Chapter 7: Repeated Games

Game Theory:

An Introduction with Step-by-Step Examples

by Ana Espinola-Arredondo and Felix Muñoz-Garcia

Introduction

- Situations in which players interact in one of the strategic setting and they play the game for several rounds, which are known as “repeated games”
- Interested in identifying if the game’s repetition provides players with more incentives to cooperate

Prisoner’s Dilemma

- We first study the finitely-repeated version of this game
- Extend players’ interaction to an infinitely repeated game
- We then apply the above tools to different settings, such as collusion in oligopoly models, where every firm chooses its output level in each period of interaction, and to stage games with more than one NE

Repeating the Game twice

		Player 2: Confess	Player 2: Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.1 The Prisoner's Dilemma Game

- Let us repeat the game twice and find its SPE since players now interact in a sequential-move game:
  - In the first stage, every player i simultaneously and independently chooses whether to *Confess* or *Not Confess*
  - In the second stage, observing the outcome of the first stage, every player i selects again simultaneously and independently whether to *Confess* or *Not Confess*
  - This sequential-game is depicted in Figure 7.1 (next slide).

Twice-repeated Prisoner's Dilemma Game

- The game tree includes information sets in the first and second stage.
- At the beginning of second stage, however, players perfectly observe the strategy profile played in the first stage (no information sets there).

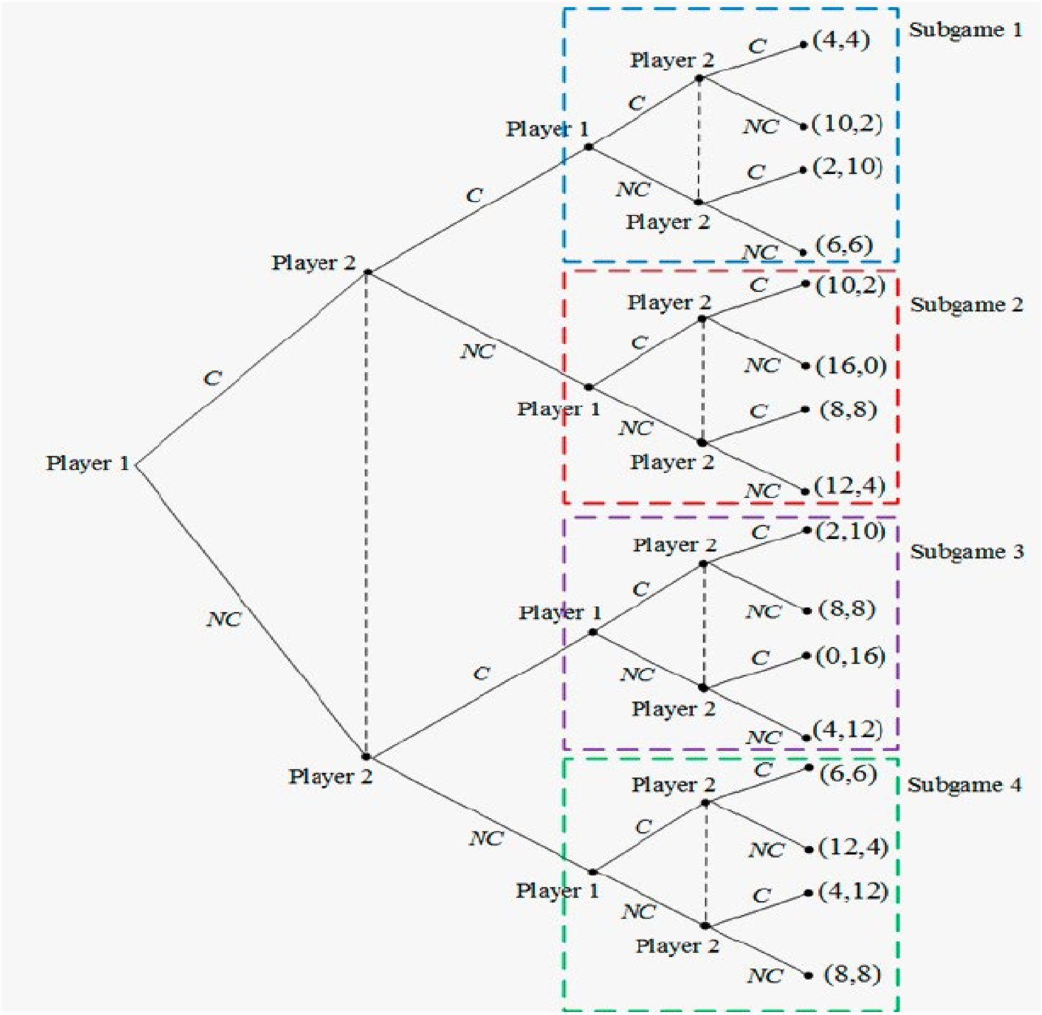

Figure 7.1. Twice-repeated prisoner's dilemma game.

Twice-repeated Prisoner's Dilemma Game

- Payoff pairs in the terminal nodes are just a sum of payoffs in the first and second stage, assuming no discounting.
- For instance, if (C, C) is played in the first stage, but (NC, C) occurs in the second stage:
  - Player 1 earns $2 + 0 = 2$, while
  - Player 2 earns $2 + 8 = 10$.
- A similar argument applies to other payoffs at the bottom of the tree. (Practice!)

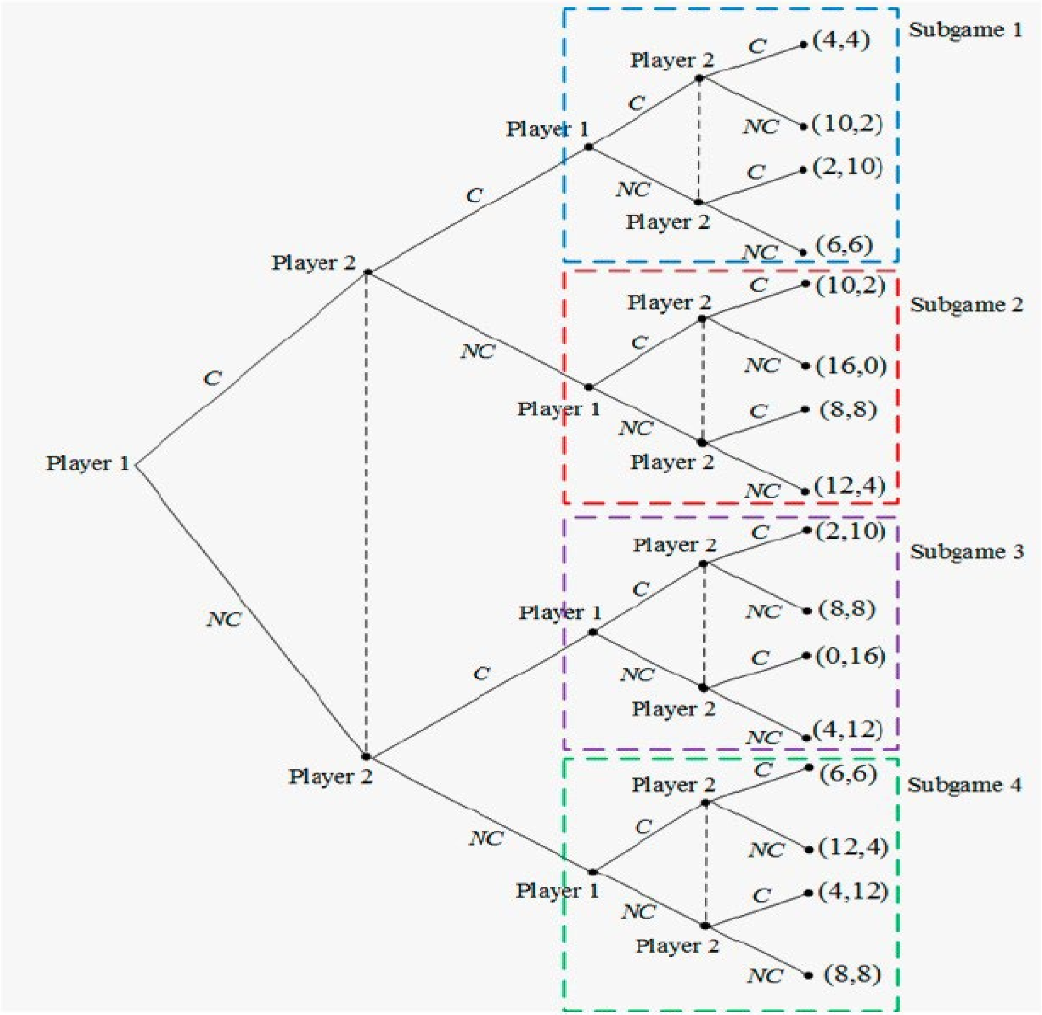

Figure 7.1. Twice-repeated prisoner's dilemma game.

Twice-repeated Prisoner’s Dilemma Game

- There are five subgames:
  1. One initiated after players choose (C, C) in the first stage.
  2. Another initiated after (NC, C)
  3. Another initiated after (C, NC)
  4. Another initiated after (NC, NC)
  5. And the game as a whole
- Operating by backward induction, we solve each of these subgames, starting with subgames 1-4.
- Since these subgames are simultaneous-move games, we can solve each of them by transforming each to its matrix form.

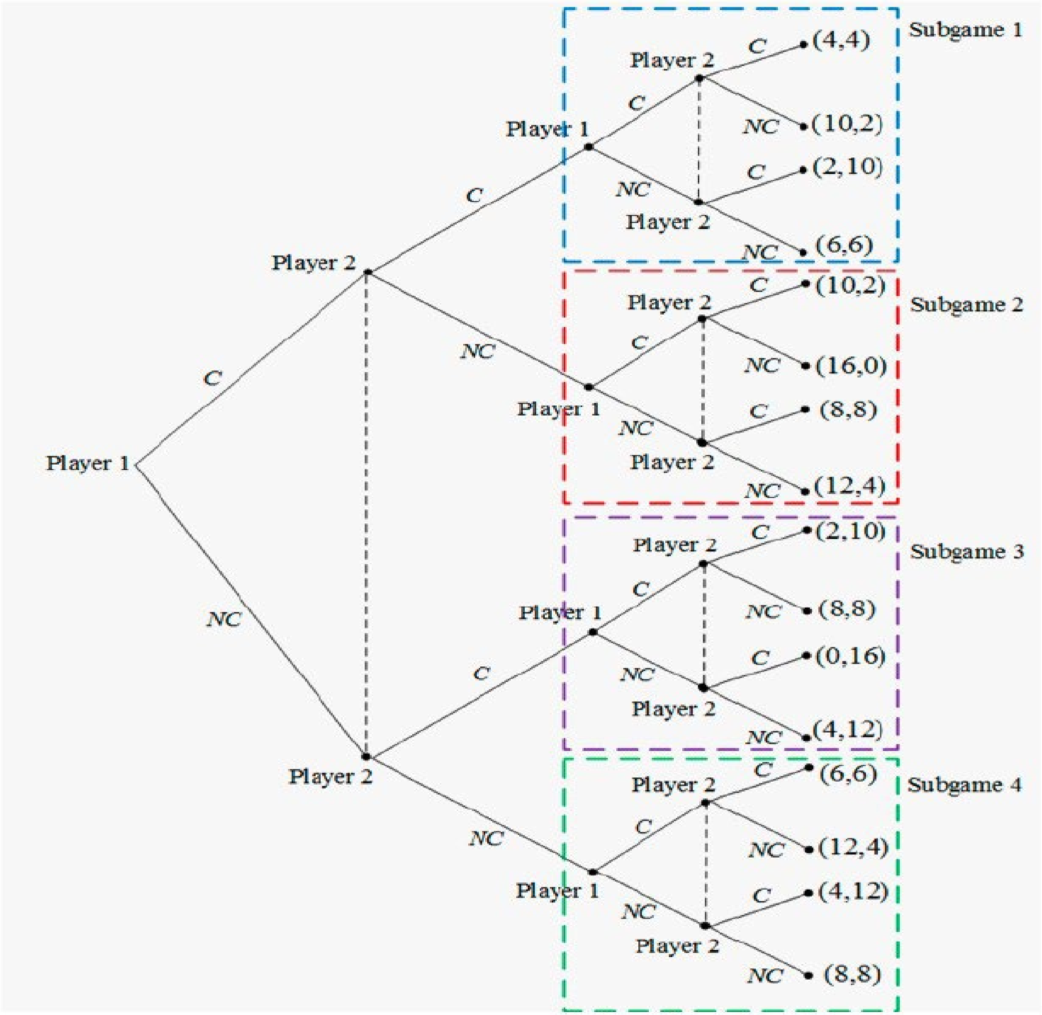

Figure 7.1. Twice-repeated prisoner’s dilemma game.

