

17.810/17.811 – Game Theory

Lecture 5: Repeated Games

Asya Magazinnik

MIT

Where We Are/Where We Are Headed

- We have developed a notion of dynamic games of complete information in which players make multiple, sequential moves
- We will now consider a special form of such games: **repeated games**, in which players repeat the same game structure again and again
- We will study **finitely** and **infinitely** repeated games

These slides will focus on the following readings:

- Finitely Repeated Games
 - Gibbons, 2.3A
- Infinitely Repeated Games
 - Gibbons, 2.3B

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Repeated Games

- The most interesting conceptual issue in repeated games is the extent to which *repetition* creates the opportunity to sustain more behavior (as Nash equilibria) than is possible in single-shot games.
- The set of Nash equilibria is much **larger** in repeated games than the corresponding static versions.
 - Repeated games have a different problem: the proliferation of equilibria is so great that generating precise predictions becomes difficult.

Some Details

Definition (Stage Game)

Let $G = \{A_1, \dots, A_n; u_1, \dots, u_n\}$ denote a static game of complete information in which players 1 through n simultaneously choose actions a_1 through a_n from the action spaces A_1 through A_n , respectively, and payoffs are $u_1(a_1, \dots, a_n)$ through $u_n(a_1, \dots, a_n)$.

The game G will be called the **stage game** of a repeated game.

Given a stage game G , let $G(T)$ denote the **finitely repeated game** in which G is played T times, with the outcomes of all preceding plays observed before the next play begins. The payoffs for $G(T)$ are simply the sum of the payoffs from the T stage games.

Some Details

As before:

- A **history** is a sequence of play defining a path through the game tree, which is also a record of prior actions and stage game outcomes for all previous interactions.
- A **strategy** is a complete, contingent plan that tells the player what to do in every situation, that is, at every possible history.
- The **equilibrium path** is the sequence of outcomes determined in each stage game that results from the interaction of the players' equilibrium strategies at each moment in time.

An Example: Subgame Perfect Nash Equilibria

1 / 2	<i>L</i>	<i>M</i>	<i>R</i>
<i>T</i>	8, 8	0, 0	1, 9*
<i>M</i>	0, 0	5*, 5*	0, 0
<i>B</i>	9*, 1	0, 0	3*, 3*

- If this game is played once there are two Nash equilibria: (M, M) and (B, R)
- Although the strategy profile (T, L) provides the highest aggregate payoff, it is not a Nash equilibrium; Player 1 unilaterally defects to B and Player 2 unilaterally defects to R .
- What happens if this game is played twice with players caring about their combined two-period payoffs? Can players ever get the (T, L) payoff?

Subgame Perfect Nash Equilibria

1 / 2	<i>L</i>	<i>M</i>	<i>R</i>
<i>T</i>	8, 8	0, 0	1, 9*
<i>M</i>	0, 0	5*, 5*	0, 0
<i>B</i>	9*, 1	0, 0	3*, 3*

- Consider the following strategies:
 - Player 1: Play *T* in period 1; if Player 2 plays *L* in period 1 play *M* in period 2; otherwise play *B* in period 2.
 - Player 2: Play *L* in period 1; if Player 1 plays *T* in period 1 play *M* in period 2; otherwise play *R* in period 2.
- Both players' equilibrium payoff is $8 + 5 = 13$. (Check that there is no deviation that leaves either player better off.) In fact, these strategies constitute a **subgame perfect Nash equilibrium**.
 - Because (M, M) and (B, R) are Nash equilibria of the one-shot game, playing them in the proper subgames is consistent with subgame perfection.

The Repeated Prisoner's Dilemma

- One of the most studied games is the repeated prisoner's dilemma.
- Consider an application focused on trade policy.
- Suppose the world economy performs better when all nations agree to free trade, but that individual countries prefer to protect their domestic economy.
- Given this tension, how are free trade regimes sustained?
- One answer is that free trade can be supported as an equilibrium in a repeated game where a trade war begins whenever a major country defects from the trade agreement.

Free Trade

Free Trade Game		
US / EU	<i>Free Trade</i>	<i>Protect</i>
<i>Free Trade</i>	10, 10	1, 12*
<i>Protect</i>	12*, 1	4*, 4*

- Obviously, if the game is played once, the unique Nash equilibrium is the strategy profile (*Protect*, *Protect*).

Free Trade

If it is played twice, then the strategy sets for each player are:

$$\{FFF,FFP,FPT,FPP,PFF,PFP,PPF,PPP\}$$

where *FFP* means “play **Free Trade** in period 1 and play **Free Trade** in period 2 if the other country plays **Free Trade** in period 1, otherwise play **Protect**.”

- Note that a **complete, contingent plan** conditions strategies on prior histories.

Normal Form

Two-Period Free Trade Game								
US / EU	FFF	FFP	FPF	FPP	PFF	PFP	PPF	PPP
FFF	20,20	20,20	11,22	11,22	11,22	11,22	2,24	2,24
FFP	20,20	20,20	11,22	11,22	13,13	13,13	5,16	5,16
FPF	22,11	22,11	14,14	14,14	11,22	11,22	2,24	2,24
FPP	22,11	22,11	14,14	14,14	13,13	13,13	5,16	5,16
PFF	22,11	13,13	22,11	13,13	14,14	5,16	14,14	5,16
PFP	22,11	13,13	22,11	13,13	16,5	8,8	16,5	8,8
PPF	24,2	16,5	24,2	16,5	14,14	5,16	14,14	5,16
PPP	24,2	16,5	24,2	16,5	16,5	8,8	16,5	8,8

Nash Equilibria

Unlike the first example, repeating the game once does not achieve cooperation, as (PPP, PPP) is the only Nash equilibrium.

