

Consider now the set of subgames in which **everyone has always chosen to cooperate in the past**.

In that case, player i expects that the other player will choose to **cooperate forever** (*one-shot deviation principle*).

If player i **sticks to the postulated strategy**, then player i will also cooperate forever and therefore expects to receive:

$$3 + 3\delta + 3\delta^2 + 3\delta^3 + \dots = 3 \sum_{t=0}^{\infty} \delta^t = \frac{3}{1-\delta},$$

evaluated in terms of present value at period t .

Note: Any series $\sum_{t=0}^{\infty} z^t$ with $z \in [0, 1)$ is convergent and its limit is equal to $\frac{1}{1-z}$.

Grim trigger strat. in the prisoner's dilemma (8)

But now look at the **best possible deviation** for player i .

- ▷ Here it is straightforward as there is only two available strategies. Only possible deviation in those subgames is playing B .

If player i **deviates** and plays B instead of C in period t :

$$5 + 1\delta + 1\delta^2 + 1\delta^3 + \dots = 5 + \sum_{t=1}^{\infty} \delta^t = 5 + \frac{\delta}{1 - \delta},$$

as deviating in t yields 5 but then players will revert to B forever.

Grim trigger strat. in the prisoner's dilemma (9)

Hence, it is not obvious whether playing the **grim trigger strategy is worth it**.

When everyone has always cooperated in the past, **deviations are not profitable** when

$$\begin{aligned}\frac{3}{1-\delta} &\geq 5 + \frac{\delta}{1-\delta} \\ \Leftrightarrow \delta &\geq \frac{1}{2}.\end{aligned}$$

What does it mean?

Grim trigger strat. in the prisoner's dilemma (10)

The condition $\delta \geq \frac{1}{2}$ can be interpreted as follows.

A grim trigger strategy **is a SPNE** of this game if

- ▷ players are **patient enough**
- ▷ or that the probability that the game ends is reasonably low

Alternatively, notice that the values of the **cooperative, punishment and deviation payoffs** matter a lot.

- ▷ If the deviation payoff were 100 instead of 5
- ▷ Then we would need $\delta \geq \frac{97}{99} \approx .9797 \dots$

Other SPNE

As previously said, the grim trigger strategy **is not the only SPNE** of this game.

There might exist some **forgiving-trigger-strategy** SPNE in which a deviation is followed by only a **finite number of periods of punishment** before turning back to cooperation.

Also notice that the strategy in which players **choose to betray in all periods** (the Nash equilibrium of the stage game) is also a SPNE.

And potentially **many other equilibria**.

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Folk theorems

There are some (technical) interesting results for infinitely repeated games: **the Folk theorems**.

The basic idea is that, in any infinitely repeated game, there exists a SPNE that can implement **some “interesting” payoff structure as long as players are patient enough**.

To introduce our folk theorem we have to define two new notions, namely, **average payoff** and **feasible payoffs**.

Average payoff

The **average payoff** is simply a **normalization** of the δ -discounted utility:

$$A_i(a) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t u_i(a_i^t, a_{-i}^t) = (1 - \delta) U_i(a),$$

Allow for comparisons with the stage game payoffs directly.

- ▷ In the prisoner's dilemma the payoff from cooperating indefinitely is $\frac{3}{1-\delta}$. Multiplying it by $(1 - \delta)$ gives 3, the stage game payoff from cooperation
- ▷ Does not change player i 's maximization problem (maximizing $U_i(a)$ or any $\alpha U_i(a)$ with $\alpha > 0$ is equivalent

Feasible payoffs

Feasible payoffs corresponds to the set of payoffs that can be obtained by a **convex combination** of the pure-strategy payoffs of the stage game G .

Formally, if we denote by $x^k = (x_1^k, \dots, x_n^k)$ a generic payoff for the stage game G , the set of feasible payoffs is given by

$$F := \{y \in \mathbb{R}^n \mid y = \sum_{k \in K} \alpha_k x^k\},$$


where $\sum_{k \in K} \alpha_k = 1$, $\alpha_k > 0$ for all $k \in K$ and $K = \prod_{i \in N} |A_i|$.

Note: K is the number of possible combinations of actions and each x^k corresponds to the vector of payoffs associated with a particular combination of actions.

Feasible payoffs in the prisoner's dilemma

In practice, with simple games, it is very easy to **find the set of feasible payoffs graphically**.

In the previous prisoner's dilemma, it suffices to represent each four couple of payoffs on the plane and connect the points as follows:



The area in the center corresponds to the set of feasible payoffs.