Twice-repeated Prisoner’s Dilemma Game - Subgames for Second Stage

Player 1 \ Player 2	Confess	Not Confess
Confess	4,4	10,2
Not Confess	2,10	6,6

Matrix 7.2a Twice-repeated Prisoner’s Dilemma game – Subgame 1

Player 1 \ Player 2	Confess	Not Confess
Confess	10,2	16,0
Not Confess	8,8	12,4

Matrix 7.2b Twice-repeated Prisoner’s Dilemma game – Subgame 2

Player 1 \ Player 2	Confess	Not Confess
Confess	2,10	8,8
Not Confess	0,16	4,12

Matrix 7.2c Twice-repeated Prisoner’s Dilemma game – Subgame 3

Player 1 \ Player 2	Confess	Not Confess
Confess	6,6	12,4
Not Confess	4,12	8,8

Matrix 7.2d Twice-repeated Prisoner’s Dilemma game – Subgame 4

Twice-repeated Prisoner's Dilemma Game – First Stage

- **Summary**
  - Every player $i = \{1,2\}$ chooses to play $Confess$ in the first stage, and
  - Players $Confess$ in the second stage regardless of the strategy profile played in the first stage

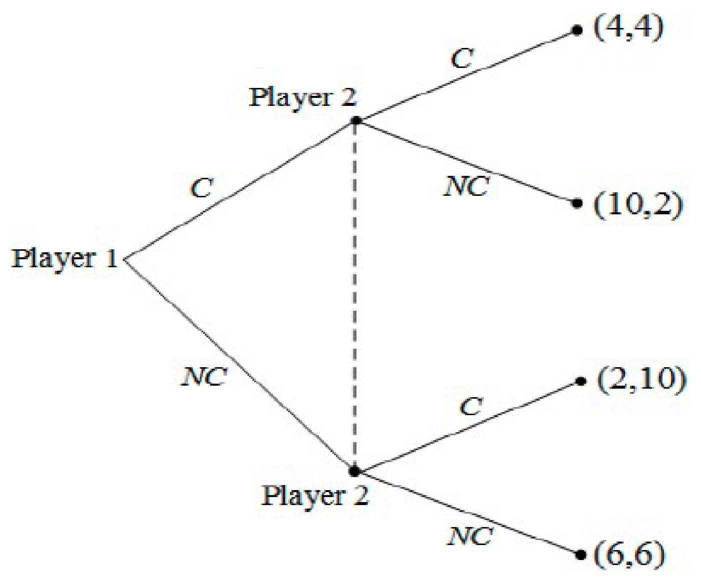

Figure 7.2. Twice-repeated prisoner's dilemma gama - First stage.

		Player 2	
		Confess	Not Confess
Player 1	Confess	4,4	10,2
	Not Confess	2,10	6,6

Matrix 7.2e Twice-repeated Prisoner's Dilemma game – Game as whole

Summary

- Intuitively, players anticipate that they will both play $Confess$ in the last stage independently of how cooperative or uncooperative they were in the previous stage
- Informally, every player thinks:
  - "It does not matter what I do today, tomorrow we will all confess.
  - I can then treat tomorrow's game as independent from today's play, as if I played two separate Prisoner's Dilemma games."
- In summary, repeating the game twice did not help us sustain cooperation in the SPE.
- Can cooperation emerge if we repeat the game for more periods, generally, for $T \geq 2$ periods?

Repeating the game $T \geq 2$ times

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2+P_1^{T-1}}, \underline{2+P_2^{T-1}}$	$\underline{8+P_1^{T-1}}, 0+P_2^{T-1}$
	Not Confess	$0+P_1^{T-1}, \underline{8+P_2^{T-1}}$	$4+P_1^{T-1}, 4+P_2^{T-1}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

- Matrix 7.3 represents the game that players face when interacting in the last round (period T)
- Notation:
  - P_1^{T-1} denotes the discounted sum of all payoffs that player 1 earned in the previous $T-1$ periods.
  - P_2^{T-1} is a similar expression for player 2.
- This presentation helps us describe the subgame that players face at period T for any previous history of play.
  - That is, for any strategy profiles occurring in previous periods.
  - *Intuitively*: regardless of how we played the game in previous periods.

Repeating the game $T \geq 2$ times

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-1}}, \underline{2 + P_2^{T-1}}$	$\underline{8 + P_1^{T-1}}, 0 + P_2^{T-1}$
	Not Confess	$0 + P_1^{T-1}, \underline{8 + P_2^{T-1}}$	$4 + P_1^{T-1}, 4 + P_2^{T-1}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

- To illustrate the generality of this matrix representation, evaluate terms P_1^{T-1} and P_2^{T-1} in the special case where the game is only repeated $T = 2$ times.
- In this context:
  - P_1^{T-1} captures the payoff that player 1 earned in the first period of interaction.
  - And similarly, with P_2^{T-1} for player 2.
- For instance, if (C, C) emerged in the first stage of the game,
  - $P_1^{T-1} = P_2^{T-1} = 2$.
  - While if (C, NC) occurred, then $P_1^{T-1} = 8$ and $P_2^{T-1} = 0$.

Repeating the game $T \geq 2$ times

- If the Prisoner's Dilemma can produce four different strategy profiles in each period, there are 4^{T-1} nodes at the beginning of period T.
- In the twice-repeated Prisoner's Dilemma game, for instance, there are $4^{2-1} = 4$ nodes at the beginning of the second period, as shown in the previous tree.
- More generally, if the unrepeated version of the game has
    - $k > 2$, different strategy profiles (cells) and
    - the game is repeated T times,
- Then, there are:
    - k^{T-1} nodes at the beginning of period T,
    - only k^{T-2} nodes at the beginning of period $T - 1$,
    - and similarly in previous periods…
    - with just $k^{T-T} = k^0 = 1$ node at the beginning of the game (first period of interaction).

Iterative Method: Period T

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2+P_1^{T-1}}, \underline{2+P_2^{T-1}}$	$\underline{8+P_1^{T-1}}, 0+P_2^{T-1}$
	Not Confess	$0+P_1^{T-1}, \underline{8+P_2^{T-1}}$	$4+P_1^{T-1}, 4+P_2^{T-1}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

Period T:

- Underlining best response payoffs, we can see that (C, C) is a NE in the T-period subgame in Matrix 7.3.
- Player 1's best response:
  - When Player 2 chooses *Confess,* player 1's best response is to *Confess* because $2 + P_1^{T-1} > 0 + P_1^{T-1}$ simplifies to $2 > 0$.
  - When Player 2 selects $Not\ Confess$, player 1's best response is *Confess* since $8 + P_1^{T-1} > 4 + P_1^{T-1}$ simplifies to $8 > 4$.

Iterative Method: Period T

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2+P_1^{T-1}}, \underline{2+P_2^{T-1}}$	$\underline{8+P_1^{T-1}}, 0+P_2^{T-1}$
	Not Confess	$0+P_1^{T-1}, \underline{8+P_2^{T-1}}$	$4+P_1^{T-1}, 4+P_2^{T-1}$

Matrix 7.3. Finitely-repeated Prisoner's Dilemma Game – Subgames at period T

Period T:

- Note that:
  - Players' best responses in period T don't depend on players' payoffs in previous periods, as captured by P_1^{T-1} and P_2^{T-1}.
  - As if they were sunk, or they couldn't be changed at this point.
  - In other words, player 1 treats P_1^{T-1} as a constant; while player 2 treats P_2^{T-1} as a constant.
  - This implies that every player finds:
    - *Confess* to be a best response to her rival's strategy *regardless* of the previous history of play.

Iterative Method: Period $T-1$

- **Period $\boldsymbol{T-1}$.**

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2+P_1^{T-2}}, \underline{2+P_2^{T-2}}$	$\underline{8+P_1^{T-2}}, 0+P_2^{T-2}$
	Not Confess	$0+P_1^{T-2}, \underline{8+P_2^{T-2}}$	$4+P_1^{T-2}, 4+P_2^{T-2}$

Matrix 7.4. Finitely-repeated Prisoner's Dilemma Game – Subgames at period $T-1$

- (C, C) is again the unique NE.
- The intuition is identical to that in the twice-repeated game:
  - Players anticipate that everyone will play *Confess* in the last round of interaction (period T).
  - Therefore, player have no incentives to cooperate today (in period $T-1$) since misbehavior will not be disciplined in the following period.

Iterative Method: Period $T - 2$

- **Period $T - 1$**.

		Player 2	
		Confess	Not Confess
Player 1	Confess	$\underline{2 + P_1^{T-3}}, \underline{2 + P_2^{T-3}}$	$\underline{8 + P_1^{T-3}}, 0 + P_2^{T-3}$
	Not Confess	$0 + P_1^{T-3}, \underline{8 + P_2^{T-3}}$	$4 + P_1^{T-3}, 4 + P_2^{T-3}$

Matrix 7.4. Finitely-repeated Prisoner's Dilemma Game – Subgames at period $T - 1$

- The same argument follows for period $T - 2$, as illustrated in the above matrix.
- Same argument applies to period $T - 3$, and so on.
- Extending that reasoning to all previous periods, we find that (C, C) is the unique NE:
  - in every subgame and
  - in every period.

Repeating the Game infinitely many times

- We could not support cooperation by repeating the game twice or $T \geq 2$ times…
  - But our above results point us to why uncooperative outcomes emerge:
  - The presence of a terminal period when players know that the game ends
- At that last stage, both players:
  - Behave as in the unrepeated version of the game,
  - Regardless of how the game was played in previous rounds.
- Using backward induction, they extend this behavior to all previous interactions.

Repeating the Game infinitely many times

- Therefore, if we seek to sustain cooperation:
  - Players cannot know with certainty when the game will end.
  - They interact today, at time t, and will keep playing with each other tomorrow, at time t+1, with a probability p, where $p \in [0, 1]$.
  - For instance, if $p = 0.7$, the chances of interacting for more than 20 rounds are less than 1 percent.
  - Probability of interacting converges fast to zero, but it is always strictly positive.

Uncooperative outcome

- We first show that (C,C) in every period can still be sustained as one of the SPEs in the infinitely-repeated game.
- Note that if a player chooses C at every period t, anticipating that his rival will choose C as well, he obtains a sequence of 2, with discounted value

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

Recall that

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

Uncooperative outcome (C, C)

- If, instead, she unilaterally deviates to $Not\ Confess$ for one period, he earns

$$\underbrace{0}_{Not\ Confess} + \underbrace{2\delta + 2\delta^2 + \cdots}_{Confess} = 0 + 2\frac{\delta}{1-\delta}$$

Hence, this player earns a higher stream of payoffs choosing $Confess$ than $Not\ Confess$ because

$$2\frac{1}{1-\delta} \geq 2\frac{\delta}{1-\delta}$$

simplifies to $1 \geq \delta$ which holds given that $\delta \in (0,1)$ by assumption.

