Repeated games

Felix Munoz-Garcia

EconS 503 - Microeconomic Theory II
Washington State University

Repeated games are very usual in real life:

1. Treasury bill auctions (some of them are organized monthly, but some are even weekly),
2. Cournot competition is repeated over time by the same group of firms (firms simultaneously and independently decide how much to produce in every period).
3. OPEC cartel is also repeated over time.

- In addition, players' interaction in a repeated game can help us rationalize cooperation...
    - in settings where such cooperation could not be sustained should players interact only once.

We will therefore show that, when the game is repeated, we can sustain:

1. Players' cooperation in the Prisoner's Dilemma game,
2. Firms' collusion:
    1. Setting high prices in the Bertrand game, or
    2. Reducing individual production in the Cournot game.
3. But let's start with a more "unusual" example in which cooperation also emerged: Trench warfare in World War I.

Trench warfare in World War I

Trench warfare in World War I

- Despite all the killing during that war, peace would occasionally flare up as the soldiers in opposing tenches would achieve a truce.
- Examples:
  - The hour of 8:00-9:00am was regarded as consecrated to "private business,"
  - No shooting during meals,
  - No firing artillery at the enemy's supply lines.
- One account in Harrington:
  - After some shooting a German soldier shouted out "We are very sorry about that; we hope no one was hurt. It is not our fault, it is that dammed Prussian artillery"
- But... how was that cooperation achieved?

Trench warfare in World War I

- We can assume that each soldier values killing the enemy, but places a greater value on not getting killed.
- That is, a soldier's payoff is

$$4 + 2 \times (\text{enemy soldiers killed}) - 4(\text{own soldiers killed})$$

- This incentive structure produces the following payoff matrix,
  - This matrix represents the so-called "stage game", i.e., the game players face when the game is played only once.

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	2, 2	6, 0
	Miss	0, 6	4, 4

Trench warfare in World War I

- **Where are these payoffs coming from?**
  - For instance, $(Miss, Kill)$ implies a payoff pair of $(0, 6)$ since

$$u_{Allied} = 4 + 2 * 0 - 4 * 1 = 0, \text{ and}$$
$$u_{German} = 4 + 2 * 1 - 4 * 0 = 6$$

  - Similarly, $(Kill, Kill)$ entails a payoff pair of $(2, 2)$ given that

$$u_{Allied} = 4 + 2 * 1 - 4 * 1 = 2, \text{ and}$$
$$u_{German} = 4 + 2 * 1 - 4 * 1 = 2$$

Trench warfare in World War I

- **If this game is played only once...**

Allied Soldiers \ German Soldiers	*Kill*	*Miss*
Kill	2, 2	6, 0
Miss	0, 6	4, 4

- (*Kill*, *Kill*) is the **unique NE** of the stage game (i.e., unrepeated game).
- In fact, "Kill" is here a **strictly dominant strategy** for both players,
  - making this game strategically equivalent to the standard PD game (where confess was strictly dominant for both players).

Trench warfare in World War I

- But we know that such a game was not played only once, but many times.
- For simplicity, let's see what happens if the game is played twice. Afterwards, we will generalize it to more than two repetitions.
  - (See the extensive form game in the following slide)

Trence warfare in World War I

- Twice-repeated trench warfare game

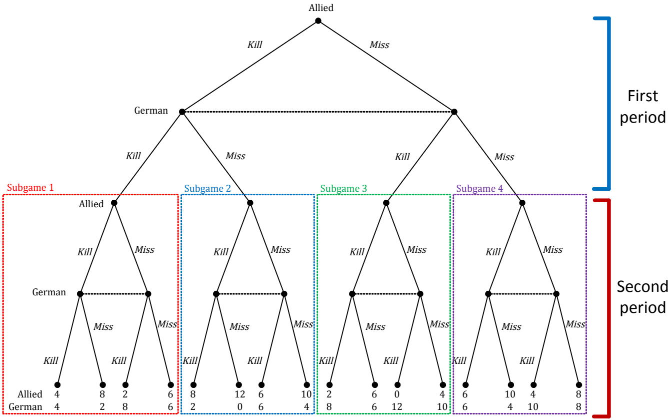

Trench warfare in World War I

- We can solve this twice-repeated game by using backward induction (starting from the second stage):
- **Second stage:**
  - We first identify the proper subgames: there are four, as indicated in the figure, plus the game as a whole.
  - We can then find the NE of each of these four subgames separately.
  - We will then be ready to insert the equilibrium payoffs from each of these subgames, constructing a reduced-form game.
- **First stage:**
  - Using the reduced-form game we can then solve the first stage of the game.

Trench warfare in World War I

- Subgame 1 (initiated after (Kill Kill) arises as the outcome of the first-stage game):

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	$\underline{4}, \underline{4}$	$\underline{8}, 2$
	Miss	$2, \underline{8}$	$6, 6$

- Only one psNE of Subgame 1: $(Kill, Kill)$.

Trench warfare in World War I

- Subgame 2 (initiated after (Kill Miss) outcome emerges the first-stage game)

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	<u>8</u>, <u>2</u>	<u>12</u>, 0
	Miss	6, <u>6</u>	10, 4

- Only one psNE of Subgame 2: $(Kill, Kill)$.

Trench warfare in World War I

- Subgame 3 (initiated after (Miss, Kill) outcome in the first stage):

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	$\underline{2}, \underline{8}$	$\underline{6}, 6$
	Miss	$0, \underline{12}$	$4, 10$

- Only one psNE of Subgame 3: $(Kill, Kill)$.

Trench warfare in World War I

- Subgame 4 (initiated after the (Miss, Miss) outcome in the first stage):

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	$\underline{6}, \underline{6}$	$\underline{10}, 4$
	Miss	$4, \underline{10}$	$8, 8$

- Only one psNE of Subgame 4: $(Kill, Kill)$.

Trench warfare in World War I

- Inserting the payoffs from each subgame, we now construct the reduced-form game:

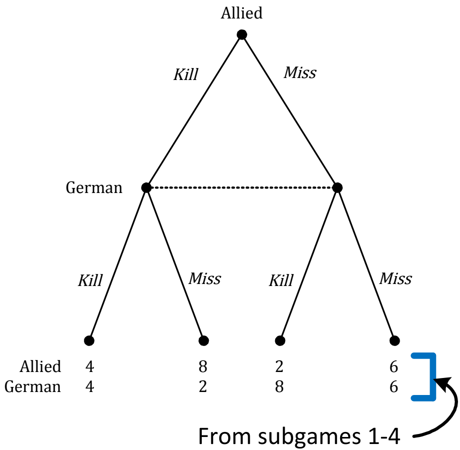

Trench warfare in World War I

- Since the above game tree represents a simultaneous-move game, we construct its Normal-form representation:

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	$\underline{4}, \underline{4}$	$\underline{8}, 2$
	Miss	$2, \underline{8}$	$6, 6$

- We are now ready to summarize the Unique SPNE:
    - Allied Soldiers: $(Kill_1, Kill_2$ regardless of what happened in period 1)
    - German Soldiers: $(Kill_1, Kill_2$ regardless of what happened in period 1)

Trench warfare in World War I

- But then the SPNE has both players shooting to kill during both period 1 and 2!!
- As Harrington puts it:
  - Repeating the game only twice "was a big fat failure!" in our goal to rationalize cooperation among players.
- Can we avoid such unfortunate result if the game is, instead, played $T > 2$ times? Let's see... (next slide)
  - Caveat: we are still assuming that the game is played for a finite T number of times.