A folk theorem for infinitely repeated games

Now we can state our **folk theorem**.

Theorem (Friedman, 1971)

Let (e_1, \dots, e_n) denote the payoffs from a Nash equilibrium of G , and let (x_1, \dots, x_n) denote any other **feasible payoffs** from G . If $x_i > e_i$ for every player i and if δ is **sufficiently close to one**, then **there exists a subgame-perfect Nash equilibrium** of the infinitely repeated game of G that achieves (x_1, \dots, x_n) as the **average payoff**.

In other words, provided that players are patient enough, all feasible payoffs that are better than a Nash equilibrium of the stage game can be sustained as a SPNE.

A folk theorem for infinitely repeated games

In the prisoner's dilemma, it corresponds to all the payoffs in the feasible payoffs area that are also in the northeast part of the red quadrant.




Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Application: Collusion in Bertrand

Recall the **Bertrand duopoly** setting:

- ▷ Two firms, say i and j ;
- ▷ Homogeneous good;
- ▷ **Competition by setting prices simultaneously**, p_i and p_j ;
- ▷ Same marginal cost $c \in \mathbb{R}_+$.

Let $Q(p_1 + p_2)$ denote the **market demand**.

- ▷ Decreasing function

Let $\pi^m := \max_p (p - c)Q(p)$ and $p^m = \arg \max_p (p - c)Q(p)$

- ▷ **Monopoly profit and price**

Application: Collusion in Bertrand

We know that if firm **compete once**, the only Nash equilibrium is such that both firms set their **price equal to marginal cost** $p_i = p_j = c$ and **make zero profit**.

If they were to **compete for infinite number of periods**, could they achieve a better outcome (from their point of view)?

We are going to investigate the **existence of a grim trigger strategy** in the infinitely repeated Bertrand model.

Application: Collusion in Bertrand

First, let us find our **candidate actions** for cooperation and punishment.

Cooperation: Let us aim high. The “best” firms could do is to both **set their price equal to the monopoly price** $p_i = p_j = p^m$ and **share the monopoly profit**.

Punishment: The natural punishment action here is to revert to standard Bertrand competition, that is **set the price equal to marginal cost**.

Let us investigate the **grim trigger strategy** in which both firms set the monopoly price **as long as no one as never done anything else** and price equals marginal cost otherwise.

Application: Collusion in Bertrand

Consider any **subgame** starting at period t such that someone has chosen **price equals marginal cost at least once in the past**.

Given that firm j chooses $p_j = c$, then **if firm i follows a grim trigger strategy**, it should set $p_i = c$ as well and expect to get 0 forever.

How could firm i **deviate**?

- ▷ Set $p_i > c$: But then no one buys its product, still zero profit.
- ▷ Set $p_i < c$: Everyone buys its product but makes negative profits.

It is then clear that **sticking to the grim trigger strategy is weakly preferred by firm i** .

Application: Collusion in Bertrand

Now consider the **subgames** starting at t for which everyone has always set their **price equal to the monopoly price for all past periods**.

Given that firm j will choose $p_j = p^m$ at period t (one-shot deviation principle),

firm i can **stick to their grim trigger strategy** and expect the following payoff:

$$\frac{\pi^m}{2} + \delta \frac{\pi^m}{2} + \delta^2 \frac{\pi^m}{2} + \dots = \frac{\pi^m}{2} \sum_{t=0}^{\infty} \delta^t = \frac{\pi^m}{2} \frac{1}{1 - \delta}.$$

Application: Collusion in Bertrand

Firm i can instead choose to **deviate**, i.e., set $p_i \neq p^m$.

What is the **best way to deviate**?

- ▷ Set p_i *very close* to p^m but smaller
- ▷ So that firm i gets all the demand
- ▷ And a profit *arbitrarily close* to π^m

Hence, in that case, firm i expects to receive

$$\pi^m + 0\delta + 0\delta^2 + \dots = \pi^m,$$

as firm i expects that firm j will revert to $p_j = c$ forever after this deviation.

Application: Collusion in Bertrand

Collusion is possible, i.e., sustainable as a SPNE of the infinitely repeated Bertrand game, whenever

$$\frac{\pi^m}{2} \frac{1}{1 - \delta} \geq \pi^m,$$

that is, whenever $\delta \geq \frac{1}{2}$.

Table of Contents

1. Introduction
2. Finitely v.s. infinitely repeated games
3. Finitely repeated games
4. Infinitely repeated games
5. Folk theorems in infinitely repeated games
6. Collusion in Bertrand
7. Supplements on trigger strategies

Supplements on trigger strategies

Let us investigate some other things about **trigger strategies**.