Cooperative Outcome: Grim-Trigger Strategy

- Can we sustain cooperation as the SPE of the infinitely-repeated game?
- For that, we need players to:
  - Cooperate and keep cooperating as long as every player cooperated in previous period, but…
  - Otherwise, threat to move the game to a "grim" outcome where players earn lower payoffs than by cooperating.
- The following Grim-Trigger Strategy (GTS) helps us achieve exactly that:
  1. In period $t = 1$, choose $Not\ Confess$
  2. In every period $t \geq 2$,
     a. Keep choosing $Not\ Confess$ if every player chose $Not\ Confess$ in all previous periods, or
     b. Choose $Confess$ thereafter for any other history of play (i.e., if either player chose Confess in any period)

Cooperative Outcome: Grim-Trigger Strategy

- The name of the GTS strategy should be clear at this point:
  - Any deviation triggers a grim punishment thereafter, where every player reverts to the NE of the unrepeated version of the game (often known as "Nash reversion").
- To show that the GTS is a SPE of the infinitely-repeated game, we need to demonstrate that it is an optimal strategy:
  - for *every* player, and
  - at *every* subgame where they are called to move.
- This means that players find the GTS to be optimal:
  - at *any* point t, and
  - after *any* previous history of play.
- A daunting task? No, we only need to examine two histories of play.

Case (1). No Cheating History

- At any period t (both $t = 1$ and $t \geq 2$), if no previous history of cheating occurs, the GTS prescribes that:
  - Every player keeps cooperating, which yields a payoff of 4 for every player i, entailing a payoff stream of

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

Case (1). No Cheating History

- If, instead, player i chooses $Confess$ today (which we interpret as her deviating from the above GTS, or more informally as "cheating"), her current payoff increases from 4 to 8.
  - This is a <u>*unilateral deviation*</u> as only player i chooses $Confess$ while her opponent still plays according to the GTS.
- Therefore, player i's discounted payoff stream from cheating is

$$\underbrace{8}_{Deviation\ to\ Confess} + \underbrace{2\delta + 2\delta^2 + \cdots}_{Punishment\ thereafter} = 8 + 2\frac{\delta}{1-\delta}$$

Case (1). No Cheating History

- ***Comparison***:
- Then, after a history with no previous cheating episodes, every player chooses to cooperate (player $Not\ Confess$) if

$$4\frac{1}{1-\delta} \geq 8 + 2\frac{\delta}{1-\delta}$$

- Multiplying both sides by $1-\delta$, yields

$$4 \geq 8(1-\delta) + 2\delta$$

or $4 \geq 8 - 6\delta$, which solving for δ yields

$$\delta \geq \frac{2}{3}$$

Case (2). Some Cheating History

- At period t, if one or both players cheated in previous periods, then the GTS dictates that:
  - Every player responds with $Confess$ thereafter, earning a discounted stream of payoffs

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

- Player i could, instead, unilaterally deviate to $Not\ Confess$, while his rival plays $Confess$ as part of the punishment in the GTS.
  - You may suspect that such deviation is not profitable:
    - Player i is the only one choosing $Not\ Confess$, decreasing his payoff in that period, rather than increasing it.
    - And his deviation does not change his rival's behavior in subsequent stages.
  - That suspicion is correct: Player i's discounted stream of payoffs from this deviation is

$$\underbrace{0}_{Deviation\ to\ Not\ Confess} + \underbrace{2\delta + 2\delta^2 + \cdots}_{Punishment\ thereafter}$$

Case (2). Some Cheating History

- In period t, players observe (NC, C) or (C, NC) being played, implying that a deviation from the fully cooperative outcome (NC, NC) occurred, triggering an infinite punishment of (C, C) thereafter, with a payoff of 2 to every player. This payoff stream simplifies to

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

- ***Comparison:*** After a history of cheating, every player i prefers to implement the punishment prescribed by the GTS if

$$2\frac{1}{1-\delta} \geq 2\frac{\delta}{1-\delta} \Rightarrow 2 \geq 2\delta$$

which holds for $\delta \in (0,1)$

Summary

- Combining our results from cases (1) and (2), we only found one condition restricting the value of δ for us to sustain the GTS as a SPE of the infinitely-repeated Prisoner's Dilemma game:

$$\delta \geq \frac{2}{3}$$

- Intuitively:
  - When players assign a sufficiently high weight to future payoffs…
  - They start cooperating in the first period and keep cooperating in all subsequent periods, yielding outcome (C, C) in every round of interaction.

Summary

- Figure 7.3 illustrates the trade-off that every player experiences when deciding whether to cooperate, playing $Not\ Confess$, after a history of cooperation.
- If she sticks to the GTS, she earns a payoff of 4 thereafter, as depicted in the flat dashed-line in the middle of the figure.

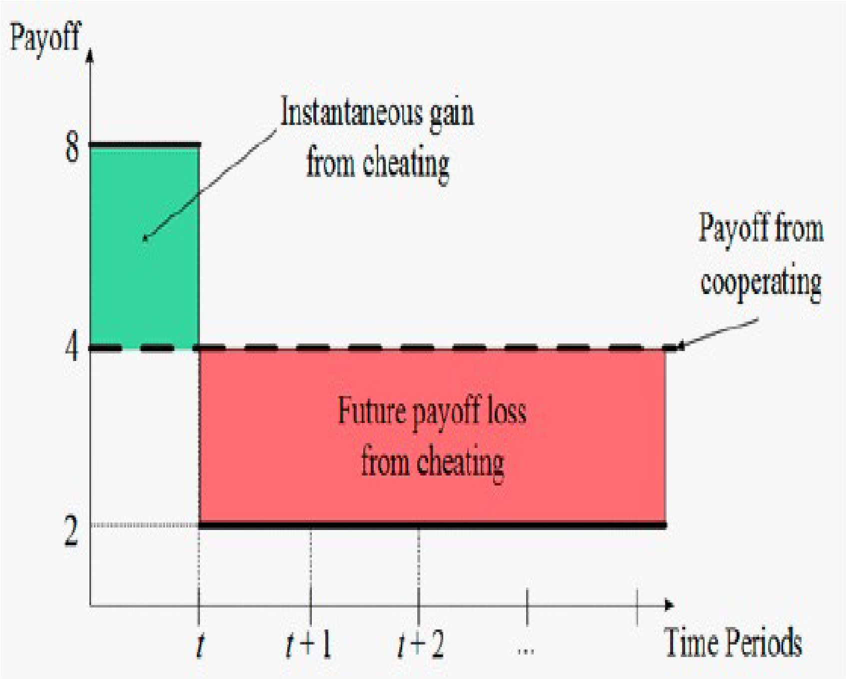

Figure 7.3. Incentives from cheating in the infinitely-repeated prisoner's dilemma game.

Summary

- If, instead, she cheats:
  - Her current payoff increases from 4 to 8 during one period.
  - But her rival detects her cheating, triggering a punishment with (C, C) thereafter, with associated payoff 2.
- Relative to what she earns by cooperating (4), cheating provides:
  - An instantaneous gain
  - But a future loss due to the punishment.

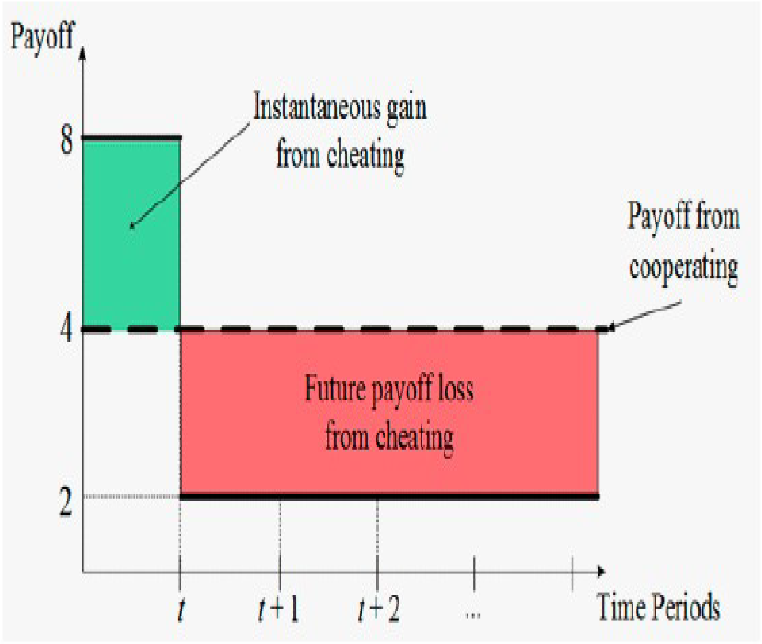

Figure 7.3. Incentives from cheating in the infinitely-repeated prisoner's dilemma game.

Extensions: Temporary Punishments

- The GTS in the above example assumes an infinite reversion to the NE of the stage game (unrepeated version of the game)
- But, what if:
  - players only revert to this NE during a finite number of periods, N,
  - moving back to cooperation once every player observes that both players implemented the punishment during N periods?
- Graphically, the right rectangle in the figure would be narrower.
- In summary, shortening the punishment, while keeping its severity, shrinks the parameter values where we can sustain cooperation.

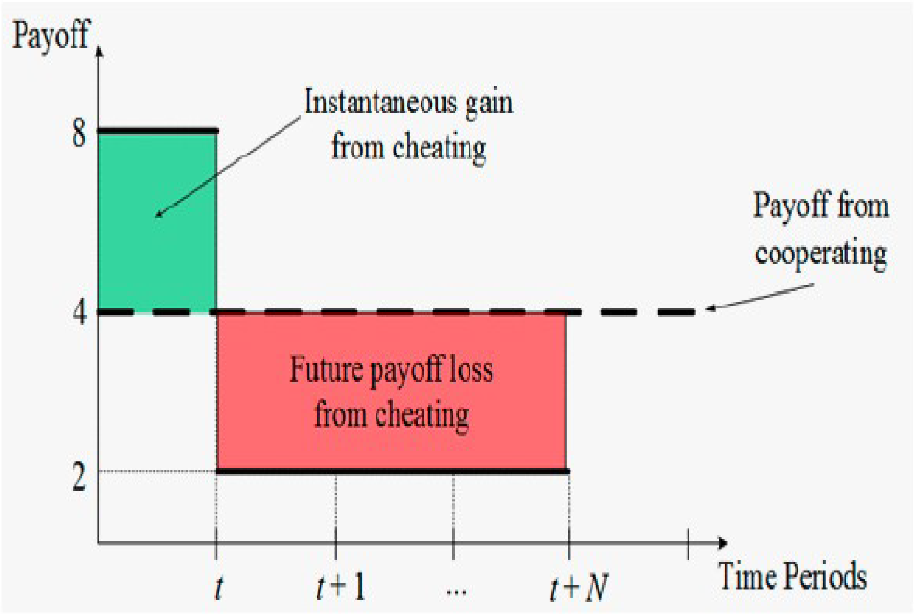

Figure 7.4. Temporary punishments.

Extensions: More Attractive Cheating

- If every player i earns 10 rather than 8 when she chooses $Confess$ while her opponent plays $Not\ Confess$…
  - the height of the instantaneous gain from cheating in Figure 7.5 increase, thus making cheating more attractive.
- Formally, this increases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
  - where the GTS can be sustained as a SPE if $\delta \geq \underline{\delta}$.

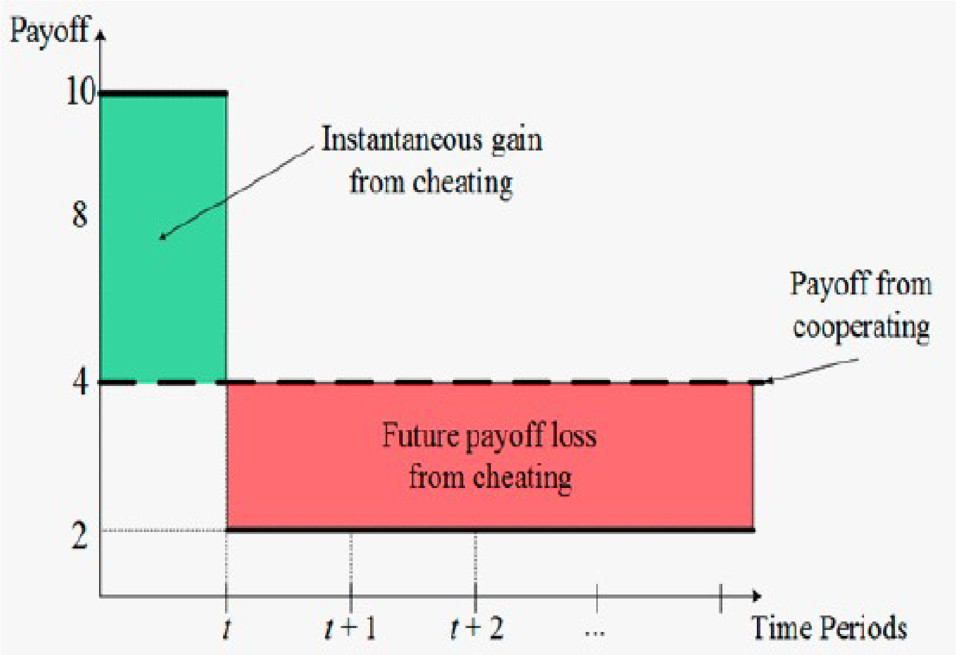

Figure 7.5. More attractive cheating.