What if the game was repeated T periods?

- This would be the normal form representation of the subgame of the last period, T.
  - A^{T-1} denotes the sum of the Allied soldier's previous $T-1$ payoffs.
  - G^{T-1} denotes the sum of the German soldier's previous $T-1$ payoffs.

		German Soldiers	
		Kill	*Miss*
Allied Soldiers	*Kill*	$\underline{A^{T-1}+2}, \underline{G^{T-1}+2}$	$\underline{A^{T-1}+6}, G^{T-1}$
	Miss	$A^{T-1}, \underline{G^{T-1}+6}$	$A^{T-1}+4, G^{T-1}+4$

- Only one psNE in the subgame of the last stage of the game: $(Kill_T, Kill_T)$.

What if the game was repeated T periods?

- Given the $(Kill_T, Kill_T)$ psNE of the stage-T subgame, the normal form representation of the subgame in the $T-1$ period is:

German Soldiers

Allied Soldiers	*Kill*	*Miss*
Kill	$\underline{A^{T-2}+4}, \underline{G^{T-2}+4}$	$\underline{A^{T-2}+8}, G^{T-2}+2$
Miss	$A^{T-2}+2, \underline{G^{T-2}+8}$	$A^{T-2}+6, G^{T-2}+6$

- Again, only one psNE in the subgame of period $T-1$.
- Similarly for any other period $T-2, T-3, \ldots, 1$.

Trench warfare in World War I

- But this is even worse news than before:
  - Cooperation among players cannot be sustained when the game is repeated a *finite number of times*, T (not for $T = 2$ or $T > 2$).

Trench warfare in World War I

- **Intuition:**
  - Sequential rationality demands that each players behaves optimally at every node (at every subgame) at which he/she is called on to move.
  - In the last period T, your action does not affect your previous payoffs, so you'd better maximize your payoff at T (how? shooting to kill).
  - In the $T-1$, your action does not affect your previous payoffs nor your posterior payoffs —since you can anticipate that the NE of the posterior subgame is $(kill_T, kill_T)$— so you'd better maximize your payoff at $T-1$ (how? shooting to kill).
  - Similarly at the $T-2$ period... and all other periods until the first.

Finitely repeated games

- This result provides us with some interesting insight:
  - **Insight:** If the stage game we face has a unique NE, then there is a unique SPNE in the finitely-repeated game in which all players behave as in the stage-game equilibrium during all T rounds of play.
  - Examples:
    - Prisoner's dilemma,
    - Cournot competition,
    - Bertrand competition (both with homogeneous and differentiated products).
    - etc.
- What about games with more than one NE in the stage game? (Let's get into them now!).

Finitely repeated games

- Consider the following game, from Tadelis (section 10.1)

		Player 2		
		m	f	r
Player 1	M	$4, 4$	$-1, \underline{5}$	$0, 0$
	F	$\underline{5}, -1$	$\underline{1}, \underline{1}$	$0, 0$
	R	$0, 0$	$0, 0$	$\underline{3}, \underline{3}$

- The unrepeated version of the game has two psNE:
  - (R, r) and (F, f), where the former is Pareto superior to the latter by both players.
- However, (M, m) does not arise in the unrepeated version of the game.
  - Can it emerge when the game is repeated a finite number of rounds, e.g., twice? Yes, through a stick and carrot strategy.

Stick-and-carrot strategies

- **Player 1:** Play M in stage 1. In stage 2, play R if (M, m) was played in stage 1, but play F otherwise.
- **Player 2:** Play m in stage 1. In stage 2, play r if (M, m) was played in stage 1, but play f otherwise.
- Before we get into the proof, note that, while the stick-and-carrot strategy might specify an outcome in the first stage that is not sustained as NE in the unrepeated version of the game, such as (M, m)...
  - it must specify an outcome that constitutes a NE of the unrepeated game in the last stage, such as (R, r).
  - Otherwise, players would have incentives to deviate in the last stage irrespective of previous history.

Stick-and-carrot strategies

- **Proof:** As the game is finite and sequential, we can operate by backward induction to test whether the above strategy is a SPNE of the twice-repeated game.

Stick-and-carrot strategies

- **Proof (cont'd):**
- If player 1 observes that the outcome in stage 1 was **different** from (M, m), then he anticipates that player 2 will be playing f. In this setting, player 1's best response is to play F, i.e., $BR_1(f) = F$ (see middle column of the payoff matrix).
  - A similar argument applies for player 2, i.e., anticipating F he best responds with f.
- If, instead, player 1 observes outcome (M, m), then he anticipates that player 2 will select r in the second stage, and his best response is $BR_1(r) = R$.
  - Similarly for player 2, after outcome (M, m), he anticipates that player 1 will choose R, and $BR_1(R) = r$.

Stick-and-carrot strategies

- **Intuition:** If a player i deviates from (M, m) in the first period (for instance, playing F for player 1 in order to see his first-period payoff increase from 4 to 5)...
  - he would induce the "punishing" outcome of (F, f) in the second period, and thus a payoff of only 1,
  - rather than outcome (R, r) and thus a payoff of 3.

Stick-and-carrot strategies

- Thus, every player *i* behaves as prescribed by the "stick-and-carrot" strategy if

$$4 + \delta 3 > 5 + \delta 1 \Longrightarrow \delta > \frac{1}{2}$$

- That is, players must assign a sufficiently high value to future payoffs.

Infinitely repeated games

- In finitely repeated games, players know when the game will end: in $T = 2$ periods, in $T = 7$ periods, etc.
- But... what if they don't?
    - This setting illustrates several strategic contexts where firms/agents simply know that there is a positive probability they will interact again in the next period
    - For instance, the soldiers know that there is a probability $p = 0.7$ that war will continue the next day, allowing for the game to be repeated an infinite number of times.
    - Example: After $T = 100$ rounds (e.g. days), the probability two soldiers interact one more round is 0.7100 (which is one in millions!)
- Let us analyze the **infinitely-repeated version** of this game.

Infinitely repeatead games - Payoffs

- **Present value:** Please read section 10.2 in Tadelis, for details about the present value of an infinite sequence of payoffs $\{v_i^t\}_{t=1}^{\infty}$

$$v_i = v_i^1 + \delta v_i^2 + \delta^2 v_i^3 + \ldots = \sum_{t=1}^{\infty} \delta^{t-1} v_i^t$$

and why do we need to impose $\delta < 1$ for such present value to be bounded.