This result can be generalized to any finite number of periods:

- In the last period, each country protects.
- This is known in the penultimate period. Thus, each country has an incentive to protect in this period as well.
- This process unravels until each country is protecting in every period.

Why could we induce first-period cooperation in first example?

- Because first-period behavior helps coordinate between multiple equilibria in the second period.

Nash Equilibria

In the first example, the good equilibrium is used as a reward whereas the bad equilibrium is used as a punishment.

Because the Prisoner's Dilemma has only one Nash equilibrium, it is impossible to encourage cooperation with the promise of coordinating on a good equilibrium or the threat of coordinating on a bad equilibrium.

We can generalize this result:

Proposition

If the stage game G has a unique Nash equilibrium then, for any finite T , the repeated game $G(T)$ has a unique subgame-perfect outcome: the Nash equilibrium of G is played in every stage.

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Payoffs in Repeated Games: Discounting

If we **discount** our payoffs, that means payoffs received today are more valuable than payoffs received in the future.



© Source unknown. All rights reserved. This content is excluded from our Creative Commons license. For more information, see <https://ocw.mit.edu/help/faq-fair-use/>.

Why?

- Impatience, inflation, death, game form may change, preferences may change, game may end...

Simple Payoffs

So if the payoffs to the stage game on the equilibrium path are $\pi_1, \pi_2, \pi_3 \dots$, then the **present value** of this infinite series of payoffs is:

$$u_i(s_i, s_{-i}) = \pi_1 + \delta\pi_2 + \delta^2\pi_3 + \dots = \sum_{t=1}^{\infty} \delta^{t-1}\pi_t$$

Since the discount factor satisfies $0 \leq \delta < 1$, this is a **convergent geometric series** (see next slide).

We will show that:

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

A Quick Convergence Proof

$$y = \pi + \delta\pi + \delta^2\pi \dots$$

$$y = \pi + \delta(\pi + \delta\pi + \delta^2\pi \dots)$$

$$y = \pi + \delta y$$

$$y - \delta y = \pi$$

$$y(1 - \delta) = \pi$$

$$y = \frac{\pi}{1 - \delta}$$

Some Other Useful Information

- The **continuation value** is the payoff stream starting from some time τ onward and is given by:

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

- Given the discount factor δ , the **average payoff** of the infinite sequence of payoffs $\pi_1, \pi_2, \pi_3, \dots$ is:

$$(1 - \delta) \sum_{t=1}^{\infty} \delta^{t-1} \pi_t$$

Since the average payoff is just a rescaling of the present value, maximizing the two quantities is equivalent.

Other Preliminaries

We are now ready to restate our familiar notions of **strategies**, **Nash equilibrium**, **subgames**, and **subgame perfection** in the context of infinitely repeated games.

But first, let's define an infinitely repeated game.

Definition

Given a stage game G , let $G(\infty, \delta)$ denote the **infinitely repeated game** in which G is repeated forever and the players share the discount factor δ . For each t , the outcomes of the $t - 1$ preceding plays of the stage game are observed before the t^{th} stage begins. Each player's payoff in $G(\infty, \delta)$ is the present value of the player's payoffs from the infinite sequence of stage games.

Strategies in Infinitely Repeated Games

Strategies are exactly the same as in finitely repeated games:

Definition

In the infinitely repeated game $G(\infty, \delta)$, a player's **strategy** specifies the action the player will take in each stage, for each possible history of play through the previous stage.

This is an infinite number of strategies for us to consider!

We won't have to enumerate them all. Rather, we will consider **strategy types** in our analysis.

Nash Equilibrium in Infinitely Repeated Games

Our general notion of **Nash equilibrium** also remains the same: given the other players are playing their best response in every period, player i has no incentive to unilaterally deviate from their strategy in any period.

Subgames and Subgame Perfection

Definition

In the infinitely repeated game $G(\infty, \delta)$, each **subgame** beginning at stage $t + 1$ is identical to the original game $G(\infty, \delta)$. As in the finite-horizon case, there are as many subgames beginning at stage $t + 1$ of $G(\infty, \delta)$ as there are possible histories of play through stage t .

Definition

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

The One-Shot Deviation Principle

But how do we check every subgame in an infinitely repeated game? Here we are helped by the **one-shot deviation principle**.

It turns out that to find an SPNE, it suffices to compare playing your equilibrium strategy to any **one-shot deviations** of the form:

- Playing your equilibrium strategy up to period $t - 1$
- Deviating to something else in period t
- Returning to your equilibrium strategy in period $t + 1$

This allows us to consider the finite **types** of subgames where we might end up, and one-shot deviations at some arbitrary period t .

The One-Shot Deviation Principle

Definition (One-shot deviation principle)

A strategy profile of an extensive-form game is a subgame-perfect equilibrium (SPE) if and only if there is no profitable one-shot deviation for any player in any subgame.

In an infinite horizon game where the discount factor is less than 1, a strategy profile is a subgame perfect equilibrium if and only if it satisfies the one-shot deviation principle.