Extensions: More Attractive Cheating

- Because the range of $\delta's$ sustaining cooperation satisfy $\delta \in [\underline{\delta}, 1)$, an increase in cutoff $\underline{\delta}$ entails that cooperation emerges in a more restricted range of $\delta's$.
- For compactness, the literature says that:
  - an increase in cutoff $\underline{\delta}$ "hinders cooperation in equilibrium",
  - while a decrease in $\underline{\delta}$ facilitates such cooperation.

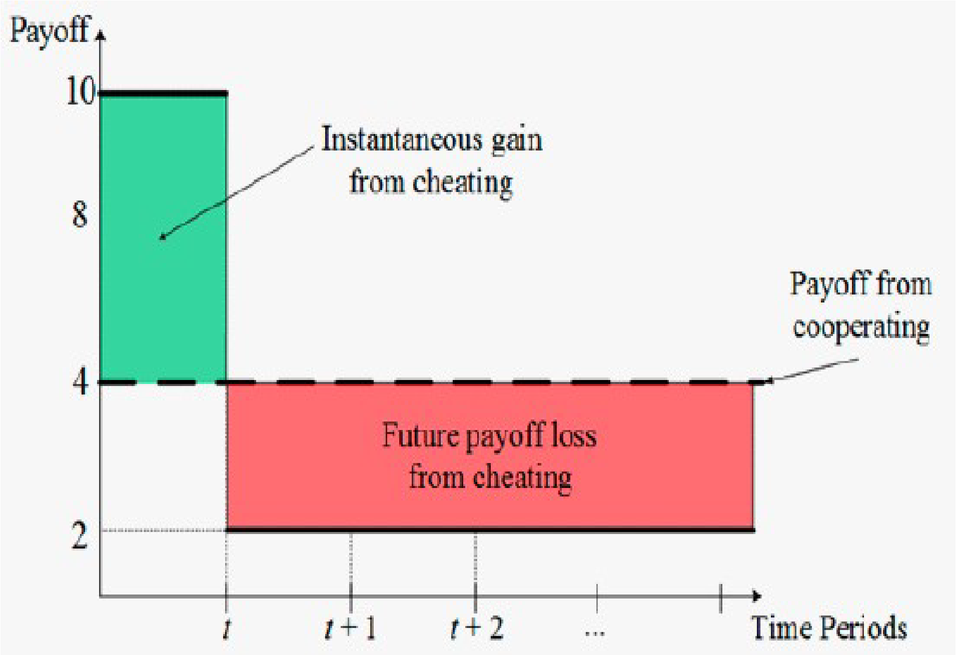

Figure 7.5. More attractive cheating.

Extensions: More Severe Punishments

- If every player earns 0 at the NE of the game, rather than 2:
  - the right rectangle in Figure 7.3 becomes deeper, as depicted in Figure 7.6,
  - indicating a more severe future payoff loss from cheating today.
- Intuitively:
  - cheating becomes less attractive,
  - which decreases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
  - implying that the GTS can be sustained under a larger range of $\delta' s$, i.e., $\delta \in [\underline{\delta}, 1)$

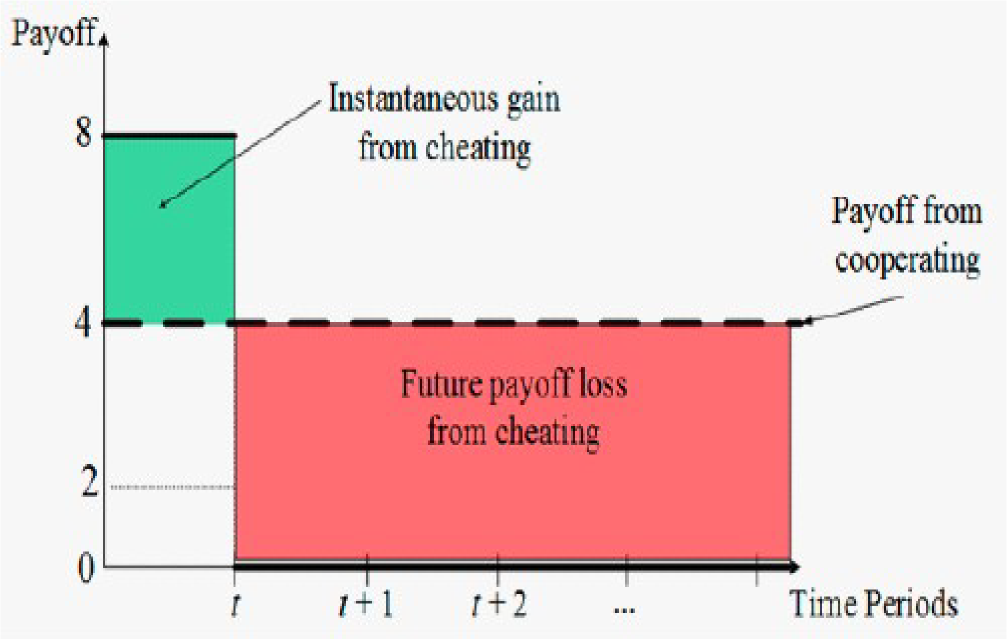

Figure 7.6. More severe punishment.

Extensions: Lag in Detecting Cheating

- However, in some real-life examples, players may detect cheating $k \in Z_+$ periods after it happened.
- If:
  - $k = 0$, we would still have immediate detection,
  - while $k > 0$ entails a lag in detecting cheating episodes.

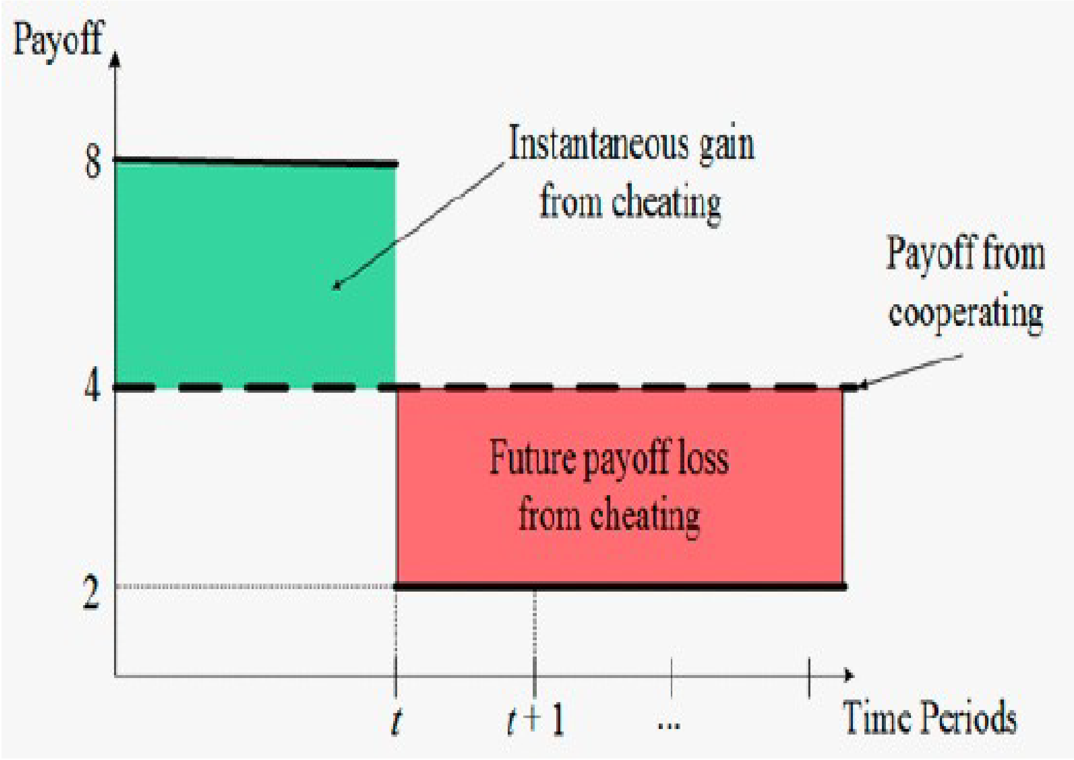

Figure 7.7. Lag at detecting cheating.

Extensions: Lag in Detecting Cheating

- This lag widens the left square in Figure 7.7 representing the gain from cheating.
  - It is not instantaneous in this setting.
  - Meaning that player i enjoys her cheating payoff during $1 + k$ periods.
- This makes cheating more attractive and:
  - increases the minimal discount factor sustaining cooperation, $\underline{\delta}$,
  - which hinders cooperation in equilibrium.

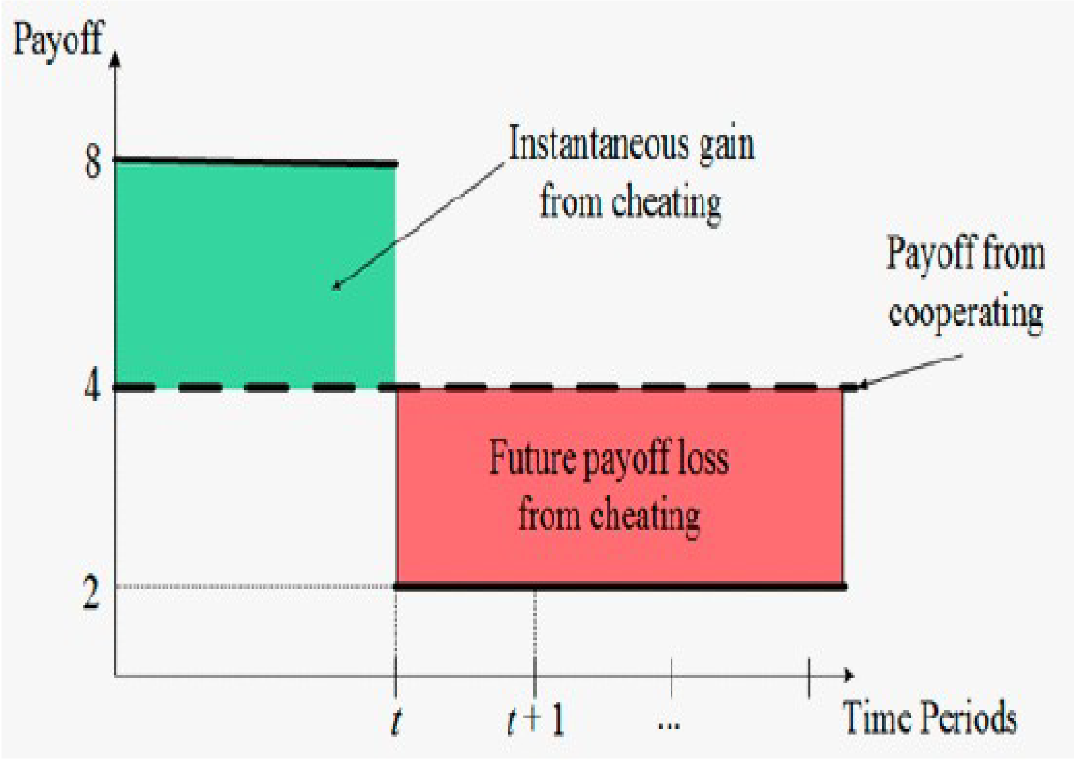

Figure 7.7. Lag at detecting cheating.