Infinitely repeatead games - Payoffs

- **Average payoff:** In addition, for an infinite sequence of payoffs $\{v_i^t\}_{t=1}^{\infty}$, we define the average payoff $\overline{v}_i$ of player i as the payoff he must obtain in every round of the game for his overall payoff to be $v_i = \sum_{t=1}^{\infty} \delta^{t-1} v_i^t$. That is,

$$\sum_{t=1}^{\infty} \delta^{t-1} \overline{v}_i = v_i \Longleftrightarrow \frac{\overline{v}_i}{1-\delta} = v_i$$

which solving for $\overline{v}_i$ yields

$$\overline{v}_i = (1-\delta) v_i = (1-\delta) \sum_{t=1}^{\infty} \delta^{t-1} v_i^t$$

Trench warfare - infinitely repeated version

- First, note that $(kill_t, kill_t)$ at every period t is still one of the SPNE of the infinitely repeated game game.
- In order to show that, note that if a player chooses $kill_t$ at every period t, he obtains

$$2 + \delta 2 + \delta^2 2 + ... = \frac{1}{1-\delta} 2$$

- If, instead, he unilaterally deviates to "miss" at a particular time period, he obtains

$$\overbrace{0}^{\text{Payoff when he misses but his opponent shoots to kill}} + \overbrace{\delta 2 + \delta^2 2 + ...}^{\text{Discounted stream of payoffs when this player reverts to kill (the NE of the stage game).}}$$

$$= \delta [1 + \delta 2 + ...] = \frac{\delta}{1-\delta} 2$$

Trench warfare - infinitely repeated version

- Hence, this player does not deviate from $kill_t$ since

$$\frac{1}{1-\delta}2 > \frac{\delta}{1-\delta}2 \Leftrightarrow 2 > 2\delta \Leftrightarrow 1 > \delta$$

is satisfied given that the discount factor is restricted by definition in the range $\delta \in (0,1)$.

Trench warfare - infinitely repeated version

- But, can we sustain **cooperation** as a SPNE of this infinitely-repeated game? Yes!
- Consider the following symmetric strategy:
  - In period $t = 1$, choose "miss" (i.e., cooperate).
  - In period $t \geq 2$,
    - keep choosing "miss" if both armies chose "miss" in all previous periods, or
    - choose "kill" thereafter for any other history of play, i.e., if either army chose "kill" in any previous period.
- This strategy is usually referred to as a **Grim-Trigger strategy**, because any deviation triggers a grim punishment thereafter. Note that the punishment implies reverting to the NE of the unrepeated version of the game $(kill_t, kill_t)$.

Trench warfare - infinitely repeated version

- We need to show that such Grim-Trigger strategy (GTS) is a SPNE of the game.
- In order to show that, we need to demonstrate that it is an optimal strategy for both players at every subgame at which they are called on to move. That is, using the GTS strategy must be optimal:
  - at *any* period t, and
  - after *any* previous history (e.g., after cooperative rounds of play and after periods of non-cooperation).
- A formidable task? Not so much!
- In fact, there are only two cases we need to consider.

Trench warfare - infinitely repeated version

- Only two cases we need to consider.
- **First case:** Consider a period t and a previous history in which every one has been cooperative (i.e., no player has ever chosen "kill.")
  - If you choose miss (cooperate), your stream of payoffs is

$$4 + \delta 4 + \delta^2 4 + \ldots = \frac{1}{1-\delta} 4$$

  - If, instead, you choose to kill (defect), your payoffs are

$$\underbrace{6}_{\text{You choose to deviate towards "kill" while your opponent behaves cooperatively by "missing"}} + \underbrace{\delta 2 + \delta^2 2 + \ldots}_{\text{Then your opponent detects your defection (one of his soldiers dies!) and reverts to kill thereafter.}} = 6 + \frac{\delta}{1-\delta} 2$$

Trench warfare - infinitely repeated version

- **Second case:** Consider now that at period t some army has previously chosen to kill. We need to show that sticking to the GTS is optimal, which in this case implies implementing the punishment that GTS prescribes after defecting deviations.
  - If you choose kill (as prescribed), your stream of payoffs is

$$2 + \delta 2 + \delta^2 2 + ... = \frac{1}{1-\delta} 2$$

  - If, instead, you choose to miss, your payoffs are

$$0 + \delta 2 + \delta^2 2 + ... = \frac{\delta}{1-\delta} 2$$

  - After this history, hence, you prefer to choose kill since $\delta < 1$.

Trench warfare - infinitely repeated version

- We can hence conclude that the GTS is a SPNE of the infinitely-repeated game if

$$\frac{1}{1-\delta}4 \geq 6 + \frac{\delta}{1-\delta}2 \longleftarrow \text{ Unique Condition.}$$

- Multiplying both sides by $(1-\delta)$, we obtain

$$4 \geq 6 + 2(1-\delta)$$

  and solving for δ, we have $\delta \geq \frac{1}{2}$.

  - that is, players must assign a sufficient high value of payoffs received in the future (more than 50%)

Trench warfare - infinitely repeated version

- This condition is graphically represented in the following figure:

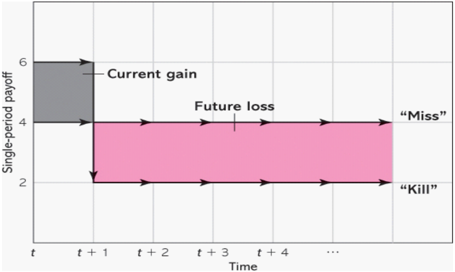

- **Intuition:** if I sufficiently care about future payoffs, I won't deviate since I have much to lose.

Finitely repeated prisoner's dilemma

		Player 2	
		Coop	*Defect*
Player 1	*Coop*	$2, 2$	$0, \underline{3}$
	Defect	$\underline{3}, 0$	$\underline{1}, \underline{1}$

- **Finitely repeated game:** Note that the SPNE of this game is (Defect, Defect) during all periods of time.
- Using backward induction, the last player to move (during the last period that the game is played) defects. Anticipating that, the previous to the last defects, and so on (unraveling result).
- Hence the unique SPNE of the finite repeated PD game has both players defecting in every round.

Infinitely repeated prisoner's dilemma

- **Infinitely repeated game:** They can support cooperation by using, for instance, Grim-Trigger strategies.
- For every player i, the Grim-Trigger strategy prescribes:
  1. Choose C at period $t = 1$, and
     Choose C at period $t > 1$ if all players selected C in previous periods.
  2. Otherwise (if some player defected), play D thereafter.

  - At any period t in which players have been cooperating in all previous rounds, every player i obtains the following payoff stream from cooperating

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

Infinitely repeated prisoner's dilemma

- And if any player i defects during a period t, while all other players cooperate, then his payoff stream becomes

$$\underbrace{3}_{\text{current gain}} + \underbrace{1\delta + 1\delta^2 + 1\delta^3 + ...}_{\text{future punishment}} = 3 + 1(\delta + \delta^2 + \delta^3 + ...)$$

$$= 3 + 1\frac{\delta}{1-\delta}$$

Infinitely repeated prisoner's dilemma

- Hence, from any period t, player i prefers to keep his cooperation (instead of defecting) if and only if

$$EU_i(\text{Coop}) \geq EU_i(\text{Defect}) \iff 2\frac{1}{1-\delta} \geq 3 + 1\frac{\delta}{1-\delta}$$

  and solving for δ, we obtain that cooperation is supported as long as $\delta \geq \frac{1}{2}$.

