REPEATED GAMES

Early PD experiments

In 1950, Merrill Flood and Melvin Dresher (at RAND) devised an experiment to test Nash's theory about defection in a two-person prisoners' dilemma.

Experimental Design

- They asked two friends to play the PD 100 times.
- They measured the success of Nash's equilibrium concept by counting the number of times the players chose {D;D}.

Flood and Dresher's results

Player 1 cooperated in 68 rounds

Player 2 cooperated in 78 rounds

Both cooperated in 60 of last 89 rounds

Flood

Dresher

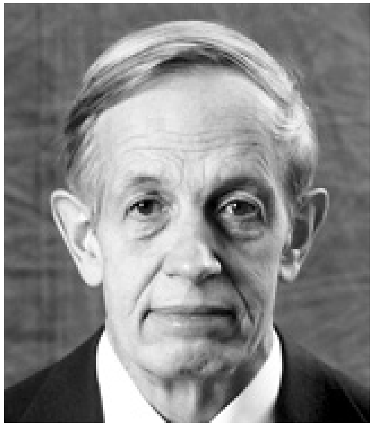

Nash

Flood and Dresher's results

Player 1 cooperated in 68 rounds

Player 2 cooperated in 78 rounds

Both cooperated in 60 of last 89 rounds

Flood Dresher Nash

Nash's response

"If this experiment were conducted with various different players rotating the competition and with *no information given to a player of what choices the others have been making until the end* of all trials, then the experimental results would have been quite different, for this modification of procedure would remove the interaction between the trials."

Nash's response

"The flaw in this experiment as a test of equilibrium point theory is that the experiment really amounts to having the players play **one large multimove game**. One cannot…think of the thing as a sequence of independent games…there is too much interaction."

In other words, Nash said that repeating the game changes the game itself. In which case, different equilibria may apply.

Repeated games

- A repeated game is sequential move game constructed from a (simultaneous move) base game. The base game is called a stage game (e.g., PD)

- Any stage game can be repeated (not just the PD). …We will study PD's here.

- Games can be repeated a finite or an infinite number of times.

…This matters.

Extra Credit

- You will be randomly paired with another member in the class. Consider yourself row.
- You will play the following stage game with your partner for exactly 5 rounds.

	C	D
C	1.5, 1.5	0, 3
D	3, 0	1, 1

- The total points you earn at the end of ten rounds will be added to your last homework assignment.
- Please put your name on a piece of paper and denote precisely how you will play each round. If there is any lack of clarity in your instructions, I will treat you as a “do not play.”

Repeated games

Length of Repetition

- Finite horizon ($T < \infty$)
  - Solve by backward induction
- Infinite horizon ($T = \infty$)
  - Cannot be solved by backward induction (since there is no end)

Goals of the analysis

- Does cooperation emerge if we repeat the PD? If so, under what conditions?
- What are the equilibria in a repeated PD?
- How do we analyze infinitely repeated games?
- Are there general results about repeated games?

Time preferences

Time preferences

- Key assumption: in many settings a payoff in the future is worth less than today.
- **Discount factor** $\delta \in (0, 1)$ parameterizes *patience*.
- Utility (present value at time t) of receiving X at time t+1 is δX.

Time preferences

Consider four periods of {C,C} in this PD ->

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Period (t)		1		2		3		4
payoff	=	*3*	*+*	*δ3*	*+*	*δ(δ3)*	*+*	*δ(δδ3)*

$$= \quad 3 + \delta^1 3 + \delta^2 3 + \delta^3 3$$

$$= \quad \sum_{t=1}^{4} \delta^{t-1} 3$$

Time preferences

Consider four periods of {C,C} in this PD ->

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Period (t)		1		2		3		4
payoff	=	*3*	*+*	*δ3*	*+*	*δ(δ3) +*		*δ(δδ3)*

$$= \quad 3 + \delta^1 3 + \delta^2 3 + \delta^3 3$$

$$= \quad \sum_{t=1}^{4} \delta^{t-1} 3$$

$$U_i(h) \quad = \quad \sum_{t=1}^{T} \delta^{t-1} u_i(x_t)$$

This is a general formula for finite repetition.

Time preferences

- Practice

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

1. What is the discounted utility for player 1 (row) in a 3 period repeat of the stage game above with play (D,D), (C,C), (D,C)? [hint: use δ].

2. What is the discounted utility for player 2 (column) in the same game from the same play?

Time preferences

- Practice

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

3. What is the discounted utility for player 1 (row) in a 20 period repeat of the stage game above with play (D,D), (C,C), followed by (D,C) for 18 rounds? [hint: use $\sum$ and δ].

