DS 574 Algorithmic Mechanism Design	**Repeated Games Worksheet**
Fall 2026	Prof. Kira Goldner

Repeated Games

In Prisoner's Dilemma, each player chooses whether to cooperate or defect. Each player gets a higher payoff by defecting, regardless of the other's choice, yet both would be better off if both cooperated than if both defected.

Payoffs are listed as (row player, column player).

	Column cooperates	**Column defects**
Row cooperates	(3, 3)	(0, 4)
Row defects	(4, 0)	(1, 1)

Fix a normal-form game, called the *one-round game*, and play it in rounds $t = 1, 2, 3, \ldots$. At the end of each round everyone observes the actions played.

A repeated-game strategy specifies play as a function of history. For each history, it says which action to choose next, or how to randomize over the available actions.

Definition 1. A *history* H_{t-1} records everyone's actions in rounds 1 through $t-1$; H_0 is empty. A (behavioral) *strategy* s_i specifies an action, or a distribution $s_i(H_{t-1})$ over actions, for each history H_{t-1}.

Definition 2. Players have a common *discount factor* $\delta \in (0, 1)$. Player i evaluates the payoffs u_i^t from successive rounds by

$$\sum_{t \geq 1} \delta^{t-1} u_i^t,$$

with expectation over any randomization. Payoffs in each round are bounded.

A payoff of 1 tomorrow is worth δ today; the same payoff two rounds from now is worth δ^2. Larger δ means that future payoffs matter more.
What is the value of receiving u every round? Of receiving d now and w every round starting tomorrow?

Subgame Perfection

A Nash equilibrium requires that no player can improve their payoff by changing their strategy, given the other players' strategies. However, it can rely on a threat that a player would not want to carry out if the time came. If nobody deviates, the threatened punishment never occurs. Changing what the player would do in that unreached situation therefore does not change their payoff, so the threat can be part of a Nash equilibrium even if carrying it out would be against their interests.

In contrast, in a subgame-perfect equilibrium, each player must want to carry out their part of a threatened punishment, given the other players' strategies, if the punishment is triggered.

Definition 3. A *subgame* is the repeated game remaining after a history H_{t-1}. A strategy profile is a *subgame-perfect equilibrium* (SPE) if it induces a Nash equilibrium after every history, including histories that never occur when everyone follows the profile.

Complete the comparison:

- Nash equilibrium checks deviations starting

- Subgame perfection checks deviations starting

Theorem 1 (One-shot deviation principle)**.** *In this discounted repeated game, a strategy profile is subgame perfect if and only if, after every history, no player can gain by changing their action for one round, then returning to their strategy in all later rounds.*

We take the principle as given. To check subgame perfection, consider deviations that last just one round. After every possible history, ask: could a player gain by choosing a different action for one round, then returning to their strategy for all future rounds? When answering, all players' strategies must respond to the history that actually occurs—including the changed action. Returning to a strategy does not necessarily mean returning to cooperation: if the changed action triggers punishment, the players follow their strategies for that punishment.

The Grim Strategy

We now return to the Prisoner's Dilemma from the beginning. Repeating this game with the same two players gives *Iterated Prisoner's Dilemma*. Defecting today might lead the other player to defect in future rounds. Can that prospect make cooperation worthwhile?

	Column cooperates	**Column defects**
Row cooperates	(3, 3)	(0, 4)
Row defects	(4, 0)	(1, 1)

Find the one-round equilibrium and each player's payoff there.

Both players would receive more from (Cooperate, Cooperate). To sustain that outcome, today's action must affect tomorrow's play. Consider the following *Grim strategy*:

For each history, compare payoffs from that round onward, treating it as today. Earlier payoffs are fixed, and the common discount factor multiplying all remaining payoffs does not change which continuation is better.

> **Exercise.** For which discount factors does this version of the Grim strategy sustain (Cooperate, Cooperate) as an SPE? Check both histories with a past deviation and histories without one.

What happens to this incentive comparison as $\delta \to 1$ and as $\delta \to 0$?

More generally, specify a target action profile and a fixed one-round Nash equilibrium as punishment. For each history, play the target if it has always been played, and the punishment otherwise. The argument uses the following steps.

Step 1: Find the punishment equilibrium. Check that repeating it makes the punishment credible.

Step 2: Name the target and its per-round payoff. Compute each player's payoff from conforming.

Step 3: Compute the best one-shot deviation. Hold the other players' current actions fixed.

Step 4: Write down conform $\geq$ deviate. Include the continuation payoff on each side.

Step 5: Solve for δ. Check the condition for every player.

General Threshold (Optional Extension)

> **Exercise.** Now fix a player and let the target payoff be R, the best one-round deviation payoff be D, and the punishment equilibrium payoff be P, where $D > R > P$. Derive the discount-factor threshold.

Holding D and R fixed, what happens when P decreases? Why can't we simply choose any small P we want?

	Column cooperates	**Column defects**
Row cooperates	(3, 3)	(0, 4)
Row defects	(4, 0)	(1, 1)

Finite Horizons

Suppose instead there are exactly T rounds, numbered $1, \ldots, T$, and everybody knows this.

Claim 1. If the one-round game has a unique Nash equilibrium (including mixed strategies), every SPE of its T-round repetition prescribes that equilibrium after every history.

Why must the same equilibrium be played after every history?

What does the claim imply for ten repetitions of our matrix? Which assumption would fail in a one-round game with two Nash equilibria?

Two Rounds, Multiple Equilibria

Now consider a different one-round game. The row and column players each choose A, B, or C. Payoffs are listed as (row player, column player).

	Column A	**Column** B	**Column** C
Row A	(2, 2)	(0, 3)	(0, 0)
Row B	(3, 0)	(1, 1)	(0, 0)
Row C	(0, 0)	(0, 0)	(3, 3)

Find all pure Nash equilibria of the one-round game. Is (A, A) a Nash equilibrium?

Exercise. Play this game for exactly two rounds, $t = 1, 2$, with payoff $u_i^1 + \delta u_i^2$ and $\delta = 3/4$. Construct an SPE in which both players choose A in round 1. Specify each player's round-2 action as a function of the round-1 history, and check both players' incentives in both rounds.

Which step of the finite-horizon proof no longer works?