- (Intuitively, players must be "sufficiently patient" in order to support cooperation along time).

- Graphical illustration of:
    1. short-run increase in profits from defecting (relative to respecting the cooperative agreement); and
    2. long-run losses from being punished forever after (relative to respecting the cooperative agreement).

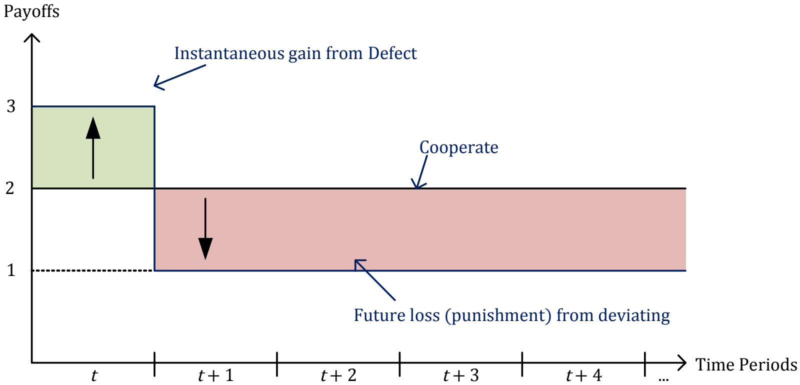

Payoffs
Instantaneous gain from Defect
3
Cooperate
2
1
Future loss (punishment) from deviating
Time Periods
t
t + 1
t + 2
t + 3
t + 4
...

- Introducing the role of δ in the previous figure:
  - A discount factor δ close to zero "squeezes" the future loss from defecting today.

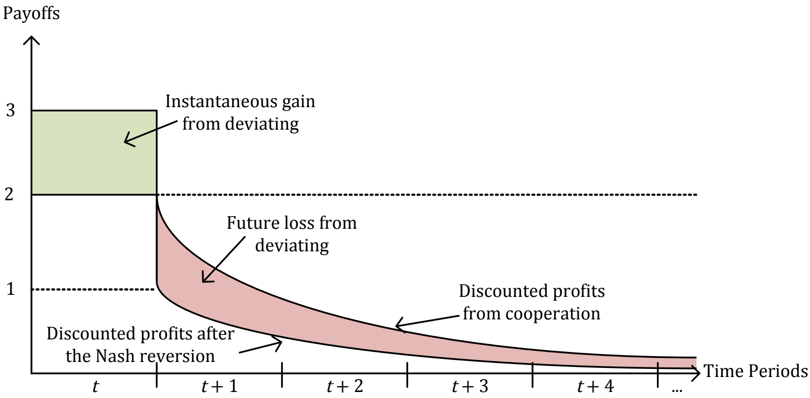

More SPNE in the repeated game

- Watson: pp. 263-271
- So far we showed that the outcome where players choose cooperation (C, C) in all time periods can be supported as a SPNE for sufficiently high discount factors, e.g., $\delta \geq \frac{1}{2}$.
- We also demonstrated that the outcome where players choose defection (D, D) in all time periods can also be sustained as a SPNE for all values of δ.
- But, can we support other partially cooperative equilibria?
  - *Example*: cooperate during 3 periods, then defect for one period, then start over, which yields an average per-period payoff lower than that in the (C, C) outcome but still higher than the (D, D) outcome.
  - Yes!

More SPNE in the repeated game

- Before we show how to sustain such a partially cooperative equilibria, let's be more general and explore all per-period payoff pairs that can be sustained in the infinitely-repeated PD game.
- We will do so with help of the so called "Folk Theorem"

The Folk Theorem

- Define the set of feasible payoffs (FP) as those inside the following diamond.$\longrightarrow$
  - (Here is our normal form game again, for reference)

Player 1 \ Player 2	*Coop*	*Defect*
Coop	2, 2	0, <u>3</u>
Defect	<u>3</u>, 0	<u>1</u>, <u>1</u>

The Folk Theorem

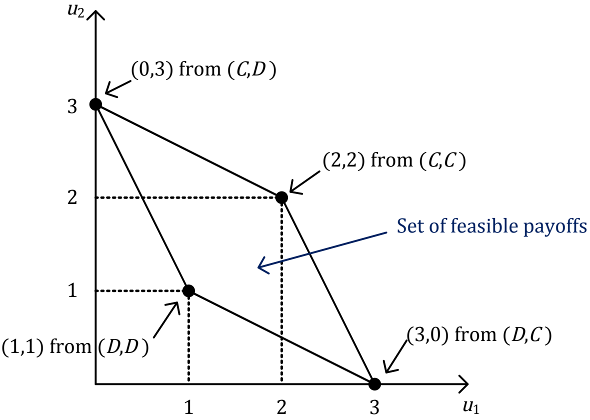

The Folk Theorem

- Why do we refer to these payoffs as feasible?
  - you can draw a line between, for instance, (2,2) and (1,1). The midpoint would be achieved if players randomize between cooperate and defect with equal probabilities.
  - Other points in this line (and other lines connecting any two vertices) can be similarly constructed to implement other points in the diamond

The Folk Theorem

- Define the set of individually rational payoffs (IR) as those that weakly improve player i's payoff from the payoff he obtains in the Nash equilibrium of the stage game, $\bar{v}_i$.
- (In this example, $\bar{v}_i = 1$ for all player $i = \{1, 2\}$).

The Folk Theorem

- Individual rational (IR) set

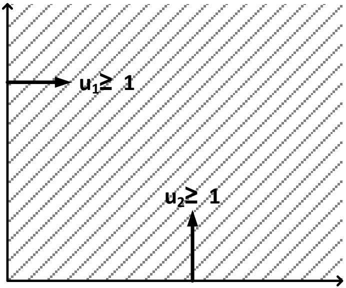

- We consider the set of feasible and individually rational payoffs, denoting it as the FIR set.
- We overlap the two sets FP and IR,and FIR is their intersection (common region).

The Folk Theorem

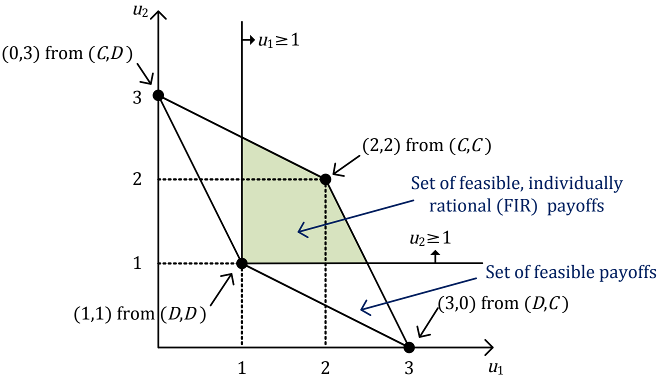

- FIR: $u_i \geq$ maximin payoff for player i, e.g., $u_1 \geq 1$ $u_2 \geq 1$
  - For simple games with a unique psNE, this payoff coincides with the psNE payoff. (We know that from the chapter on maximin strategies.)