Extensions: Lag in Starting Punishments

- A similar argument as in point (4) applies if players:
  - Despite detecting cheating immediately after it occurs…
  - Need k periods to revert to the NE of the stage game.
- This happens, for instance, when “cooperation” means producing few units while “cheating” indicates producing a significantly larger number of units.
- If the cheated firm needs several periods to expand its production process, the beginning of the punishment phase is, essentially, delayed.
- This lag expand the payoff gain from cheating in the left square of Figure 7.7 during $1 + k$ periods.

Folk Theorem

- Can we also support other, partially cooperative outcomes, where, for instance, players choose $Confess$ during only some periods?
- In particular, we seek to identify the *per-period payoffs* that players earn at different SPEs of this game.
- Definition.
- **Per-period payoff.** If player i's present value from an infinite payoff stream is defined as $PV_i = \sum_{t=0}^{\infty} \delta^t v_i^t$, her per-period payoff $\overline{v_i}$ is the constant payoff that solves

$$PV_i = \frac{\overline{v_i}}{1-\delta}$$

or, after solving $\overline{v_i}$, $\overline{v_i} = (1-\delta)PV_i$.

- Intuitively, when player i receives that constant sum $\overline{v_i}$ in every period, she is indifferent between that (flat) payoff stream and her (potentially variable) stream of payoffs, as they both yield the same present value, PV_i.

Feasible and Individually Rational Payoffs

- We now seek to find all SPEs that can be sustained in the infinitely-repeated version of the game
  - and, as a consequence, predict which per-period payoffs players earn in each SPE.
- Afterwards, we restrict this set of payoffs to those where every player earns a higher payoff than at the NE of the stage game.

- Definition.
- **Feasible Payoffs (FP).** A feasible payoff vector $v = (v_1, v_2, \ldots, v_N)$ can be achieved by convex combinations of any two payoff vectors v' and v''.

- Therefore, in payoffs vector v, where player i earns v_i, is found by a convex combination $v_i = \alpha v_i' + (1-\alpha)v_i''$ where $\alpha \in [0,1]$ represents the weight on v_i'.
- Intuitively, the FP set captures all possible payoff vectors that players can earn if they play the game in different ways.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- First, each vertex depicts one of the payoff pairs that players earn by playing one of the pure strategy profiles every period.
- But as Figure 7.8 suggests, FP includes more than just the four vertices.
- Convex combinations of these vertices can yield other, still feasible, per-period payoffs.
- Examples next.

		Player 2: Confess	Player 2: Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game

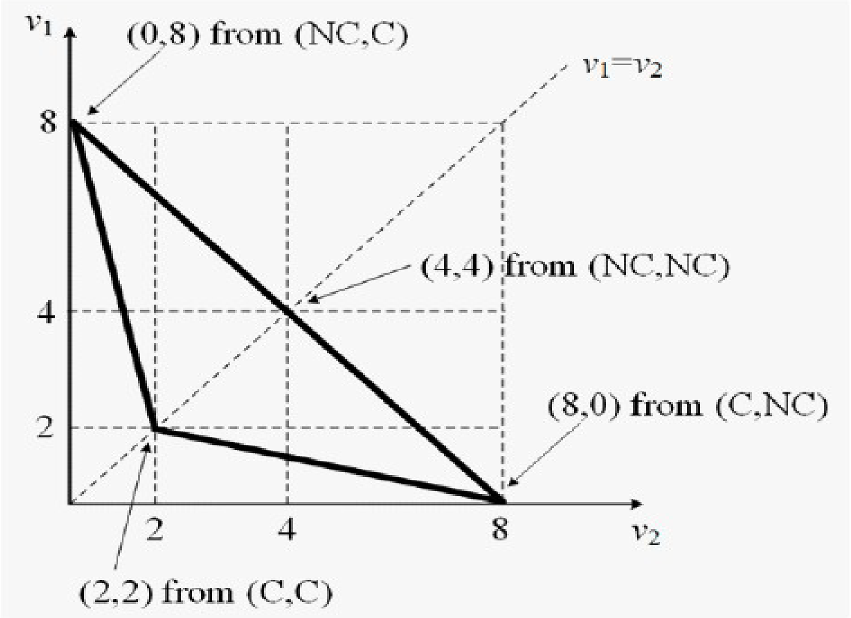

Figure 7.8. FP set in the prisoner's dilemma game.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- For instance, if players alternate between (C,C) and (NC,NC), the per-period payoff is 0.5*2+0.5*4=3 to each player.
  - Graphically positioned in the diagonal connecting (2,2) and (4,4).
- If, instead, players alternate between (C,C) and (NC,C), the per-period payoff for:
  - player 1 is 0.5*2+0.5*0=1,
  - while that of player 2 is 0.5*2+0.5*8=5.
- Similarly, for all payoff pairs inside the FP set.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game

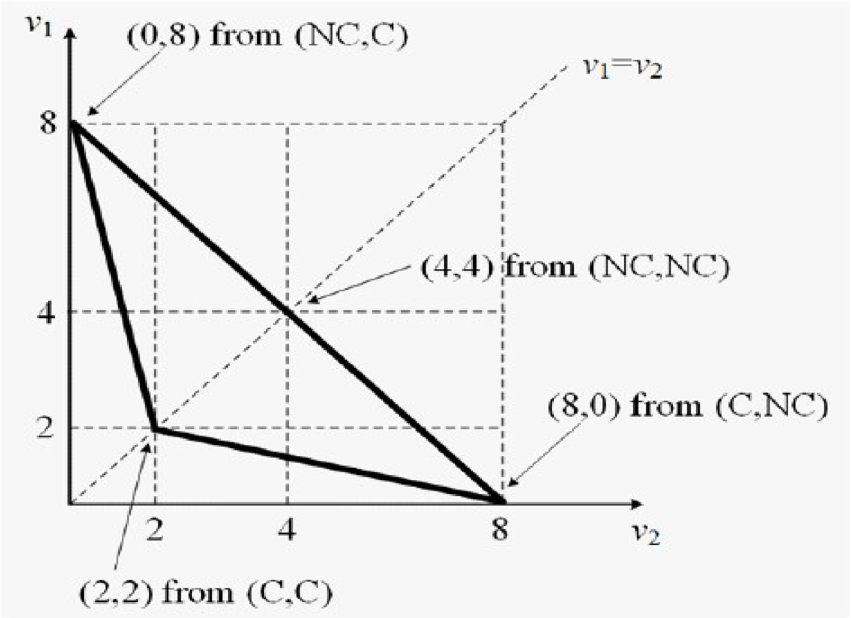

Figure 7.8. FP set in the prisoner's dilemma game.

Example 7.1. Finding FP in the Prisoner's Dilemma Game

- Recall that the FP set does not mean that players earn any payoff pair in the FP every period.
- Instead, it means that players can play the game in such a way that, even if their payoffs vary across time, their per-period payoff would lie inside the FP.

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.5. The Prisoner's Dilemma Game

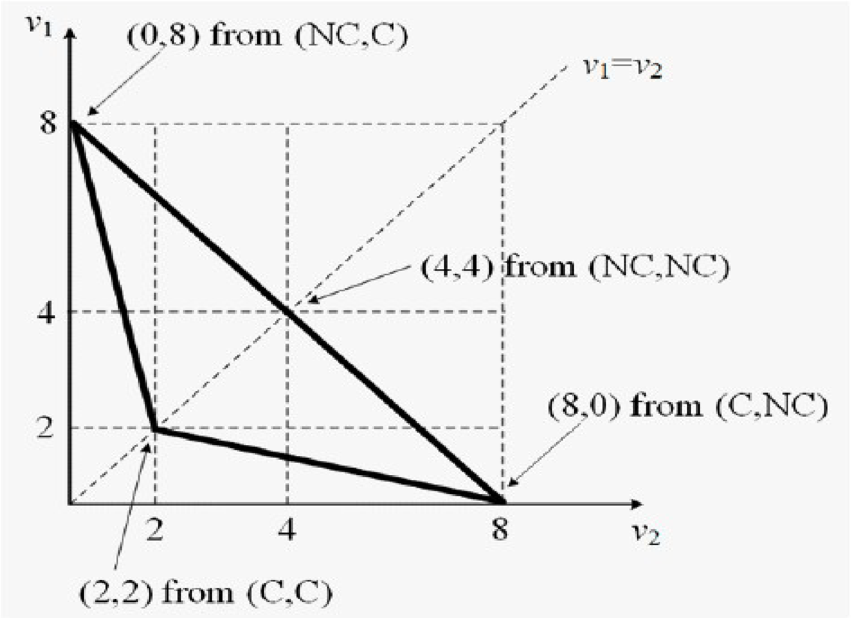

Figure 7.8. FP set in the prisoner's dilemma game.

Individually Rational Payoffs (FIR)

- Definition.
  - **Individually Rational Payoffs (IR).** An individually rational (IR) payoff vector $v = (v_1, v_2, \dots, v_N)$ satisfies

    $$v_i \geq v_i^{NE} \text{ for every player } i,$$

  - where v_i^{NE} denotes her NE payoff in the unrepeated version of the game.

- Combining FP and IR, we obtain FIR set (feasible individually rational payoffs).
- Example next.

Individually Rational Payoffs (FIR)

- Example 7.2. **Finding FIR payoffs.**
- Consider the Prisoner's Dilemma game of Example 7.1.
- A per-period payoff of player i is individually rational if $v_i \geq 2$, since every player earns 2 in the NE of the stage game.

Feasibly Individually Rational Payoffs (FIR)

- Figure 7.9 depicts:
  - a vertical line representing $v_1 \geq 2$ for player 1, which holds for all payoffs to the right-hand side of 2, and
  - a horizontal line capturing $v_2 \geq 2$ for player 2, which occurs for all payoffs above 2.
- The FIR diamond is the shaded area of the FP set, indicating that both players earn a higher per-period payoff than at the NE of the stage game.

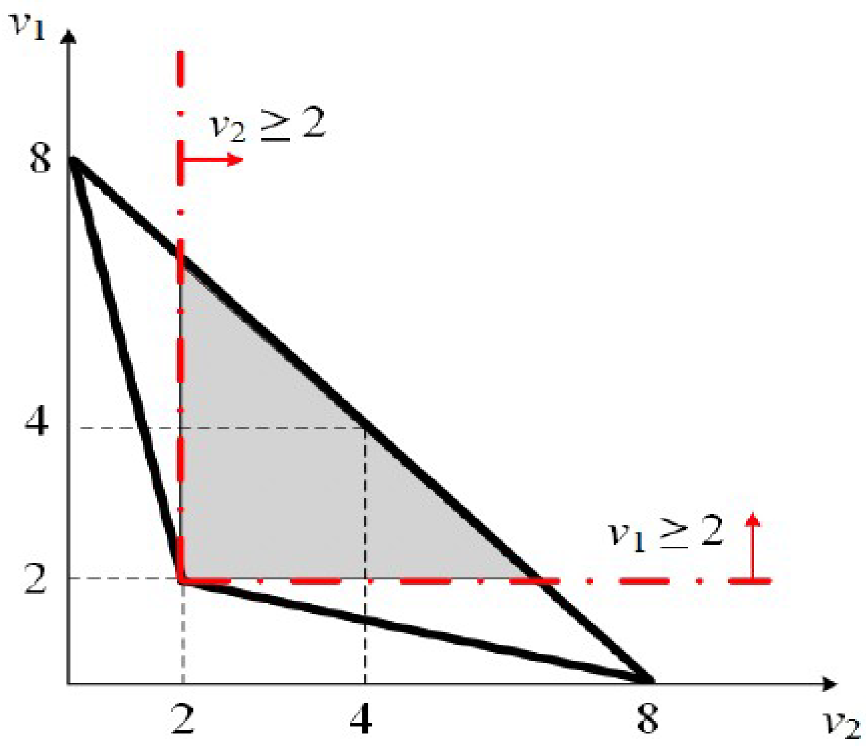

Folk Theorem and Cooperation

- Definition. **Folk Theorem.**
- Every per-period payoff vector $v = (v_1, v_2, \dots, v_N)$ in the FIR set can be sustained as the SPE of the infinitely-repeated game for a sufficiently high discount factor, δ, where $\delta \geq \underline{\delta}$.