Note: More broadly, Nash equilibria have no profitable one-shot deviations on their equilibrium paths, but may have profitable one-shot deviations off their equilibrium paths.

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Back to PD: The Grim Trigger

- Consider the following strategy: “Play free trade in every period until the other country protects. If the other country protects, then protect forever after.”
- This is known as the **grim trigger strategy**, because any failure to cooperate leads to the non-cooperative equilibrium in all future periods.
- Does (grim trigger, grim trigger) constitute a **Nash equilibrium**? Does it constitute a **subgame perfect Nash equilibrium**?
 - The answer to both questions is yes, under some conditions.

Nash Equilibrium: Infinitely Repeated PD

Free Trade Game		
US / EU	<i>Free Trade</i>	<i>Protect</i>
<i>Free Trade</i>	10, 10	1, 12
<i>Protect</i>	12, 1	4, 4

To show that (grim trigger, grim trigger) is a Nash equilibrium, we must show that neither player has a profitable deviation along the equilibrium path.

- If each country plays grim trigger, both receive 10 in every period
- If both countries discount the future at a common factor of δ , the total utility of this strategy is $\frac{10}{1-\delta}$

Nash Equilibrium: Infinitely Repeated PD

Now let's try to formulate the strongest possible alternative strategy for US, given EU is playing grim trigger.

- Suppose the US defects in some period t .
- Since EU is playing grim trigger, they protect forever starting in $t + 1$.
- Then the US should also protect forever starting in $t + 1$.
- Thus we consider the strategies:

	1	2	3	...	t	$t+1$	$t+2$...
US	F	F	F	...	P	P	P	...
EU	F	F	F	...	F	P	P	...

Note: this is **not** a one-shot deviation; we're not yet solving for SPNE. (What would a one-shot deviation look like for the US?)

Nash Equilibrium: Infinitely Repeated PD

- This strategy yields for the US a payoff of 12 in the first defection period (t) and a stream of 4 forever after
- The US is better off playing grim trigger if and only if:

$$\begin{aligned}\frac{10}{1-\delta} &\geq 12 + \frac{\delta 4}{1-\delta} \\ 10 &\geq 12(1-\delta) + 4\delta \\ 8\delta &\geq 2 \\ \delta &\geq \frac{1}{4}\end{aligned}$$

- Thus, as long as the players are sufficiently patient (δ is sufficiently large), both players playing the grim trigger strategy is a Nash equilibrium.

Generalized Prisoner's Dilemma

Generalized Prisoner's Dilemma		
1 / 2	Cooperate	Don't cooperate
Cooperate	a, a	d, c
Don't cooperate	c, d	b, b

where $c > a > b > d$. Using exactly the same arguments, the grim trigger strategy is a Nash equilibrium if and only if $\frac{a}{1-\delta} \geq c + \frac{\delta b}{1-\delta}$. Rearranging yields the condition:

$$\delta \geq \frac{c - a}{c - b}.$$

Generalized Prisoner's Dilemma

Thus, cooperation is harder to sustain (requires a higher discount factor) when:

- 1 c is large relative to a and b
- 2 a and b are roughly equal.

Generalized Prisoner's Dilemma: SPNE

We have derived the condition on δ under which (grim trigger, grim trigger) is a Nash equilibrium. But is it an SPNE?

We now have to check every subgame, including ones off the equilibrium path. **Fortunately, we only have to check one-shot deviations.**

Only one additional type of subgame is relevant to us: one in which Player 1 has defected at time $t - 1$, which is off the equilibrium path. Then both players playing grim trigger would dictate:

	1	2	...	$t - 1$	t	$t + 1$	$t + 2$...
Player 1	C	C	...	<i>D</i>	C	<i>D</i>	<i>D</i>	...
Player 2	C	C	...	C	<i>D</i>	<i>D</i>	<i>D</i>	...

Generalized Prisoner's Dilemma: SPNE

Is defecting forever a credible threat for Player 2?

- Her payoffs (starting from period t) from sticking to grim trigger are:

$$c + \delta b + \delta^2 b + \delta^3 b \dots = c + \frac{\delta b}{1 - \delta}$$

- But what if she just turned a blind eye to Player 1's defection and cooperated in period t instead?

	1	2	...	$t-1$	t	$t+1$	$t+2$...
Player 1	C	C	...	<i>D</i>	C	C	C	...
Player 2	C	C	...	C	C	C	C	...

This would yield the payoff stream $\frac{a}{1-\delta}$.

Generalized Prisoner's Dilemma: SPNE

Player 2 is better off sticking to grim trigger if and only if:

$$\begin{aligned}c + \frac{\delta b}{1 - \delta} &\geq \frac{a}{1 - \delta} \\c(1 - \delta) + \delta b &\geq a \\ \delta(b - c) &\geq a - c \\ \delta &\leq \frac{a - c}{b - c}\end{aligned}$$

Where we have flipped the inequality because $b - c < 0$.

Thus, (grim trigger, grim trigger) is subgame perfect under the **knife-edge condition** $\delta = \frac{c-a}{c-b}$, since the previously derived condition for Nash equilibrium also has to hold.