The Folk Theorem

- Therefore, any point on the edge or interior of the shaded FIR diamond can be supported as a SPNE of the infinitely-repeated game as long as:
    - The discount factor δ is close enough to 1 (players care about the future).

The Folk Theorem (more formally)

- Consider any infinitely-repeated game.
- Suppose there is a Nash equilibrium that yields an equilibrium payoff vector $\bar{v}_i$ for every player i in the unrepeated version of the game.
- Let $v = (v_1, v_2, ..., v_n)$ be any feasible average per-period payoff such that every player i obtains a weakly higher payoff than in the Nash equilibrium of the unrepeated game, i.e., $v_i \geq \bar{v}_i$ for every player i.
- Then, there exists a sufficiently high discount factor $\delta \geq \bar{\delta}$ for which the payoff vector $v = (v_1, v_2, ..., v_n)$ can be supported as a SPNE of the infinitely-repeated game.

Another example:

- Here is another version of the repeated prisoner's dilemma game:

		Player 2: *Coop*	Player 2: *Defect*
Player 1	*Coop*	$3, 3$	$0, \underline{5}$
	Defect	$\underline{5}, 0$	$\underline{1}, \underline{1}$

- (FP set on next slide)$\longrightarrow$

Another example:

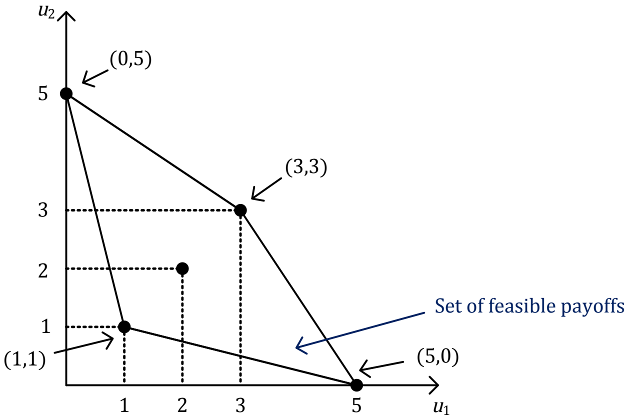

Another example:

Since the NE of the unrepeated game is (Defect, Defect), with equilibrium payoffs (1,1), then we know that the IR set must be to the northeast of (1,1) for both players to be weakly better.

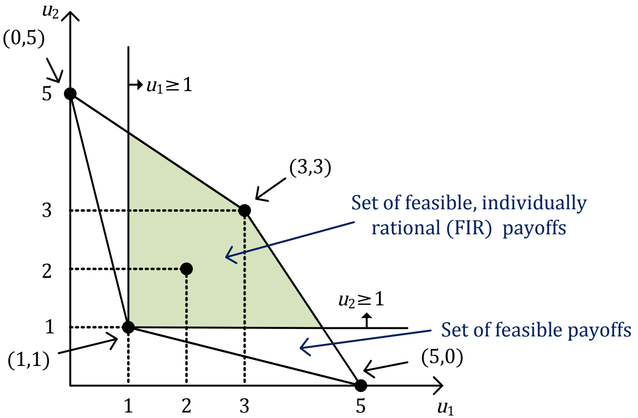

Can (C,C) be supported as a SPNE of the game?

- In any given time period t in which cooperation has been always observed in the past, if player *i* cooperates, he *i* obtains

$$3 + \delta 3 + \delta^2 3 + \ldots = \frac{3}{1-\delta}$$

- If, instead, he deviates his stream of discounted payoffs become

$$\underbrace{5}_{\text{Current}} + \underbrace{\delta 1 + \delta^2 1 + \ldots}_{\text{Future punishment}} = 5 + \frac{\delta}{1-\delta}$$

Can (C,C) be supported as a SPNE of the game?

- Hence, comparing the two payoff streams and solving for δ,

$$\frac{3}{1-\delta} \geq 5 + \frac{\delta}{1-\delta} \implies 3 \geq 5(1-\delta) + \delta$$

$$\implies 3 \geq 5 - 4\delta \implies 4\delta \geq 2 \implies \delta \geq \frac{1}{2}$$

Partial cooperation

- So far we just showed that the upper right-hand corner of the FIR diamond can be sustained as a SPNE of the infinitely repeated game.
  - What about other payoff pairs that belong to the FIR set, such as the points on the edges of the FIR diamond?
- Take, for instance, the average per-period payoff (4,1.5) in the frontier of the set of FIR payoffs.

Partial cooperation

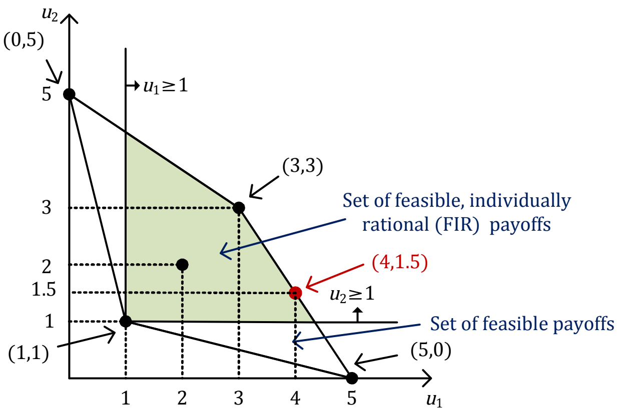

Partial cooperation

- Intuitively, we must construct a randomization between outcome (C,C) and (D,C) in order to be at a point in the line connecting the two outcomes in the FIR diamond.

1. Let us consider the following "modified grim-trigger strategy":
    1. players alternate between (D,C) and (C,C) over time, starting with (C,C) in the first period.
    2. If either or both players has deviated from this prescription in the past, players revert to the stage Nash profile (D,D) forever.

Partial cooperation

Modified Grim Trigger Strategy that alternates between (C,C) and (D,C) outcomes

	Player 1		Player 2		Resulting
	Action	Payoff	Action	Payoff	outcome
$t = 1$	C $\longrightarrow$	3	C $\longrightarrow$	3	(C,C)
$t = 2$	D $\longrightarrow$	5	C $\longrightarrow$	0	(D,C)
$t = 3$	C $\longrightarrow$	3	C $\longrightarrow$	3	(C,C)
$\vdots$					

Partial cooperation

2. To determine whether this strategy profile is a SPNE, we must compare each player's short-run gain from deviating to the associated punishment he would suffer.
   - Since the actions that this modified GTS prescribes for each player are asymmetric (player 2 always plays C as long everyone cooperated in the past, whereas player 1 alternates between C and D), we will have to separately analyze player 1 and 2.
   - Let's start with player 2.