- Graphically, the Folk theorem says that any point in the FIR diamond (on the edges or strictly inside) can be supported as a SPE of the infinitely-repeated game as long as players care enough about the future (high δ).
- *Examples*:
  - The uncooperative outcome, where (C, C) emerges in every period, can be sustained for all values of δ, yielding per-period payoffs $(2,2)$ as depicted on the southwest corner of the FIR diamond.
  - Similarly, the fully cooperative outcome, where (NC, NC) arises in every period, yielding per-period payoffs $(4,4)$, can be supported if $\delta \geq \frac{2}{3}$.
  - What about partial cooperation, with other per-period payoffs? Next slide.

Example 7.3. Supporting Partial Cooperation

- Consider the following modified GTS where players:
  - Alternate between (NC, C) and (C, NC) over time,
  - starting with (NC, C) in the first period.
- If either or both players deviates, both players revert to the NE of the stage game, (C, C), forever.
- To determine whether this modified GTS can be sustained as a SPE, we must show that no player benefits by unilaterally deviating (cheating).

	Outcome	Payoffs
Period 1	NC, C	0,8
Period 2	C, NC	8,0
Period 3	NC, C	0,8
Period 4	C, NC	8,0
…		

Table 7.1. Modified GTS inducing partial cooperation

Example 7.3. Supporting Partial Cooperation

- *Odd-numbered periods.* When player 1 cooperates in this GTS, her stream of discounted payoffs starting at an odd-numbered period (e.g. period 1) is

$$
\begin{aligned}
0 + \delta 8 + \delta^2 0 + \delta^3 8 + \cdots &= 0(1 + \delta^2 + \cdots) + 8(\delta + \delta^3 + \cdots) \\
&= 8\delta(1 + \delta^2 + \cdots) \\
&= \frac{8}{1 - \delta^2}
\end{aligned}
$$

since $\sum_{t=0}^{\infty} \delta^{2t} = \delta^0 + \delta^2 + \delta^4 + \cdots = \frac{1}{1-\delta^2}$.

Example 7.3. Supporting Partial Cooperation

- If, instead, player 1 deviates to C in an odd-numbered period, her current payoff increases from 0 to 2, yielding a stream of discounted payoffs of

$$\underbrace{2}_{Deviation} + \underbrace{\delta 2 + \delta^2 2 + \delta^3 2 + \cdots}_{Reversion\ to\ NE} = 2(1 + \delta + \delta^2 + \cdots) = \frac{2}{1-\delta}$$

- ***Comparison:*** Comparing player 1's payoff streams, we find that she sticks to the GTS in every odd-numbered period if

$$\frac{8\delta}{1-\delta^2} \geq \frac{2}{1-\delta}$$

and since $1-\delta^2 = (1-\delta)(1+\delta)$, we can rearrange the above equality to obtain $8\delta \geq 2(1+\delta)$, yielding $\delta \geq \frac{1}{3}$.

Example 7.3. Supporting Partial Cooperation

- *Even-numbered periods.* When player 1 cooperates in this GTS, her stream of discounted payoffs starting at an odd-numbered period (e.g. period 1) is

$$8 + \delta 0 + \delta^2 8 + \delta^3 0 + \cdots = 8(1 + \delta^2 + \cdots) + 0(\delta + \delta^3 + \cdots) = \frac{8}{1 - \delta^2}$$

- If, instead, player 1 unilaterally deviates from C to NC, in an even-numbered period, her current payoff actually decreases (from 8 to 4), yielding a stream of discounted payoffs

$$\underbrace{4}_{Deviation} + \underbrace{\delta 2 + \delta^2 2 + \delta^3 2 + \cdots}_{Reversion\ to\ NE} = 4 + 2\delta(1 + \delta + \delta^2 + \cdots) = 4 + \frac{2\delta}{1 - \delta}$$

Example 7.3. Supporting Partial Cooperation

- ***Comparison:*** Therefore, player 1 sticks to the GTS instead of deviating at even-numbered periods if

$$\frac{8}{1-\delta^2} \geq 4 + \frac{2\delta}{1-\delta}$$

and since $1-\delta^2 = (1-\delta)(1+\delta)$, we can rearrange this equality to obtain

$$8 \geq 4(1-\delta)(1+\delta) + 2\delta(1+\delta),$$

which further simplifies to $2\delta^2 - 2\delta + 4 \geq 0$.

Example 7.3. Supporting Partial Cooperation

- ***Comparison:***
  - This inequality further simplifies to $2\delta^2 - 2\delta + 4 \geq 0$.
  - Plotted here for illustration purposes:
    - Showing that $2\delta^2 - 2\delta + 4$ lies above the horizontal axis for all δ.

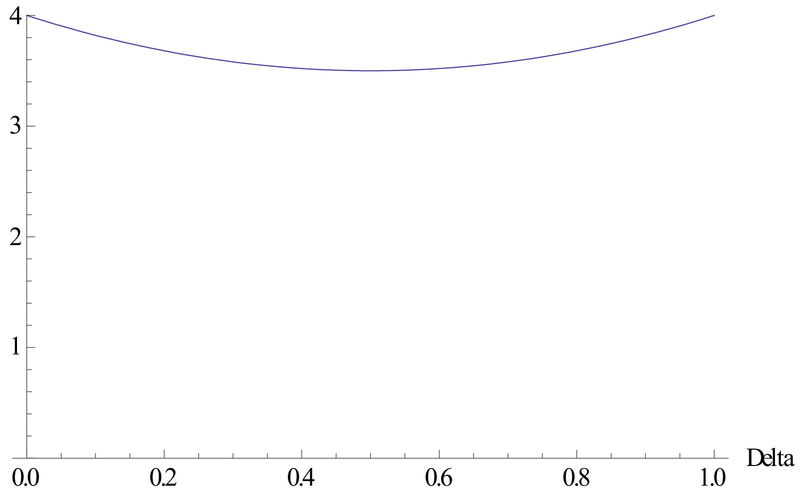

- This inequality holds for all admissible values of $\delta \in [0,1]$
- This means that player 1 sticks to the GTS in every even-numbered period regardless of her discount factor.

Folk Theorem

- The Folk theorem provides us with a positive and a negative result.
- Positive result:
  - We can sustain cooperation in games where such cooperation couldn't be supported in their unrepeated version or in their finitely-repeated versions.
  - In other words, players can reach Pareto-improving outcomes.
- Negative result:
  - We can reach *any* per-period payoff in the FIR diamond.
  - That's a large set of SPEs!
  - Limited predictive power.

Application to Collusion in Oligopoly

- Consider the duopoly with two firms competing à la Cournot.
- Firms face the same marginal cost of production c, where $1 > c \geq 0$
- Inverse demand function is $p(Q) = 1 - Q$, where $Q \geq 0$ denotes aggregate output
- We know:
  - NE is $(q_i, q_j) = \left(\frac{1-c}{3}, \frac{1-c}{3}\right)$
  - $\pi_i = \frac{(1-c)^2}{9}$
- As in the Prisoner's Dilemma game, we study next if a GTS can be sustained as a SPE of the infinitely-repeated version of this Cournot competition game.

Application to Collusion in Oligopoly: Collusion

$$\max_{q_1,q_2\geq 0} \pi(q_1,q_2) = \underbrace{[(1-q_1-q_2)q_1 - cq_1]}_{\pi_1} + \underbrace{[(1-q_1-q_2)q_2 - cq_2]}_{\pi_2}$$

Since $Q = q_1 + q_2$

$$\max_{Q\geq 0} \quad \pi(Q) = (1-Q)Q \; - cQ$$

which leads to

$$Q^C = \frac{1-c}{2}$$

$$q_i^C = \frac{Q^C}{2} = \frac{1-c}{4}$$

$$p^C = 1 - Q^C = 1 - \frac{1-c}{2} = \frac{1+c}{2}$$

$$\pi_i^C = (p^C - c)q_i^C = \frac{(1-c)^2}{8}$$

GTS in Collusion

- We are now ready to specify the GTS in this game:
  1. In period $t = 1$, every firm i chooses the collusive output $q_i^C = \frac{1-c}{4}$.
  2. In all periods $t > 1$, every firm i chooses the collusive output $q_i^C = \frac{1-c}{4}$ if both firms produced q_i^C in every previous period. Otherwise, every firm i reverts to the NE of the stage game, choosing $q_i = \frac{1-c}{3}$, thereafter.
- To show that this GTS can be sustained in equilibrium, we need to show that firms have no incentives to deviate from it:
  - at every period t, and
  - after any previous history of play.
- We separately analyse whether firms have incentives to deviate from the GTS:
  - after observing that firms cooperated in previous periods and
  - after observing some uncooperative episodes.

GTS in Collusion

- **After a history of cooperation.** If both firms cooperated in all previous periods, choosing q_i^C, every firm i can collude in period t, as specified by the GTS, or deviate from this collusive output.
- *Profits from collusion.* If firm i colludes, it earns $\pi_i^C = \frac{(1-c)^2}{8}$ in every period, entailing a discounted present value of

$$\pi_i^C + \delta\pi_i^C + \delta^2\pi_i^C + \cdots = \frac{(1-c)^2}{8(1-\delta)}$$

Profits from Deviation

If, instead, firm i deviates from the collusive output q_i^C, we must first determine firm $i's$ *most profitable unilateral deviation*, which means that firm j still produces the collusive output, q_j^C.

$$\max_{q_i \geq 0} \; \pi = \left(1 - q_i - q_j^C\right) q_i \; - c q_i$$

which leads to

$$q_i\left(q_j^C\right) = \frac{1-c}{2} - \frac{1}{2} q_j^C$$

since $q_j^C = \frac{1-c}{4}$,

$$q_i^{Dev} = \frac{3(1-c)}{8}$$

(This can be directly done by inserting $q_j^C = \frac{1-c}{4}$ into firm *i*'s BRF.)

Therefore, equilibrium price is

$$p^{Dev} = 1 - q_i^{Dev} - q_j^C = \frac{3+5c}{8}$$

Profits from Deviation

And the profits,

$$\pi_i^{Dev} = (p^{Dev} - c)q_i^{Dev} = \frac{9(1-c)^2}{64}$$

Therefore, if firm i deviates at any period t, it earns deviation profits $\pi_i^{Dev} = \frac{9(1-c)^2}{64}$ in that period, but Cournot profits in every period afterward

$$\pi_i^{Dev} + \delta\pi_i + \delta^2\pi_i + \cdots = \pi_i^{Dev} + \delta\pi_i(1 + \delta + \delta^2 + \cdots)$$
$$= \pi_i^{Dev} + \pi_i\left(\frac{\delta}{1-\delta}\right)$$
$$= \frac{9(1-c)^2}{64} + \frac{\delta(1-c)^2}{9(1-\delta)}$$

Comparison: We can say that at any period t after a history of collusion, firm i keeps colluding if

$$\pi_i^C\left(\frac{1}{1-\delta}\right) \geq \pi_i^{Dev} + \pi_i\left(\frac{\delta}{1-\delta}\right) \Rightarrow \delta \geq \delta \equiv \frac{9}{17}$$

Therefore, collusion can be sustained if firms assign a sufficiently high value to future profits. Otherwise, collusion cannot be sustained.

After a History of No Cooperation

- If one firm deviated from collusive output $q_i^C = \frac{1-c}{4}$ in any previous period, the GTS prescribes that every firm i reverts to the NE of the stage game.
- To confirm that the GTS is a SPE of the infinitely-repeated game, we need to show that every firm i has incentives to, essentially, implement this punishment.
- If, upon observing a deviation, firm i behaves as prescribed by the GTS, producing the Cournot output, it earns a discounted profit stream of

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

- If, instead, firm i unilaterally deviates from the Cournot output, producing $q_i \neq \frac{1-c}{3}$, while firm j behaves as prescribed by the GTS, its best response function prescribes firm i to choose:

$$q_i\left(\frac{1-c}{3}\right) = \frac{1-c}{2} - \frac{1}{2}\frac{1-c}{3} = \frac{1-c}{3}$$

- In other words, even if firm i wanted to deviate from the GTS, its best response would be to behave as prescribed by the GTS.