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Tit-for-Tat Strategies

- The grim trigger strategy is not the only equilibrium of the infinitely repeated prisoner's dilemma that sustains the cooperative outcome.
- The grim trigger equilibrium may be undesirable because cooperation disappears forever following a single defection.
 - It is not robust to mistakes by the players.
 - Following a breakdown of cooperation the players cannot renegotiate to return to the cooperative phase, something they clearly have an incentive to do.
- An alternative Nash equilibrium is based on “tit-for-tat” strategies of the form: “cooperate in the first period; then, in any subsequent period, play the action that the other player chose in the previous period.”

Tit-for-Tat Strategies: Subgame Perfection

There are two types of subgames to consider:

- ① A subgame where, following cooperation in the previous period, both players are expected to cooperate in the future. This is the **cooperation phase**.
 - We compare the utility of tit-for-tat to the utility of a unilateral, one-period deviation to **not cooperating**, then returning to tit-for-tat.
- ② A subgame following cooperation by one player and defection by the other. This is the **punishment phase**.
 - We compare the utility of tit-for-tat to the utility of **not punishing** and cooperating in the period immediately following defection, then returning to tit-for-tat.

Tit-for-Tat Strategies: Subgame Perfection

① **The cooperation phase.** Consider Player 1:

- Following tit-for-tat yields the present value utility $\frac{a}{1-\delta}$.
- A one-shot defection in the cooperation phase (while Player 2 adheres to tit-for-tat) looks like:

	1	2	...	$t-1$	t	$t+1$	$t+2$	$t+3$...
Player 1	C	C	...	C	<i>D</i>	C	<i>D</i>	C	...
Player 2	C	C	...	C	C	<i>D</i>	C	<i>D</i>	...

This payoff stream yields the present value utility:

$$\begin{aligned} & c + \delta d + \delta^2 c + \delta^3 d + \delta^4 c + \delta^5 d \dots \\ &= c + \delta d + \delta^2(c + \delta d) + \delta^4(c + \delta d) \dots \\ &= \frac{c + \delta d}{1 - \delta^2} \end{aligned}$$

Tit-for-Tat Strategies: Subgame Perfection

(Referring back to the generalized payoff matrix:)

Generalized Prisoner's Dilemma		
1 / 2	Cooperate	Don't cooperate
Cooperate	a, a	d, c
Don't cooperate	c, d	b, b

where $c > a > b > d$.

Thus, tit-for-tat is an SPNE if and only if:

$$\frac{a}{1 - \delta} \geq \frac{c + \delta d}{1 - \delta^2}$$

Rearranging allows us to express this as a condition on the discount rate.

Tit-for-Tat Strategies

$$\frac{a}{1-\delta} \geq \frac{c+\delta d}{1-\delta^2}$$

$$a \geq \frac{c+\delta d}{1+\delta}$$

$$(1+\delta)a \geq c+\delta d$$

$$\delta a - \delta d \geq c - a$$

$$\delta \geq \frac{c-a}{a-d}$$

Subgame Perfection

② **The punishment phase.** Again consider it from Player 1's point of view:

- After Player 2 has defected, both players playing tit-for-tat looks like:

	1	2	...	$t-1$	t	$t+1$	$t+2$	$t+3$...
Player 1	C	C	...	C	D	C	D	C	...
Player 2	C	C	...	D	C	D	C	D	...

Which, as above, yields the payoff stream $\frac{c+\delta d}{1-\delta^2}$.

Alternatively, Player 1 could cooperate in period t and get a in every period. Player 1 prefers to play tit-for-tat as long as:

$$\frac{c + \delta d}{1 - \delta^2} \geq \frac{a}{1 - \delta} \rightarrow \delta \leq \frac{c - a}{a - d}$$

Subgame Perfection

Thus, tit-for-tat is only a **subgame perfect Nash equilibrium** if:

- 1 $\delta \geq \frac{c-a}{a-d}$, and
- 2 $\delta \leq \frac{c-a}{a-d}$

This is satisfied when $\delta = \frac{c-a}{a-d}$, a very **knife-edge** condition.

Subgame Perfect Tit for Tat (“Adjusted Tit-for-Tat”)

An alternative version of tit-for-tat avoids the problem of oscillation in the punishment phase. Milgrom, North and Weingast (1990) argue for the strategy:

- Start out playing *Cooperate*
- Always play *Cooperate* at time t unless these two conditions both hold:
 - ① The other player defected in $t - 1$
 - ② You cooperated in $t - 2$

This strategy **punishes defection** but **rewards punishment for defection**, allowing players to get back to a cooperative equilibrium.

1	2	...	$t - 2$	$t - 1$	t	$t + 1$...
C	C	...	C	C	D	C	...
C	C	...	C	D	C	C	...

Subgame Perfect Tit for Tat (“Adjusted Tit-for-Tat”)

Again, need to check two subgames: cooperation and punishment phase.

- 1 In a cooperation phase the condition under which a player continues to cooperate is:

$$\frac{a}{1-\delta} \geq c + \delta d + \frac{\delta^2 a}{1-\delta}$$

- 2 In a punishment phase, we have to check that the defector wants to return to cooperating rather than defect another period:

$$d + \frac{\delta a}{1-\delta} \geq b + \delta d + \frac{\delta^2 a}{1-\delta}$$

and that the punisher is better off punishing:

$$c + \frac{\delta a}{1-\delta} \geq \frac{a}{1-\delta}$$

(which is true by definition, $c > a$)

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Intermediate Punishment Strategies

- The grim trigger and tit-for-tat strategies represent just two of the possible strategies that sustain cooperative outcomes.
- These strategies can be generalized to include strategies that involve punishment phases of intermediate length.