Partial cooperation

1. **Player 2:** Starting with player 2, his sequence of discounted payoffs (starting from any *odd-numbered* period, in which players select (C,C)) is:

$$3 + 0\delta + 3\delta^2 + 0\delta^3 + ... =$$

$$= 3[1 + \delta^2 + \delta^4 + ...] + 0\delta[1 + \delta^2 + \delta^4 + ...]$$

$$= \frac{3}{1 - \delta^2}$$

And starting from any *even-numbered* period (in which players select (D,C)) player 2's sequence of discounted payoffs is:

$$0 + 3\delta + 0\delta^2 + 3\delta^3 + ... =$$

$$= 0[1 + \delta^2 + \delta^4 + ...] + 3\delta[1 + \delta^2 + \delta^4 + ...]$$

$$= \frac{3\delta}{1 - \delta^2}$$

Partial cooperation

1. Incentives to cheat for player 2 in an *odd-numbered* period:
    1. By cheating player 2 obtains an payoff of 5 (instantaneous gain of 2), but
    2. His defection is detected, and punished with (D,D) thereafter. This gives him a payoff of 1 for every subsequent round, or $\frac{\delta}{1-\delta}$ thereafter.
    3. Instead, by respecting the modified GTS, he obtains a payoff of 3 during this period (odd-numbered period, when they play (C,C)).
        1. In addition, the discounted stream of payoffs from the next period (an even-numbered period) thereafter is $\frac{3\delta^2}{1-\delta^2}$.
    4. Hence, player 2 prefers to stick to this modified GTS if

$$3 + \frac{3\delta^2}{1-\delta^2} \geq 5 + \frac{\delta}{1-\delta} \iff \delta \geq \frac{1+\sqrt{33}}{8} \simeq 0.84 \qquad (1)$$

Partial cooperation

1. Incentives to cheat for player 2 in an *even-numbered* period:
   1. By cheating player 2 obtains an payoff of 1 (instantaneous loss of 2), moreover...
   2. His defection is detected, and punished with (D,D) thereafter. This gives him a payoff of 1 for every subsequent round, or $\frac{\delta}{1-\delta}$ thereafter.
   3. Instead, by respecting the modified GTS, he obtains a payoff of 0 during this period (even-numbered period, when they play (D,C) and he is player 2). In addition, the discounted stream of payoffs from the next period (an odd-numbered period) thereafter is $\frac{3\delta}{1-\delta^2}$.
   4. Hence, player 2 prefers to stick to this modified GTS if

$$0 + \frac{3\delta}{1-\delta^2} \geq 1 + \frac{\delta}{1-\delta} \iff \delta \geq \frac{1}{2} \tag{2}$$

Partial cooperation

- And because $\frac{1+\sqrt{33}}{8} \simeq 0.84$ (for odd-numbered period) is larger than $\frac{-1+\sqrt{3}}{2} \simeq 0.37$ (for even-numbered period),
- Thus, player 2 cooperates in any period (odd or even) as long as $\delta \geq \frac{1+\sqrt{33}}{8} \simeq 0.84$.

Partial cooperation

1. **Player 1:** If, after a history of cooperation, he cooperates in an odd-numbered period, where he plays C, in a (C,C) outcome, his payoff stream is

$$3 + 5\delta + 3\delta^2 + 5\delta^3 + ... =$$

$$\begin{aligned} &= 3[1 + \delta^2 + \delta^4 + ...] + 5\delta[1 + \delta^2 + \delta^4 + ...] \\ &= \frac{3}{1-\delta^2} + \frac{5\delta}{1-\delta^2} \end{aligned}$$

If, instead, he deviates to D, yielding (D,C) today, his payoff increases to 5 today but decreases to 1 thereafter (punishment), yielding $5 + \frac{\delta}{1-\delta}$.

Partial cooperation

1. Therefore, player 1 continues cooperation in every odd-numbered period if

$$\frac{3}{1-\delta^2} + \frac{5\delta}{1-\delta^2} \geq 5 + \frac{\delta}{1-\delta}$$

which simplifies to

$$\frac{\delta^2(5+3\delta) - \delta(1+\delta)}{1-\delta^2} \geq 2$$

or $5\delta + 2\delta^2 - \delta \geq 2 - 2\delta^2$. Solving for δ, we find two roots, but only $\delta \geq \frac{-1+\sqrt{3}}{2} \simeq 0.36$ lies in the admissible range of $\delta \in [0,1]$. (The other root is $\delta \leq -1.36$.)

Partial cooperation

1. **Player 1:** If, after a history of cooperation, he cooperates in an even-numbered period, where he plays D, in a (D,C) outcome, his payoff stream is

$$5 + 3\delta + 5\delta^2 + 3\delta^3 + ... =$$

$$\begin{aligned} &= 5[1 + \delta^2 + \delta^4 + ...] + 3\delta[1 + \delta^2 + \delta^4 + ...] \\ &= \frac{5}{1-\delta^2} + \frac{3\delta}{1-\delta^2} \end{aligned}$$

If, instead, he deviates to C, yielding (C,C) today, his payoff would actually decrease to 3 today, and further decrease to 1 thereafter (punishment), yielding $3 + \frac{\delta}{1-\delta}$.

Partial cooperation

1. Therefore, player 1 continues cooperation in every even-numbered period if

$$\frac{5}{1-\delta^2} + \frac{3\delta}{1-\delta^2} \geq 3 + \frac{\delta}{1-\delta}$$

which simplifies to

$$\delta(3+5\delta) - \delta(1+\delta) \geq 2\delta^2 - 2$$

or $\delta^2 + \delta + 1 \geq 0$. This equation lies above zero for all $\delta \in [0, 1]$, starting at a height of 1, and reaching a height of 3, so it holds for all δ.

2. In summary, we only need condition $\delta \geq \frac{-1+\sqrt{3}}{2} \simeq 0.36$ for player 1 to cooperate in this GTS.

The Folk Theorem

- Therefore, any payoff vector within the diamond of FIR payoffs can be supported as a SPNE of the game for sufficiently high values of δ.
    - Advantages and disadvantages.

Advantages and Disadvantages of the Folk Theorem:

- **Good:** efficiency is possible
  - Recall that any improvement from (D,D) in the PD game constitutes a Pareto superior outcome.
- **Bad:** lack of predictive power
  - Anything goes!
  - Any payoff pair within the FIR shaded area can be supported as a SPNE of the infinitely repeated game.

Incentives to cooperate in the PD game:

- Our results depend on the individual incentives to cheat and cooperate.
- When the difference between the payoffs from cooperate and not cooperate is sufficiently large, then δ doesn't have to be so high in order to support cooperation.
  - Intuitively, players have stronger per-period incentives to cooperate (mathematically, the minimal cutoff value of δ that sustains cooperation will decrease).
  - Let's show this result more formally.

Incentives to cooperate:

- Consider the following simultaneous-move game

		Player 2	
		Coop	*Defect*
Player 1	*Coop*	a, a	$c, \underline{b}$
	Defect	$\underline{b}, c$	$\underline{d}, \underline{d}$

1. To make this a Prisoner's Dilemma game, we must have that D, "defect," is strictly dominant for both players.
2. That is, D must provide every player a higher payoff, both:
   1. when the other player chooses C, "cooperate" (given that $b > a$), or
   2. when the other player defects as well (since $d > c$).