Minimal Discount Factor supporting collusion

- Our analysis above considered that firms compete à la Cournot, sell a homogeneous good, and face the same marginal cost of production, c
- However, we could allow for firms to compete in prices, sell differentiated products, and/or face different production costs.
- Here, we follow a more general approach.
- In particular, after a history of cooperation, every firm i's present value from cooperating in period t is

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

Minimal Discount Factor supporting collusion

- If, instead, firm i deviates at period t, it earns π_i^{Dev} during that period but π_i^{NE} thereafter, yielding:

$$\pi_i^{Dev} + \delta\pi_i^{NE} + \delta^2\pi_i^{NE} + \cdots = \pi_i^{Dev} + \pi_i^{NE}(1 + \delta + \delta^2 + \cdots)$$
$$= \pi_i^{Dev} + \pi_i^{NE}\left(\frac{1}{1-\delta}\right)$$

Minimal Discount Factor supporting collusion

- Therefore, firm i cooperates if and only if

$$\pi_i^C\left(\frac{1}{1-\delta}\right) \geq \pi_i^{Dev} + \pi_i^{NE}\left(\frac{1}{1-\delta}\right)$$

which we can rearrange as $\pi_i^C \geq (1-\delta)\pi_i^{Dev} + \delta\pi_i^{NE}$.

- Solving for δ, yields a minimal discount factor:

$$\delta \geq \frac{\pi_i^{Dev} - \pi_i^C}{\pi_i^{Dev} - \pi_i^{NE}} \equiv \underline{\delta}$$

Minimal Discount Factor supporting collusion

$$\delta \geq \frac{\pi_i^{Dev} - \pi_i^{C}}{\pi_i^{Dev} - \pi_i^{NE}} \equiv \underline{\delta}$$

- Minimal discount factor $\underline{\delta}$ increases in $\pi_i^{Dev} - \pi_i^{C}$.
  - As deviations become more attractive, cooperation can only be sustained for higher discount factors.
- In contrast, $\underline{\delta}$ decreases in $\pi_i^{Dev} - \pi_i^{NE}$.
  - When profit loss from reverting to the NE of the stage game become more severe, firms have less incentives to cheat, expanding the range of discount factors sustaining cooperation.

Minimal Discount Factor supporting collusion

- The above ratio is convenient, as it can be readily applied to different market structures:
  - markets with two or more firms, selling homogeneous products, with symmetric or asymmetric costs, etc.
- We only need to find:
  - the NE in the stage game, and its associated profits, π_i^{NE}
  - the collusive output (or prices) that maximize firms' profits, yielding π_i^C to firm i; and
  - firm i's optimal deviation, fixing firm $j's$ collusive behaviour, so firm i's profits are π_i^{Dev}
- Inserting profits π_i^{NE}, π_i^C, and π_i^{Dev} in the above expression for $\underline{\delta}$, we obtain the minimal discount factor sustaining collusion.

Other Collusive GTS

- We said that the above GTS with perfect collusion, where firms maximize their joint profits, can be sustained:
  - “as a SPE,”
  - rather than saying “as the unique SPE,” of the infinitely repeated game.
- From the Folk theorem, we know that the other per-period payoffs can be supported as SPEs too.

Other Collusive GTS

- The above GTS can be interpreted as the most northwest payoff pair in the FIR set
  - since firms earn the highest symmetric payoff relative to what they earn in the NE of the Cournot model.
- More generally, we can characterize the line connecting profits in the above GTS and profits in the NE of the stage game as follows:

$$\pi_i = \alpha \pi_i^C + (1-\alpha)\pi_i^{NE}$$

where weight $\alpha \in [0,1]$ can be understood as how close firms are to the perfect collusion.

- In particular, firm i produces

$$q_i = \alpha q_i^C + (1-\alpha) q_i^{NE} = \alpha \frac{1-c}{4} + (1-\alpha)\frac{1-c}{3} = \frac{(1-c)(4-\alpha)}{12}$$

Implying that this firm's output coincides with the collusive output when $\alpha = 1$, but becomes the NE output when $\alpha = 0$.

Other Collusive GTS

- Consider then, the following GTS
  1. In period $t = 1$, every firm i chooses the output $q_i = \alpha q_i^C + (1 - \alpha) q_i^{NE}$
  2. In all periods $t > 1$:
     - Every firm i chooses output q_i if both firms produced q_i in every previous period.
     - Otherwise, every firm i reverts to the NE of the stage game, choosing $q_i = \frac{1-c}{3}$, thereafter.

What if the stage game has more than one NE?

- So far, we considered games with a unique NE in its unrepeated version.
  - *Examples*: the Prisoner's Dilemma game, Cournot quantity competition, or Bertrand price competition.
- The NE in these games was rather uncooperative.
  - Yet, we showed that cooperation can be supported as a SPE of the infinitely-repeated game.
- A natural question is whether cooperation can be sustained in games with more than one NE in its unrepeated version as in Matrix 7.6.

What if the stage game has more than one NE?

		Player 2		
		A	B	C
Player 1	A	8,8	-2,<u>10</u>	2,3
	B	<u>10,</u>-2	<u>2,2</u>	1,1
	C	1,1	1,1	<u>6,6</u>

Matrix 7.6. A game with two NEs

- The game has two psNEs:
  1. A "bad" NE, (B,B), where every player earns a payoff of 2.
  2. A "good" NE, (C,C), where every player earns a payoff of 6.
     - Formally, we say that (C,C) Pareto dominates (B,B).
  3. Yet, there is an even better outcome, where every player earns higher payoffs:
     - At (A,A) every player earns 8.
     - How can we sustain the cooperative outcome (A,A) being played?

What if the stage game has more than one NE?

Player 2

		A	B	C
Player 1	A	8,8	-2,10	2,3
	B	10,-2	2,2	1,1
	C	1,1	1,1	6,6

Matrix 7.6. A game with two NEs

- Consider the following GTS:
  1. In the first period, every player chooses A.
  2. In the second period:
     - Every player chooses C if (A, A) was played in the first period.
     - Otherwise, every player chooses B.

Stick-and-carrot

- Informally, this is known as a “stick-and-carrot” strategy, because in the second period:
  - It prescribes the Pareto dominant NE, (C, C), if players cooperated in the first period.
  - But prescribes the Pareto dominated NE, (B, B), if either player was uncooperative.
- Important lesson:
  - While a GTS can specify an outcome that is not a NE of the stage game during the first period, such as (A, A),…
  - It cannot prescribe an outcome that is not a NE in the last period.
  - Otherwise, if the GTS specified outcome (A, A) in both periods, every player would have incentives to deviate from A in the second (last) period.
- In summary, when designing GTSs in finitely-repeated games, we can:
  - “Shoot for the stars” during *T*-1 periods, but…
  - must settle for a NE in the last round of interaction.

Stick-and-carrot, Solution

Second period:

- Technically, there are 9 outcomes in each stage, for a total of $9^2 = 81$ terminal nodes
- Out of the 9 second-period subgames, there is only one initiated after outcome (A, A) was played in the first period.
  - The remaining 8 second-period subgames emerge because one or both players did not select A in the first period.
- *Non-cooperative history.* Upon observing any outcome <u>different</u> from (A, A) in the first period:
  - Player i anticipates that player j will play B in the second period.
  - Player i does not have incentives to deviate from the GTS in the second period, since B is a best response to B (see middle column, for instance).
- *Cooperative history.* If, instead, player i observes that outcome (A, A) was played in the first period:
  - She expects that player j will choose C in the second period.
  - Therefore, player i does not want to deviate from the GTS in the second period, since C is a best response to C (see right column, for instance).

Stick-and-carrot, Solution

- **First period:**
  - If player i chooses A, as prescribed by the GTS, her payoff stream is
  $$8 + \delta 6,$$
  since he anticipates that (C, C) will be played in the second period (the Pareto dominant NE, or "carrot")

  - If, instead, player i unilaterally deviates:
    - Her best deviation is to B, which yields a payoff of 10 today,
    - but that triggers outcome (B, B), the Pareto dominated NE or "stick," in the second period,
    - Payoff stream from deviation is, then:
    $$10 + \delta 2.$$

  - ***Comparison:*** player i behaves as prescribed by the GTS in the first period if
  $$8 + \delta 6 \ \geq 10 + \delta 2 \ \Rightarrow \ \delta \ \geq \frac{1}{2}$$

Stick-and-carrot, Solution

- Cooperation can, then, be sustained if players assign a sufficiently high weight to future payoffs.
- Players condition last-period play
  - Choosing the good or bad NE.
  - on whether players were cooperative or uncooperative in previous interactions.
- In stage games with a unique NE, we couldn’t condition last-period play.
  - Players will just behave according to the unique NE.
  - This inability to condition future play lead to the “unraveling” result, where players behave as in the stage game during every period of interaction.

Modified GTSs: An eye for an eye

- We now consider modified GTSs.
- “An eye for an eye”:
  - Code of Hammurabi, written around 1750 BC.
- It prescribes that individuals start cooperating, as in standard GTSs.
- However, if an individual is uncooperative:
  - the victim can be uncooperative, inflicting the same damage to the cheater,
  - and then players return to cooperation.
- In an infinitely-repeated game context, this adage means that, if a player cheats:
  - the cheater player gets to cheat during one period while the cheating player cooperates in that same period (so punishment phase only lasts one period),
  - afterwards players return to cooperation

Modified GTSs: An eye for an eye

- To implement this cheating sequence in a GTS, consider the game in Matrix 7.7:
    1. In the first period, every player i chooses C
    2. In subsequent periods:
        1. Every player i chooses C if (C, C) was played in all previous periods.
        2. Otherwise, every player reverts to the NE of the stage game, (B, B).
- This GTS can be sustained if $\delta \geq \frac{2}{3}$ (see Exercise 7.17).

		Player 2		
		A	B	C
Player 1	A	0,0	4,$\underline{2}$	5,-2
	B	$\underline{2}$,4	$\underline{8}$,$\underline{8}$	$\underline{14}$,3
	C	-2,5	3,$\underline{14}$	10,10

Matrix 7.7. An eye for an eye

Modified GTSs: An eye for an eye

- But what about using a GTS that follows the "an eye for an eye" adage?
  - The punishment phase now prescribes that the victim gets to cheat her opponent for a period.
  - Then they return to the cooperative outcome (C, C) thereafter.
- As an illustration, if player 1 cheats, deviating to B while player 2 still cooperates choosing C, table 7.2 summarizes the sequence of outcomes.

	Period $t-1$	Period t	Period $t+1$
Outcome	B, C	C, B	C, C
Payoffs	14,3	3,14	10,10
Who cheats?	$Player$ 1	$Player$ 2	No one

Table 7.2. Timing for an eye-for-an-eye GTS

After a History of Cooperation

- If no player has cheated on previous periods, player i's payoff steam from cooperating is
$$10 + 10\delta + 10\delta^2 + \cdots$$
- If, instead, player 1 deviates, we first find that, conditional on player 2 cooperating, player 1's BR is to play B, earning a payoff of 14.
- Therefore, player 1's deviation:
  - Increases her current payoff from 10 to 14,
  - But then triggers a punishment, which reduces her payoff to 3 during one period.
  - Afterwards, players return to cooperation.
- In summary, the stream of payoffs that player 1 earns when deviating is
$$\underbrace{14}_{Cheat} + \underbrace{3\delta}_{Cheated} + \underbrace{10\delta^2 + 10\delta^3 + \cdots}_{Back\ to\ cooperation}$$

Therefore, player cooperates if and only if
$$10 + 10\delta + 10\delta^2 + \cdots \geq 14 + 3\delta + 10\delta^2 + 10\delta^3 + \cdots$$
which simplifies to $10 + 10\delta \geq 14 + 3\delta$, or $\delta \geq \frac{4}{7}$.