Two Examples

- ① **Similar to grim trigger:** Cooperate until your opponent defects. If your opponent defects, do not cooperate for the next k periods but then return to cooperation; if you defect, do not cooperate for the next k periods but then return to cooperation. Once you have returned to cooperation, cooperate until a defection occurs.
- ② **Similar to tit-for-tat:** Cooperate until your opponent defects. If your opponent defects, do not cooperate for k periods. If she cooperates in any of the k periods, return to cooperation, ending the punishment phase. If she fails to cooperate in any period of the punishment phase, then the punishment phase starts over i.e. don't cooperate for k more periods. If your own failure to cooperate caused the punishment phase then cooperate during the punishment phase.

Strategy 1

- The payoff stream from defecting from mutual cooperation (and then defecting while you're being punished) consists of the one-period gain from defecting, b for k periods, and an infinite stream of a beginning $k + 1$ periods in the future.
- Recall that the present value of an infinite payoff stream of π beginning at a future time τ is:

$$\frac{\delta^\tau \pi}{1 - \delta}$$

- Thus the utility of defecting after cooperation is:

$$c + \frac{\delta - \delta^{k+1}}{1 - \delta} b + \frac{\delta^{k+1}}{1 - \delta} a.$$

Strategy 1

- Consequently, sustaining cooperation requires that:

$$\frac{a}{1-\delta} \geq c + \frac{\delta - \delta^{k+1}}{1-\delta} b + \frac{\delta^{k+1}}{1-\delta} a$$

Multiplying through by $1 - \delta$ yields the condition:

$$a(1 - \delta^{k+1}) \geq (1 - \delta)c + (\delta - \delta^{k+1})b$$

- We cannot generate a closed form for the critical value of δ , but we can rewrite this expression as:

$$\delta > \frac{c - a}{c - b} + \delta^{k+1} \frac{a - b}{c - b}.$$

- The first term on the right side is the critical value for the grim trigger strategy, and second term is positive for any k . Thus it is harder to sustain cooperation with a finite punishment phase. But if players make mistakes, this equilibrium may be preferred to the grim trigger.

Strategy 2

- A defection from the cooperation phase generates a payoff consisting of a one period benefit c , a punishment payoff of d for k periods, and a return to cooperative payoffs a at the end of the punishment.
- Summing these up generates $c + \frac{\delta - \delta^{k+1}}{1 - \delta} d + \frac{\delta^{k+1}}{1 - \delta} a$.
- This payoff is lower than the payoff from defection in the tit-for-tat equilibrium ($c + \delta d + \frac{\delta^2 a}{1 - \delta}$) by $\frac{\delta^2 - \delta^{k+1}}{1 - \delta} (a - d)$. Thus, increasing the length of the punishment phase **decreases the incentive to defect from the cooperative phase**.

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

The Folk Theorem

- A common theme of these examples is that so long as the agents are sufficiently patient, outcomes that are not Nash equilibria in static games may be Nash equilibria or subgame perfect equilibria of infinitely repeated games.
- In fact, any **feasible** payoff vector to an infinitely repeated game that satisfies **individual rationality** can be sustained as a SPNE so long as agents are sufficiently patient.
- This result has been around a long time in many different forms, so it's called a **Folk theorem**.

Individually Rational Payoffs

Definition

The payoff vector $v = (v_1, \dots, v_i, \dots, v_n)$ is **individually rational** if $v_i \geq \min_{s_{-i}} \{ \max_{s_i} u_i(s_i, s_{-i}) \}$ for each $i \in N$.

The value $\min_{s_{-i}} \{ \max_{s_i} u_i(s_i, s_{-i}) \}$ is the minimum stage game utility that player i attains from any strategy profile in which she plays a best response to s_{-i} . This value is identified by letting players $-i$ select s_{-i} , so as to minimize the utility to i of playing a best response to s_{-i} .

Feasibility

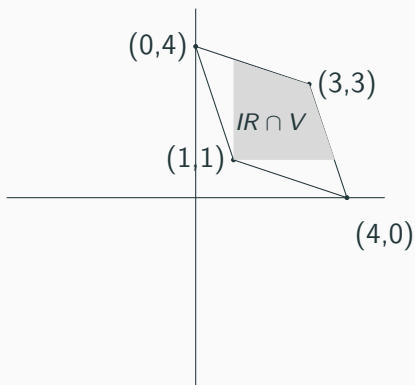
Definition (1)

The payoff vector $v = (v_1, \dots, v_i, \dots, v_n)$ is **feasible** if there is some pure strategy profile s such that for each $i \in N$,
$$u_i(s) = (1 - \delta_i)v_i.$$

Recall that $(1 - \delta)v_i$ can be understood as the discounted average of a stream of payoffs from a repeated game.

Alternatively, we call the payoffs in some stage game G feasible if they are a convex combination (i.e. a weighted average) of the pure-strategy payoffs of G (where the weights are non-negative and sum to one). We call this the **convex hull** of the pure-strategy payoffs.

Individual Rationality \cap Feasibility



	c	d
c	3,3	0,4
d	4,0	1,1

The Folk Theorem

Theorem

For every feasible and individually rational payoff vector v there is a vector of discount rates δ' (i.e. one δ'_i for each player) such that the payoff vector v occurs in a Nash equilibrium of the repeated game if $\delta_i \geq \delta'_i$ for all i .