Incentives to cooperate:

- Hence, the unique NE of the unrepeated game is (D,D).
- What if we repeat the game infinitely many times?
  - We can then design a standard GTS to sustain cooperation.

In the infinitely repeated game...

- At any period t, my payoff from cooperating is...

$$a + \delta a + \delta^2 a + \ldots = \frac{1}{1-\delta} a$$

- If, instead, I deviate my payoff becomes...

$$\underbrace{b}_{\text{current gain}} + \underbrace{\delta d + \delta^2 d + \ldots}_{\text{future loss}} = b + \frac{\delta}{1-\delta} d$$

In the infinitely repeated game...

- Hence, players cooperate if

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

Rearranging,

$$a \geq b(1-\delta) + \delta d, \text{ or } \delta \geq \frac{b-a}{b-d}$$

Intuition behind this cutoff for delta...

- $(b - a)$ measures the instantaneous gain you obtain by deviating from cooperation to defection. (more temptation to cheat!)
- $(b - d)$ measures the loss you will suffer thereafter as a consequence of your deviation.

Intuition behind this cutoff for delta...

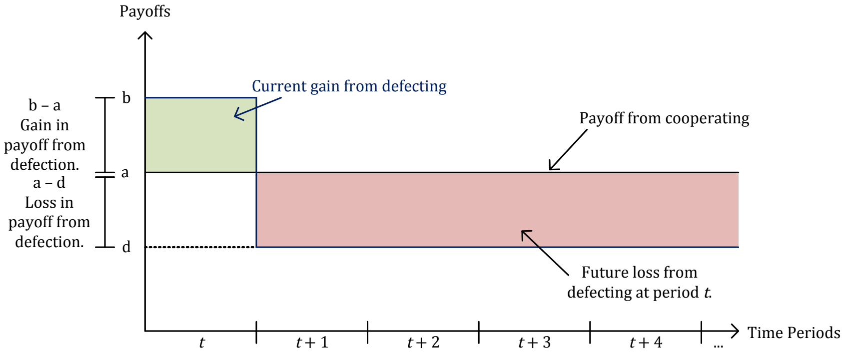

Intuition behind this cutoff for delta...

- Therefore,
  - An increase in $(b-a)$ or a decrease in $(b-d)$ implies an increase in $\bar{\delta}=\frac{b-a}{b-d}$, i.e., cooperation is more difficult to support.
  - A decrease in $(b-a)$ or an increase in $(b-d)$ implies a decrease in $\bar{\delta}=\frac{b-a}{b-d}$, i.e., cooperation is easier to support.

Intuition behind this cutoff for delta...

- When $(b-a)\uparrow$ or $(b-d)\downarrow$ the cutoff $\bar{\delta}=\frac{b-a}{b-d}$ becomes closer to 1.

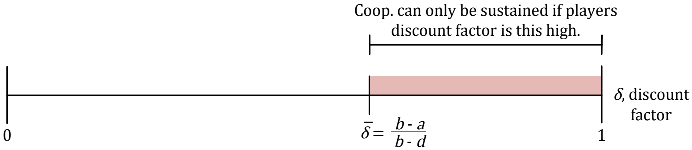

- When $(b-a)\downarrow$ or $(b-d)\uparrow$ the cutoff $\bar{\delta}=\frac{b-a}{b-d}$ becomes closer to zero.

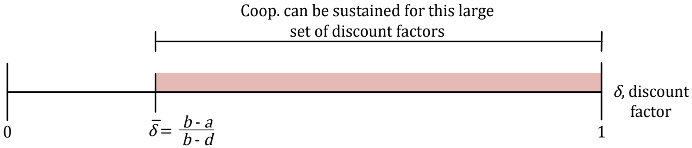

What if we have 2 NE in the stage game...

- Note that the games analyzed so far had a unique NE in the stage (unrepeated) game.
- What if the stage game has two or more NE?

What if we have 2 NE in the stage game...

- Consider the following stage game:

		Player 2		
		x	y	z
	x	5, 5	2, $\underline{7}$	1, 3
Player 1	y	$\underline{7}$, 2	$\underline{3}$, $\underline{3}$	0, 1
	z	3, 1	1, 0	$\underline{2}$, $\underline{2}$

- There are indeed 2 psNE in the stage game: (y, y) and (z, z).
- Outcome (x, x) is the socially efficient outcome, since the sum of both players' payoffs is maximized.
  - How can we coordinate to play (x, x) in the infinitely repeated game? Using a "modified" GTS.

A modified grim-trigger strategy:

1. Period $t = 1$: choose x ("Cooperate")
2. Period $t > 1$: choose x as long as no player has ever chosen y,
    1. If y is chosen by some player, then revert to z forever.
        - (This implies a big punishment, since payoffs decrease to those in the worst NE of the unrepeated game $2, rather than those in the best NE of the unrepeated game, $3.)

- **Note:** If the other player deviates from x to z while I was cooperating in x, I don't revert to z (I do so only after observing he played y).
    - Later on, we will see a more restrictive GTS, whereby I revert to z after observing any deviation from the cooperative x, which can also be sustained as a SPNE.

A modified grim-trigger strategy:

- At any period t in which the history of play was cooperative, my payoffs from sticking to the cooperative GTS (selecting x) are

$$5 + \delta 5 + \delta^2 5 + ... = \frac{1}{1-\delta} 5$$

- If, instead, I deviate towards my "best deviation" (which is y), my payoffs are

$$\underbrace{7}_{\text{current gain}} + \underbrace{\delta 2 + \delta^2 2 + ...}_{\text{Punishment thereafter}} = 7 + \frac{\delta}{1-\delta} 2$$

- **One second!** Shouldn't it be

$$7 + \delta 0 + \delta^2 2 + \delta^3 + ... = 7 + \frac{\delta^2}{1-\delta} 2$$

  - *No.* My deviation to y in any period t, also triggers my own reversion towards z in period $t+1$ and thereafter.

A modified grim-trigger strategy:

- Hence, every player compares the above stream of payoffs, and choose to keep cooperating if

$$\frac{1}{1-\delta}5 \geq 7 + \frac{\delta}{1-\delta}2$$

- Rearranging...

$$5 \geq 7(1-\delta) + 2\delta, \text{ or } \delta \geq \frac{2}{5}$$

ANOTHER modified grim-trigger strategy:

- What if the modified GTS was more restrictive, specifying that players revert to z as soon as they observe any deviation from the cooperative outcome, x.
  - That is, I revert to z (the "worst" NE of the unrepeated game) as soon as you select either y or z.
  - In our previous "modified GTS" I only reverted to z if you deviated to y.
- That is, the GTS would be of the following kind:
  1. At $t = 1$, choose x (i.e., start cooperating).
  2. At $t > 1$, continue choosing x if all players previously selected x. Otherwise, deviate to z thereafter.