Consider the following prisoner's dilemma.

1\2	Cooperate	Betray
Cooperate	x, x	$0, 3$
Betray	$3, 0$	$1, 1$

where $x \in \mathbb{R}_+$.

We investigate the relationship between the **discount factor** and the **value of cooperation** when players use grim trigger strategies.

Supplements on trigger strategies

First, no matter the value of x , we know that in all subgames such that someone has played C in the past **it is a best-response for both players to stick to the grim trigger strategy** and betray forever.

Hence, let us consider any subgame such that **all past histories contain only the cooperative action**.

If player i **sticks to the grim trigger strategy** they expect to receive:

Supplements on trigger strategies

First, no matter the value of x , we know that in all subgames such that someone has played C in the past **it is a best-response for both players to stick to the grim trigger strategy** and betray forever.

Hence, let us consider any subgame such that **all past histories contain only the cooperative action**.

If **player i sticks to the grim trigger strategy** they expect to receive:

$$x\delta^0 + x\delta^1 + x\delta^2 + \dots = x \sum_{t=0}^{\infty} \delta^t = \frac{x}{1-\delta}$$

Supplements on trigger strategies

If otherwise, player i chooses to **betray now** (and hence forever), they get:

$$3\delta^0 + 1\delta^1 + 1\delta^2 + \dots = 3 + \sum_{t=1}^{\infty} \delta^t = 3 + \frac{\delta}{1 - \delta}$$

Therefore, the grim trigger strategy **can be sustained as a SPNE** if and only if

$$\begin{aligned} \frac{x}{1 - \delta} &\geq 3 + \frac{\delta}{1 - \delta} \\ \Leftrightarrow 2\delta + x - 3 &\geq 0. \end{aligned}$$

Supplements on trigger strategies


The condition $2\delta + x - 3 \geq 0$ expresses the idea that x and δ

- ▷ must be **large enough**
- ▷ can be somewhat **substituted one to another**

Here a low value of δ can be **compensated** by a large value of x (and reciprocally)


Supplements on trigger strategies

We can represent this condition **graphically**.



Supplements on trigger strategies

The green region represents all the pairs (x, δ) such that the grim trigger strategy is a SPNE of the infinitely repeated game.



Supplements on trigger strategies: limited punishment

Let us now investigate the **limited punishment case**.

Let $x = 2$ so that we investigate the **infinitely repeated version** of the following game:

1\2	Cooperate	Betray
Cooperate	2, 2	0, 3
Betray	3, 0	1, 1

Notice that when $x = 2$, if we compute the discount factor such that the grim trigger strategy is a SPNE we get $\delta \geq 1/2$.

Supplements on trigger strategies: limited punishment

Consider the following **forgiving trigger strategy**:

- ▷ When a deviation is observed in t , play the punishment action for k periods starting at $t + 1$;
- ▷ If no deviation is observed in t , play the cooperative action in $t + 1$.

Let us see under which conditions this can be a SPNE.

Supplements on trigger strategies: limited punishment

Assume that the **cooperative action has been played in the previous period**.

Player i therefore expects player j to continue to play C .

If player i **decides to deviate**, we can compute their payoff for the next $k + 1$ periods

$$3 + \left[\delta^1 + \delta^2 + \dots + \delta^k \right] = 3 + \frac{\delta(1 - \delta^k)}{1 - \delta},$$

as for k periods, players j will punish player i by playing B and it is a best-response to i to also play B for k periods.

Note: Any series $\sum_{t=0}^k z^t$ with $z \in [0, 1)$ is equal to $\frac{1-z^{k+1}}{1-z}$.

Supplements on trigger strategies: limited punishment

Now assume that player i **sticks to the trigger strategy**.

Player i 's payoff for the next $k + 1$ periods is:

$$2\delta^0 + 2\delta^1 + \dots + 2\delta^k = \frac{2(1 - \delta^{k+1})}{1 - \delta}.$$

We now simply have to compare the payoffs from deviation and from the grim strategy for the next $k + 1$ periods.

Supplements on trigger strategies: limited punishment

The **forgiving trigger strategy** with k periods of punishment is a SPNE if and only if

$$\frac{2(1 - \delta^{k+1})}{1 - \delta} \geq 3 + \frac{\delta(1 - \delta^k)}{1 - \delta}.$$

For $k = 2$, it is equivalent to $\delta \geq 0.62$.

For $k = 3$, to $\delta \geq 0.55$.

As k increases δ goes to $1/2$.