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Bargaining Theory

- If political science is the study of “who gets what, when and how” then bargaining theory lies at its foundation.
- Legislators and executives bargain over budgets and new legislation.
- States bargain to reach new international agreements and to settle crises (e.g. refugee resettlement, climate accords...).
- Political parties bargain over coalition governments.

⋮

Bargaining Model of War

Two states are in conflict over a unit good.

- Divisible: an area of territory or an allocation of resources.
- Country 1 presents country 2 with a proposal to share the resource $(x, 1 - x)$.
- Country 2 can accept this offer (leading to peace) or reject this offer (leading to war).

War in the Bargaining Model of War

The expected payoff to war depends on the probability that a country will win, the utility of victory and defeat, and the inefficiencies of fighting.


(Normalized) payoffs for war:

- victory = 1
- defeat = 0
- cost of war $c_i > 0$
- p_i , probability of i winning: $p_2 = 1 - p_1$
- $u_i(x_i)$ is increasing and weakly concave (weakly risk-averse)

Expected utility of war to state i :

$$p_i(1) + (1 - p_i)(0) - c_i.$$

Extensive Form of One-Period Bargaining



Nash Equilibria of the Bargaining Game

What are the Nash equilibria of the one-shot bargaining game?

Anything goes.

- Pick any cutoff strategy for Player 2 (i.e., accept if $x \geq \hat{x}$ and reject otherwise)
- Have Player 1 propose the cutoff value.
- This is an equilibrium.

Proposition

In the unique subgame perfect equilibrium of this game the probability of war is zero.

Proof

- In the final stage Player 2 will accept any offer such that:

$$u_2(1 - x) \geq 1 - p - c_2.$$

- In the first stage country 1 chooses among all the x .
- It knows that for any x , it gets
 - ① either $u_1(x)$ if Player 2 accepts
 - ② or $p - c_1$ if rejects
- Of all the possible offers, one that makes Player 1's payoff largest is:

$$u_2(1 - x) = 1 - p - c_2$$

$$x = 1 - u_2^{-1}(1 - p - c_2)$$

or with linear utility in shares $x = p + c_2$

- This offer is **always accepted** and there is no war in equilibrium.

Things to Observe from Bargaining

- A conflict of interest is not sufficient for conflict.
- All else equal, stronger countries (higher p) get better deals
- When preferences are known and “fixed” bargaining produces efficient outcomes (Coase theorem)
- Bargaining power is a function of payoffs to war and the process by which agreements are reached.

What About Infinite-Horizon Bargaining Models?

- But there are many questions one might have:
 - ① What if Player 2 got to make a counter-offer instead of rejecting?
 - ② What if players have varying degrees of patience when it comes to long, drawn-out bargaining?
 - ③ Does bargaining ever end exogenously and if not, what will agreements look like?

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Rubinstein Bargaining

- Suppose that two players try to decide how to divide \$1.
- The players take turns making offers so that Player 1 proposes in periods 0, 2, 4, etc. and Player 2 makes proposals in the other periods.
- The game continues (possibly infinitely) until a proposal is accepted by the other player.
- In each period that she is the proposer, Player 1 makes an offer (x_1, x_2) where x_1 is Player 1's share and x_2 is Player 2's share where $x_1 + x_2 \leq 1$.
- If Player 2 accepts, the game ends and the dollar is divided accordingly.

Rubinstein Bargaining

- If Player 2 rejects, then she gets to make an offer (x_1, x_2) with $x_1 + x_2 \leq 1$, and the game continues if Player 1 rejects.
- To simplify matters, assume both players have linear utility functions $u_1(x_1, x_2) = x_1$ and $u_2(x_1, x_2) = x_2$.
- Each player has a discount factor δ_i ; players value proposal (x_1, x_2) accepted t periods in the future at $(\delta_1^t x_1, \delta_2^t x_2)$.
- A strategy has to consist of (1) the offers you accept when the other player proposes, and (2) the offers you make when you are the proposer.

Subgame Perfect Equilibria

- Rubinstein shows that there is a unique SPNE to this game based on playing the following strategies in every period:
 - ① Player 1 proposes $\left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}\right)$ and accepts Player 2's offer if and only if $x_1 \geq \frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}$.
 - ② Player 2 proposes $\left(\frac{\delta_1(1-\delta_2)}{1-\delta_1\delta_2}, \frac{1-\delta_1}{1-\delta_1\delta_2}\right)$ and accepts Player 1's offer if and only if $x_2 \geq \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2}$.
- Prove that these strategies constitute a subgame perfect Nash equilibrium to the alternating offers bargaining game.

Computing the Equilibrium

- Let v_1 and v_2 be the utilities of Player 1 and 2 for subgames in which they are the proposer.
- Given that the postulated strategies are the same in every period, these values are independent of t .
- These are *continuation values* because they also reflect the utility of rejecting a proposal and moving to the next subgame.