ANOTHER modified grim-trigger strategy:

- At any period t in which the previous history of play is cooperative, my payoffs from sticking to the cooperative GTS (selecting x) are

$$5 + \delta 5 + \delta^2 5 + ... = \frac{1}{1-\delta}5$$

- If, instead, I deviate towards my "best deviation" (which is y), my payoffs are

$$\underbrace{7}_{\text{current gain}} + \underbrace{\delta 2 + \delta^2 2 + ...}_{\text{Punishment thereafter}} = 7 + \frac{\delta}{1-\delta}2$$

- Hence, cooperation in x can be sustained as SPNE of the infinitely-repeated game as long as

$$\frac{1}{1-\delta}5 \geq 7 + \frac{\delta}{1-\delta}2, \text{ or } \delta \geq \frac{2}{5}$$

(Same cutoff as with the previous "modified GTS").

Summary:

- When the unrepeated version of the game has more than one NE, we can still support cooperative outcomes as SPNE of the infinitely repeated game whereby all players experience an increase in their payoffs.
- **Usual trick:** make the punishments really nasty!
  - For instance, the GTS can specify that we start cooperating...
  - but we will both revert to the "worst" NE (the NE with the lowest payoffs in the unrepeated game) if any player deviates from cooperation.
- The analysis is very similar to that of unrepeated games with a unique NE.

Many things still to come...

Note that so far we have made several simplifying assumptions...

- **Observability of defection:** When defection is more difficult to observe, I have more incentives to cheat.
  - Then, δ needs to be higher if we want to support cooperation.
- **Starting of punishments:** When the punishment is only triggered after two (or more) periods of defection, then the short run benefits from defecting become relatively larger.
  - Then, δ needs to be higher if we want to support cooperation.
- **Thereafter punishments:** Punishing you also reduces my own payoffs, why not go back to our cooperative agreement after you are disciplined?

Many things still to come...

- We will discuss many of these extensions in the next few days.
- **But let's finish with some fun!**
- Let's examine how UCLA undergraduates actually behaved when asked to play the PD game in an experimental lab:
  - One period (unrepeated game)
  - Two to four periods (finitely repeated game)
  - Infinite periods (How can we operationalize that in an experiment?? chaining them to their desks?)

Recall our general interpretation of the discount factor

- δ represents players' discounting of future payoffs, but also...
- The probability that I encounter my opponent in the future, or
  - Probability that the game continues one more round.
- This can help us operationalize the infinitely repeated PD game in the experimental lab...
  - by simply asking players to roll a die at the end of each round to determine whether the game continues,
  - i.e., probability of continuation p (equivalent to δ) can be, for instance, 50%.

Experimental evidence for the PD game

- Consider the following PD game presented to 390 UCLA undergraduates...

		Player 2	
		Mean	*Nice*
Player 1	*Mean*	2, 2	4, 1
	Nice	1, 4	3, 3

where "Mean" is the equivalent of "defect" and "Nice" is the equivalent of "cooperate" in our previous examples.

Experimental evidence for the PD game

- The PD game provides us very sharp testable predictions:
    1. If the PD game is played once, players will choose "mean."
    2. If the PD game is played a *finite* number of times, players will choose "mean" in every period.

Experimental evidence for the PD game

- More testable predictions from the PD game...
  1. If the PD game is played an indefinite (or infinite) number of times, players are likely to choose "nice" some of the time.
     1. Why "some of the time"? Recall that the folk theorem allows for us to cooperate all the time, yielding a payoff in the northeast corner of the FIR diamond, or...
     2. cooperate every other period, yielding payoffs in the interior of the FIR diamond, e.g., at the boundary but not at the northeast corner, as in the partially cooperative GTS we described
  2. If the PD game is played an indefinite (or infinite) number of times, players are more likely to choose "nice" when the probability of continuation (or the discount factor) is higher.

Experimental evidence for the PD game

- Frequency of cooperative play in the PD game:

Not zero, but close.

		1	2	3	4	5	6	7	8	9	10	11	12
Unrepeated	One-shot	9%											
Finitely Repeated	$T = 2$	13%	7%										
	$T = 4$	35%	22%	19%	11%								
Infinitely Repeated with $p \simeq \delta$	$p = \frac{1}{2}$	31%	26%	20%	13%	13%							
	$p = \frac{3}{4}$	46%	41%	39%	35%	33%	27%	25%	26%	29%	26%	32%	31%

- In the last round of the finitely repeated game, players play "as if" they were in an unrepeated (one-shot) game.
- They are not capable of understanding the SPNE of the game in the finitely repeated game (second and third row), but...
  - Their rates of cooperation increase in p ($\simeq \delta$), as illustrated in the last two rows.

Experimental evidence for the PD game

- A common criticism to experiments is that stakes are too low to encourage real competition.
    - e.g., average payoff was about $19 per student at UCLA.
- What if we increase the stakes to a few thousand dollars?
    - Is cooperation less supported than in experiments, as theory would predict?
    - Economists found a natural experiment: "Friend or Foe?" TV show.
    - Check at YouTube http://www.youtube.com/watch?v=SBgalflgx2U&feature=related

Friend or Foe?

- Two people initially work together to answer trivia questions.
  - Answering questions correctly results in contributions of thousands of dollars to a trust fund.

Friend or Foe?

- Afterwards, players are separated and asked to simultaneously and independently choose "Friend" (i.e., evenly share the trust fund) or "Foe" (i.e., get it all if the other player is willing to share), with these resulting payoffs...

		Player 2: *Foe*	Player 2: *Friend*
Player 1	*Foe*	<u>0, 0</u>	<u>V, 0</u>
Player 1	*Friend*	0, V	$\frac{V}{2}$, $\frac{V}{2}$

- Note that choosing "Foe" is a dominant strategy for each player, although it is weakly (not strictly) dominant. [Close enough to the PD]

Friend or Foe?

- A lot at stake!

	1st stage	2nd stage: Play Fried or Foe	
	Trivia Earnings	Cooperation Rate	Take-Home Earnings
Overall	$3,705	50%	$1,455
Men	$4,247	45% <	$1,834
Women	$3,183	56%	$1,088
White	$3,957	53% >	$1,417
Nonwhite	$2,825	42%	$1,587
Young (< age 31)	$3,603	41% <	$1,592
Mature (≥ age 31)	$3,839	63%	$1,276

- But the details in these results are even more intriguing!

Friend or Foe?

	Men	Women
Men	48%, 48%	43%, 55%
	$>$	$\simeq$
Women	55%, 43%	56%, 56%
	Young	**Mature**
Young	40%, 40%	42%, 63%
	$<$	$=$
Mature	63%, 42%	63%, 63%
	White	**Nonwhite**
White	51%, 51%	58%, 44%
	$<$	$>$
Nonwhite	44%, 58%	25%, 25%

Interpretation:

1. **Gender:**
    1. Men are more cooperative when his opponent is also a man, than when she is a woman.
    2. Women, in contrast, are as cooperative with men as they are with women.
2. **Age group:**
    1. Young contestants are slightly more cooperative with mature than with young contestants.
    2. Mature contestants are as cooperative with other mature contestants as they are with young opponents.
3. **Race:**
    1. White contestants are more cooperative with a non-white contestant than with another white contestant, but...
    2. Non-white contestants are less cooperative with another non-white contestant than with a white contestant.