Finite repetition

What is the sub-game perfect equilibrium (SPE)… in a finitely repeated PD (i.e., a PD repeated T times)?

Temporarily, assume δ =1…

Finitely repeated PD

The *unique SPE (and NE)* is for both players to defect after every possible history (i.e., “always defect”)

Intuition behind SPE (backward induction):

- In the last period of play, cooperation cannot engender reciprocity. Hence, (D,D) is rational in period T.
  - Because it is always better to defect in the stage game.
- If you know I’m going to defect in the last round and I know you are going to defect in the last round, then it is best for each of us to defect in T-1 because it cannot engender reciprocity in T.
- Similarly, (D,D) must be the action profile played at every period t < T-1.
- Anticipation makes cooperation “unvravel.”

Infinitely repeated game

Maybe we can engender cooperation if the game is played an infinite number of periods.

- After all, it was the last period that made defection rational and caused the game to unravel.

Geometric Progression

Consider a *constant* payoff of c for T finite periods:

$$S_T = \sum_{t=1}^{T} \delta^{t-1} c = c\left(1+\delta+\delta^2+\ldots+\delta^{T-1}\right)$$

We now use a trick to simplify the above equation. Note…

$$\delta S_T = c\delta\left(1+\delta+\delta^2+\delta^3+\ldots+\delta^{T-1}\right)$$

$$S_T - \delta S_T = c\left(1+\delta \quad +\delta^2+\ldots+\delta^{T-1}\right) - c\left(\delta+\delta^2+\delta^3+\ldots+\delta^{T}\right)$$

$$S_T = \frac{c(1-\delta^T)}{1-\delta}$$

For infinite periods: As $T \rightarrow \infty$,

$\delta^T \rightarrow 0$ and for $T = \infty$

$S_T = c / (1 - \delta)$.

Time preferences

Discounted sum of streams of constant payoff c:

$$\sum_{t=1}^{T} \delta^{t-1} c = \frac{c(1+\delta^{T})}{1-\delta} \qquad \sum_{t=1}^{\infty} \delta^{t-1} c = \frac{c}{1-\delta}$$

Mathematically, this is a **geometric series**, so discounting each future period by a constant discount factor of δ is called **geometric discounting**.

Cardinality matters (just like it did for expected utility)

Strategies

- A strategy specifies an action for *every* period of the game.
- In an infinitely repeated game, the set of strategies is infinite.
- We will restrict attention to a few strategies that are easy to describe:
  - Always defect – D in every period.
  - Always cooperate – C in every period.
  - Grim trigger: cooperate in first period, defect forever if other player has defected in a previous period.
  - Tit-for-tat: cooperate in first period, copy other player’s action in next period.

Nash equilibrium and SPE

- Sequential Equililbrium
  - How does one apply backward induction to a game that has no end?
  - Answer: you don't. Hence you would study sequential equilibria (i.e. sub-game perfect equilibria) differently.

- We will focus on Nash equilibrium
  - Because analyzing sub-game perfect equilibria in repeated games does not give us any additional insights. Furthermore, N.E. are much easier.

- Nash equilibrium
  - Set of strategies such that no player has an incentive to deviate
  - Check for deviations from something we suspect is Nash.

Useful Formulas

Formula:	$\boxed{x^a\ x^b = x^{(a+b)}}$	$\boxed{x^a / x^b = x^{(a-b)}}$
Examples:	$\delta^3\ \delta^2 = \delta^5$	$\delta^5 / \delta^2 = \delta^3$
	$\delta^6\ \delta^1 = \delta^7$	$\delta^8 / \delta^4 = \delta^4$

Steps in Equilibrium Analysis

1. Determine the play implied by the stated strategies.
2. Compute discounted sum of payoffs.
3. Find best possible deviation for one player (usually all defect, or defect in first period).
4. Set up the Nash equilibrium condition (i.e., the inequality needed for deviation to be rational) and solve to determine if there is a feasible value of δ (between 0 and 1), where equilibrium can be sustained.

Always defect

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Assume common discount factor δ from here forward.

- Step 1:
  - Player 1: D, D, D, …
  - Player 2: D, D, D, …
  - Payoffs 1: 1, 1δ, 1δ^2, …
  - Payoffs 2; 1, 1δ, 1δ^2, …
- Step 2:
  - Sum of payoffs: c / (1 - δ) = 1 / (1 - δ).

Always defect

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

Assume common discount factor δ from here forward.