Computing the Equilibrium

- Consider a subgame where Player 1 is the proposer.
- She must offer Player 2 at least $\delta_2 v_2$, Player 2's discounted continuation value.
- She keeps the rest for herself: $x_1 = 1 - \delta_2 v_2$. Because this offer is accepted, x_1 is exactly Player 1's continuation value:
 $v_1 = x_1 = 1 - \delta_2 v_2$.
- Consider a subgame where Player 2 is the proposer. She must offer at least $\delta_1 v_1$ so that $v_2 = 1 - \delta_1 v_1$.
- Now we simply solve this system of equations:

$$v_1 = 1 - \delta_2 v_2 \text{ and } v_2 = 1 - \delta_1 v_1$$

$$\rightarrow v_2 = 1 - \delta_1(1 - \delta_2 v_2)$$

$$v_2(1 - \delta_1 \delta_2) = 1 - \delta_1$$

$$76 \quad v_2 = \frac{1 - \delta_1}{1 - \delta_1 \delta_2}$$

Computing the Equilibrium

Solving for v_1 ,

$$v_1 = 1 - \delta_2 \left(\frac{1 - \delta_1}{1 - \delta_1 \delta_2} \right)$$
$$v_1 = 1 - \frac{\delta_2 - \delta_1 \delta_2}{1 - \delta_1 \delta_2}$$
$$v_1 = \frac{1 - \delta_2}{1 - \delta_1 \delta_2}$$

Thus we have:

$$v_1 = \frac{1 - \delta_2}{1 - \delta_1 \delta_2} \text{ and } v_2 = \frac{1 - \delta_1}{1 - \delta_1 \delta_2}.$$

Plugging in these values yields the SPNE above:

- Player 1 proposes $(v_1, 1 - v_1)$ and accepts if $x_1 \geq \delta_1 v_1$
- Player 2 proposes $(v_2, 1 - v_2)$ and accepts if $x_2 \geq \delta_2 v_2$

Implications

- The model suggests a very simple path of play: in period zero, Player 1 proposes $\left(\frac{1-\delta_2}{1-\delta_1\delta_2}, \frac{\delta_2(1-\delta_1)}{1-\delta_1\delta_2} \right)$, Player 2 accepts, and the game ends.
- Because the whole dollar is allocated and there is no delay, the subgame perfect Nash equilibrium is efficient.
- If both players have the same discount factor, there is a first mover advantage because $\frac{1-\delta}{1-\delta^2} > \frac{\delta(1-\delta)}{1-\delta^2}$. Intuitively, because Player 2 discounts the future, Player 1 only needs to offer her a fraction of what she gets for being the proposer next period. Because both players are identical, Player 2 is getting only a fraction of what Player 1 gets.

Implications

- We can also compute a **comparative static**: how do equilibrium outcomes change as a function of key parameters of interest?
 - How does your equilibrium offer change as a function of your discount factor? Your opponent's discount factor?
- To answer this question, simply take the first derivative of the equilibrium outcome with respect to the parameter of interest.

$$\frac{\partial}{\partial \delta_1} \left(\frac{1 - \delta_2}{1 - \delta_1 \delta_2} \right) = \frac{(1 - \delta_1 \delta_2)(0) - (1 - \delta_2)(-\delta_2)}{(1 - \delta_1 \delta_2)^2} = \frac{\delta_2 - \delta_2^2}{(1 - \delta_1 \delta_2)^2}$$

This is **positive**, so Player 1 extracts a bigger share of the dollar the more patient she is.

Implications

What about Player 1's equilibrium share with respect to Player 2's discount rate?

$$\frac{\partial}{\partial \delta_2} \left(\frac{1 - \delta_2}{1 - \delta_1 \delta_2} \right) = \frac{(1 - \delta_1 \delta_2)(-1) - (1 - \delta_2)(-\delta_1)}{(1 - \delta_1 \delta_2)^2} = \frac{\delta_1 - 1}{(1 - \delta_1 \delta_2)^2}$$

This is **negative**, so Player 1 extracts a smaller share of the dollar the more patient Player 2 is.

If $\delta_1 = \delta_2 = \delta$, then both players' shares converge to $\frac{1}{2}$ as δ converges to 1. As both players become perfectly patient, they are less willing to accept offers that are less than what they can get as the proposer next period. In the limit, they demand exactly what they expect to get next period.

Finitely Repeated Games

Infinitely Repeated Games

Discounting and Definitions

The Grim Trigger Strategy

Tit-for-Tat Strategy

Intermediate Punishment Strategies

Folk Theorem

Examples

Example 1: Bargaining Model of War

Example 2: Rubinstein Bargaining

Example 3: Bargaining Under a Closed Rule

Majority Rule Bargaining Under A Closed Rule

- A key feature of the Rubinstein model is that unanimous consent is required to reach an agreement on the allocation.
- This rules out a number of important political settings where only a simple or supermajority is required for agreement.
- Baron and Ferejohn (1989) have extended Rubinstein's model to simple majority rule with more than two bargainers.
- Suppose that there are N (odd) players bargaining and any proposal requires $n = (N + 1)/2$ votes.
- Instead of assuming alternating offers, Baron and Ferejohn consider a bargaining protocol with a **random recognition rule**.
- According to this protocol, in each period, every player is chosen to make a proposal with an equal probability ($1/N$).

Majority Rule Bargaining Under A Closed Rule

- We focus on bargaining under a *closed rule* where the proposer makes a take-it-leave-it offer for the current legislative session.
 - The proposer in each period makes an offer (x_1, x_2, \dots, x_N) such that x_i is the share for player i .
 - Feasibility requires that $\sum x_i \leq 1$.
 - If this proposal is rejected, the session ends, discounting occurs, and a new proposer is randomly chosen at the beginning of the next session.
 - To simplify, we assume that each player has the same discount factor δ .
- This game has lots of subgame perfect equilibria. In fact for large enough N and δ , there is a SPNE that can support any feasible division of the dollar.