After a History of Cheating

- When one player cheats, the GTS prescribes that the cheating player cooperates for one period, playing C, while her opponent cheats in that period, choosing B.
- If player 1 cheats in period $t-1$, the GTS prescribes that (C, B) is played in period t, and (C, C) occurs in all subsequent periods, yielding

$$\underbrace{3}_{Cheated} + \underbrace{10\delta + 10\delta^2 + 10\delta^3 + \cdots}_{Back\ to\ cooperation}$$

- If player 1 does not choose C in period t, being cheated, her best deviation is to B since, conditional on player 2 choosing B in period t (see Table 7.2), player 1's best response is B, leading to a payoff stream

$$\underbrace{8}_{Cheat} + \underbrace{3\delta}_{Cheated} + \underbrace{10\delta^2 + 10\delta^3 + \cdots}_{Back\ to\ cooperation}$$

After a History of Cheating

- Intuitively, if the cheating player does not follow the GTS, not allowing the victim to cheat during one period, he only postpones the return to cooperation.
- Therefore, player 1 allows player 2 to cheat her after player 1 cheated in the first place, if and only if

$$3 + 10\delta + 10\delta^2 + 10\delta^3 + \cdots \geq 8 + 3\delta + 10\delta^2 + 10\delta^3 + \cdots$$

which coincide after the third term, $10\delta^2$, and for all subsequent periods. We can then rewrite this inequality as

$$3 + 10\delta \geq 8 + 3\delta \Rightarrow \delta \geq \frac{5}{7}$$

Summary

- Comparing this condition on δ with that found after a history of cooperation, $\delta \geq \frac{4}{7}$, we see that condition $\delta \geq \frac{5}{7}$ is more demanding than $\delta \geq \frac{4}{7}$ since $\frac{5}{7} \geq \frac{4}{7}$.
- In other words, if condition $\delta \geq \frac{5}{7}$ holds, condition $\delta \geq \frac{4}{7}$ must also be satisfied, so $\delta \geq \frac{5}{7}$ is a sufficient condition for cooperation.

Comparing different GTS: An-eye-for-an-eye GTS vs Standard GTS

- **Equity of payoffs:**
  - We find that the punishment phase is more equitable in the former than the latter since the cheater party is compensated in the following period, when she cheats.
  - Under the standard GTS, however, the cheated party sees its payoff decrease during the cheating period, and then both players revert to the NE of the stage game, yielding a lower per-period payoff than the cheater's.
- **Minimal discount factors:**
  - An-eye-for-an-eye GTS can be sustained if $\delta \geq \frac{5}{7}$,
  - Standard GTS can be sustained if $\delta \geq \frac{2}{3}$,
  - Cooperation is, then, more difficult to arise in the former than the latter since $\frac{5}{7} \geq \frac{2}{3}$.
- *Intuition*: The punishment phase in this GTS requires the cheating party to be cheated during one period, which is not very attractive for this player, generating strong incentives to deviate from the GTS at precisely that moment.

Short and Nasty Punishments

- Consider the same payoff matrix as before.
- The punishment is now (A, A), yielding the lowest symmetric payoff in the matrix, $(0,0)$, and players return to cooperation immediately after.
- Consider the GTS as follows:
  1. In the first period, every player i chooses C
  2. In subsequent periods:
     a. Every player i chooses C if (C, C) was played in all previous periods
     b. Otherwise, every player chooses A during one period. If (A, A) was played in the last period, every player selects C thereafter

Short and Nasty Punishments

- If player 1 cheats, deviating to B while player 2 still cooperates choosing C
- Table 7.3 summarizes the sequence of outcomes if players behave according to this GTS

	Period $t-1$	Period t	Period $t+1$
Outcome	B, C	A, A	C, C
Payoffs	14,3	0,0	10,10

Table 7.3. Timing of the short-and-nasty punishments GTS

Short and Nasty Punishments

- Cooperation can be sustained with this GTS if $\delta \geq \frac{2}{5}$.
  - (Exercise 7.1, as a practice)
- Table 7.4 reports the minimal discount factor $\underline{\delta}$ under different GTS in the game in Matrix 7.7, where $\frac{5}{7} > \frac{2}{3} > \frac{2}{5}$.

	Minimal discount factor, $\underline{\delta}$
Standard GTS, section 7.4.2	$\underline{\delta} = \frac{2}{3}$
An-eye-for-an-eye GTS, section 7.8.1	$\underline{\delta} = \frac{5}{7}$
Short-and-nasty-punishments GTS, section 7.8.2	$\underline{\delta} = \frac{2}{5}$

Table 7.4. Minimal Discount Factors sustaining cooperation in the game in Matrix 7.7 when using different GTS

Imperfect Monitoring

- Players may imperfectly monitor her rival's actions, observing which action she chose with a certain probability.
- Consider Matrix 7.8:
  - (C, C) is the NE of the stage game,
  - but (NC, NC) would increase both players' payoffs from 2 to 4.

		Player 2: Confess	Player 2: Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.8. The Prisoner's Dilemma Game

Imperfect Monitoring

		Player 2	
		Confess	Not Confess
Player 1	Confess	2,2	8,0
	Not Confess	0,8	4,4

Matrix 7.8. The Prisoner's Dilemma Game

- Under imperfect monitoring, is player j chooses to cooperate, NC, we assume that, for simplicity, player i cannot perfectly detect deviation to C.
- If, instead, player j chooses C, the probability that player i observes a deviation from NC is $p \in [0,1]$.
- Intuitively:
  - when probability $p \to 0$, deviations are rarely observed,
  - while when $p \to 1$, deviations are almost always observed, as in a context with perfect monitoring.

Imperfect Monitoring

- Consider the standard GTS in this context:
  1. In the first period, every player chooses NC.
  2. In every subsequent period, every player chooses NC if (NC, NC) was *observed* in all previous periods. Otherwise, revert to the NE of the stage game, (C, C), thereafter.
- That is, the cooperative outcome (NC, NC) is now observed but it may not have been played.
  - This suggests that outcome (NC, NC) may not have been played, because at least one person deviated, but cooperation continues if such outcome is observed.
- The opposite also applies:
  - Outcome (NC, NC) may have been played, but cooperation stops if one of both players did not observe such outcome in a previous period.

After a history of observed cooperation

- If both players observe outcome (NC, NC) in all previous periods, player i's payoff stream from cooperating is

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

- If, instead player i unilaterally deviates to C, her payoff increases to 8 in the current period, but in the next period:
  - Player j detects her deviation with probability p, reverting to the NE of the stage game with a payoff of 2 forever, or…
  - Player j doesn't detect this deviation with probability $(1-p)$, which lets player i return to cooperation (choosing NC tomorrow) as if no deviation ever happened.
- In summary, player i's payoff stream from unilaterally deviating to C is

$$\underbrace{8}_{Current\ gain\ from\ dev.to\ C} + \delta\left[p \underbrace{\frac{2}{1-\delta}}_{Detected} + (1-p) \underbrace{\frac{4}{1-\delta}}_{Undetected}\right]$$

After a history of observed cooperation

- For player i to cooperate after a history of cooperation, we need that

$$\frac{4}{1-\delta} \geq 8 + \delta\left[p\frac{2}{1-\delta} + (1-p)\frac{4}{1-\delta}\right]$$

And solving for δ, yields

$$\delta \geq \frac{2}{2+p} \equiv \underline{\delta}(p)$$

- Figure 7.9 depicts the minimal discount factor $\underline{\delta}(p)$.

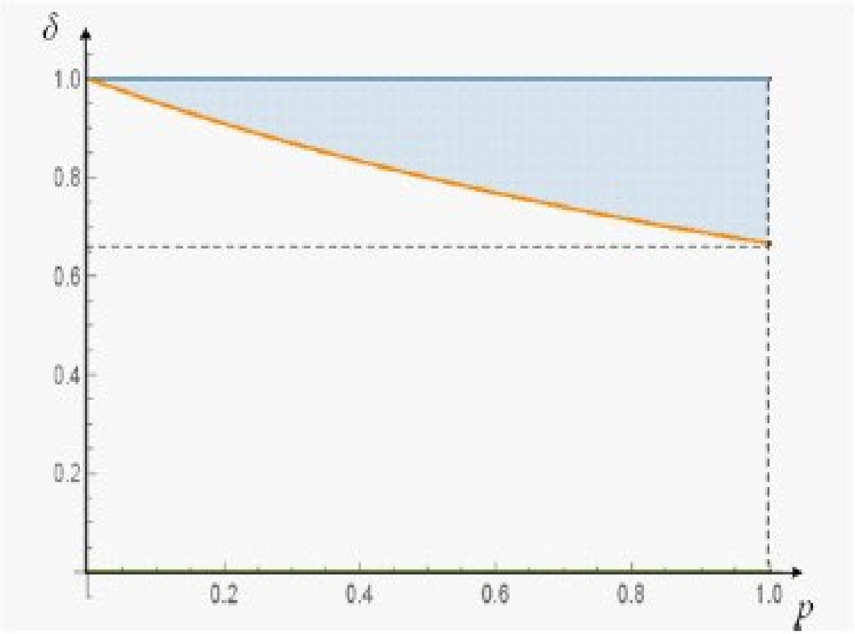

After a history of observed cooperation

$$\Rightarrow \delta \geq \frac{2}{2+p} \equiv \underline{\delta}(p)$$

- As deviations are more likely detected (higher p):
  - the expected punishment increase,
  - expanding the range of $\delta's$ sustaining cooperation.
- In contrast, when deviation cannot be detected $(p = 0)$:
  - the minimal discount factor becomes $\underline{\delta}(0) = 1$,
  - implying that collusion cannot be supported in equilibrium.

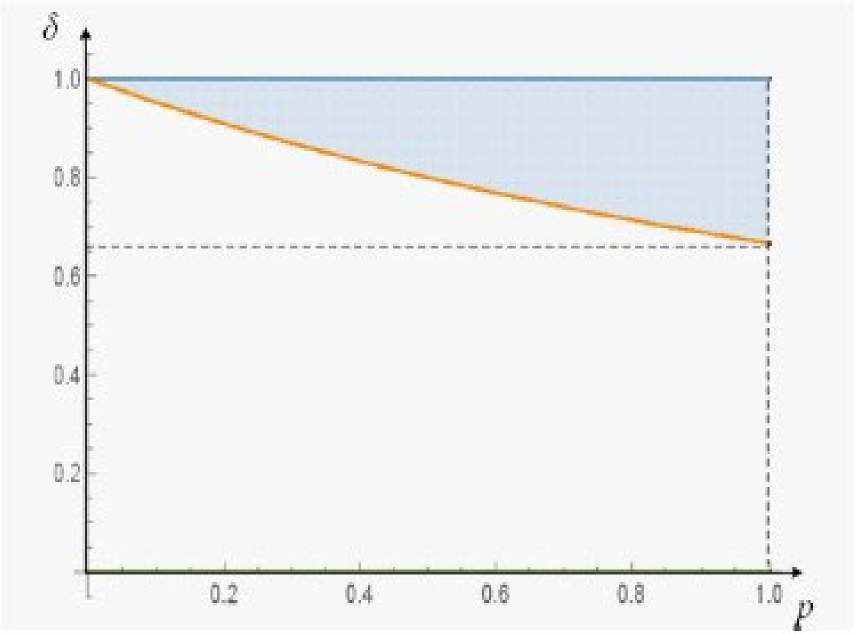

After a history of observed deviation

- If player observes that (NC, NC) was not played in a previous period, this means that player j did not cooperate.
- If player i behaves as prescribed by the GTS, she should revert to the NE of the stage game, (C, C), thereafter, earning a payoff stream

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

After a history of observed deviation

- If, instead, player i deviates, choosing NC while player j chooses C (recall that this is a unilateral deviation from the GTS, upon observing no cooperation in previous periods) player i earns 0 in this period.
- Players then play (C, C) in all subsequent periods yielding

$$0 + 2\delta + 2\delta^2 + \cdots = 2\delta(1 + \delta + \delta^2 + \cdots) = \frac{2\delta}{1 - \delta}$$

- Therefore, upon observing a deviation, player i prefers to behave as prescribed by the GTS than deviating from it since

$$\frac{2}{1-\delta} \geq \frac{2\delta}{1-\delta} \Rightarrow \delta \leq 1, \text{ which holds by assumption since } \delta \in [0,1).$$

- Overall, the only condition that we need to support this GTS as SPE is $\delta \geq \underline{\delta}(p)$, as found above.