Majority Rule Bargaining Under A Closed Rule

- These strategies require, however, that each player know the whole (possibly infinite) history of the game in order to know which actions are consistent with the prescribed punishment
- Following Baron and Ferejohn, we analyze only **stationary equilibria**, meaning those in which:
 - ① A proposer proposes the same division every time she is recognized regardless of the history of the game.
 - ② Voters vote only on the basis of the current proposal and expectations about future proposals, not on prior histories.
- Does there exist an equilibrium with this property?

Majority Rule Bargaining Under A Closed Rule

- Let v_i be the continuation value (i.e. the discounted expected utility from playing the rest of the game) for player i .
- We focus on **symmetric equilibria** (in which every player is playing the same strategy), so that $v_i = v$ for all i .
- Any voter who gets $x_i \geq \delta v$ votes in favor of the proposal, whereas any voter who receives less than δv votes against it.
- Given these voting strategies, an optimal proposer must propose:
 - δv to $n - 1$ other players
 - 0 to the rest of the other players
 - the remainder, $z = 1 - (n - 1)\delta v$ to himself

Majority Rule Bargaining Under A Closed Rule

- In this class of stationary symmetric equilibria, the proposer chooses her coalition partners randomly.
- The continuation value is then:

$$v = \underbrace{\left(\frac{1}{N}\right)z}_{\text{being proposer}} + \underbrace{\left(\frac{n-1}{N}\right)\delta v}_{\text{being in winning coalition}} + \underbrace{\left(\frac{N-n}{N}\right)0}_{\text{being left out}}$$

- Substituting for z and simplifying yields:

$$v = \frac{1 - (n-1)\delta v}{N} + \frac{(n-1)\delta v}{N} = \frac{1}{N}$$

- The continuation value is a proportional share of the dollar.

Majority Rule Bargaining Under A Closed Rule

Finally, given our solution for v , we have to compute the proposer's share and make sure it makes the proposer better off than punting (i.e. making a proposal that won't be accepted) to get to the next period.

Recalling that $z = 1 - (n - 1)\delta v$ and plugging in $v = 1/N$,

$$z = 1 - (n - 1) \left(\frac{\delta}{N} \right)$$

Let's put n in terms of N to make the analysis clearer:

$$z = 1 - \left(\frac{N + 1}{2} - 1 \right) \left(\frac{\delta}{N} \right) = 1 - \frac{(N - 1)\delta}{2N}$$

Majority Rule Bargaining Under A Closed Rule

For z to be optimal for the proposer, we must check that:

$$\begin{aligned}1 - \frac{(N-1)\delta}{2N} &\geq \frac{\delta}{N} \\1 &\geq \frac{2\delta + (N-1)\delta}{2N} \\2N &\geq \delta(N+1) \\(2-\delta)N &\geq \delta \\N &\geq \frac{\delta}{2-\delta}\end{aligned}$$

The right-hand side is maximized at 1 when $\delta = 1$. Thus, this condition is always satisfied for more than one player.

Some Takeaways

- Because v is also the expected utility of the game, this result implies that bargaining is efficient because the sum of player utilities is maximized.
- We can compute a measure of **proposal power**: the difference between the utility of being the proposer (z) and the discounted continuation value (δv):

$$\pi = z - \delta v = 1 - \frac{(N-1)\delta}{2N} - \frac{\delta}{N} = 1 - \frac{(N+1)\delta}{2N}$$

Comparative statics: How does proposal power vary with δ and N ?

Comparative Statics

Proposal power decreases as players become more patient:

$$\frac{\partial \pi}{\partial \delta} = -\frac{N+1}{2N}$$

What about as N grows large?

$$\frac{\partial \pi}{\partial N} = -\left(\frac{2N\delta - \delta(N+1)(2)}{4N^2}\right) = -\left(-\frac{\delta}{2N^2}\right) = \frac{\delta}{2N^2}$$

Since this is positive, proposal power grows as the size of the legislature increases.

Supermajority Rule (Try this at home)

- Now assume that $k > n$ votes are required.
- Repeating the steps above, you can easily derive the proposer's share as:

$$z = 1 - (k - 1)\delta v$$

and the continuation values as:

$$v = \frac{z}{N} + \frac{k - 1}{N}\delta v$$

- Algebra reveals that once again $v = \frac{1}{N}$
- The proposer's equilibrium share is now lowered to $z = 1 - \frac{\delta(k-1)}{N}$ (from $z = 1 - \frac{\delta(n-1)}{N}$). Thus, going from majority to supermajority rule mitigates the proposer's advantage.

Implications for Institutional Design

- Proposal power: can be interpreted as committee membership
→ how does the game change when proposal power is non-random?
- What features control proposal power?
 - Bigger legislature → more proposal power
 - Supermajoritarian rules → less proposal power
 - Probability of being recognized is another lever
- Do we want to increase or decrease proposal power from a normative standpoint?
 - More unequal offers may be less fair
 - But, high payoffs to committee membership may incentivize good policy

MIT OpenCourseWare
<https://ocw.mit.edu>

17.810 / 17.811 Game Theory
Spring 2021

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>.