- Step 3:
  - Any deviation from (all D, all D) leads to a lower payoff in the deviating period. Hence, there is no rational deviation.
- Steps 4: skip
- Conclude
  - This is a NE because there is no incentive to unilaterally deviate to another (repeated) strategy.

Grim trigger (GT)

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

C in first period.
C as long as other plays C.
D forever if other plays D in any round.

- Step 1
  - Player 1: C, C, C, …
  - Player 2: C, C, C, …
  - Payoffs 1: 3, 3δ, $3\delta^2$, …
  - Payoffs 2; 3, 3δ, $3\delta^2$, …
- Step 2
  - Sum of payoffs: $c / (1 - \delta) = 3 / (1 - \delta)$.

Grim trigger (GT)

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

C in first period.
C as long as other plays C.
D forever if other plays D in any round.

- Step 3
  - If player 1 deviates to “always D” (or identically grim trigger with D in the first round), then the two will get:
    - Player 1: D, D, D, …
    - Player 2: C, D, D, …

Grim trigger (GT)

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

C in first period.
C for any history such that no player has ever played D.
D if either player has ever played D.

- **Step 4: It is rational for player 1 to deviate to "always D" iff:**

$\underline{EU_1\text{(always D)} \qquad > \quad EU_1\text{(GT, GT)}}$

$$
\begin{aligned}
5 + 1(\delta + \delta^2 + \delta^3 + ...) &> \frac{3}{(1-\delta)} \\
5 + \delta(1 + \delta + \delta^2 + ...) &> \frac{3}{(1-\delta)} \\
5 + \sum_{t=1}^{\infty} \delta^{t-1} &> \frac{3}{(1-\delta)} \\
5 + \frac{\delta}{(1-\delta)} &> \frac{3}{(1-\delta)} \\
5 - 5\delta + \delta &> 3 \\
2 &> 4\delta \\
\delta &< 1/2
\end{aligned}
$$

Conclude: If δ < ½, then this deviation (and other deviations) are rational.

If δ ≥ ½, then (GT, GT) is a Nash Equilibrium, generating the outcome (C,C) in every period.

Why? Because deviating to "always defect," in a later period produces the same condition. See attached.

Always cooperate

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

C all periods.

- Step 1
  - Player 1: C, C, C, …
  - Player 2: C, C, C, …
  - Payoffs 1: $3, 3\delta, 3\delta^2, \ldots$
  - Payoffs 2; $3, 3\delta, 3\delta^2, \ldots$
- Step 2
  - Sum of payoffs: $c / (1 - \delta) = 3 / (1 - \delta)$.

Always cooperate

Stage game

	C	D
C	3, 3	0, 5
D	5, 0	1, 1

C all periods.

- Step 3
  - If player 1 deviates to "always D," then he will get $5 / (1 - \delta)$.
- Step 4
  - This deviation is rational if $5 / (1 - \delta) > 3 / (1 - \delta)$, which is true for all δ.
- Conclude
  - Hence, {always C; always C} is <u>not</u> a N.E. for any value of δ.

Intuition

- Cooperation requires
  1. *Infinite horizon*.
  2. *Threat of future punishment*.
     - If you won't punish defection (like always C), then your partner will defect on you.
  3. *Sufficient patience* among both players.
     - long-term gain from cooperating must exceed short-term gain from defection minus the long-term cost of defection.

Steps in Equilibrium Analysis

1. Determine the play implied by the stated strategies.
2. Compute discounted sum of payoffs.
3. Find best possible deviation for one player (usually all defect, or defect in first period).
4. Set up the Nash equilibrium condition (i.e., the inequality needed for deviation to be rational). Solve to determine if there is a feasible value of δ (between 0 and 1), where equilibrium can be sustained.

Tit for tat

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

- Start with C
- Play C if other player played C in previous period
- Play D if other player played D in previous period

U_1(TFT,TFT) = ?

Practice: Do first two previous steps on this PD (new payoffs).

Tit for tat

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

U_1(TFT,TFT) = $\frac{2}{(1-\delta)}$

- Start with C
- Play C if other player played C in previous period
- Play D if other player played D in previous period

Practice: Do first two previous steps on this PD (new payoffs).

Step 1:

- Player 1: C, C, C, …
- Player 2: C, C, C, …

Step 2:

- Payoff 1: $2 + 2\delta + 2\delta^2 + \ldots = \frac{2}{(1-\delta)}$

Tit for tat

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

- Start with C
- Play C if other player played C in previous period
- Play D if other player played D in previous period

$U_1(\text{TFT,TFT}) = \frac{2}{(1-\delta)}$

Practice: Do third step on this PD.

Step 3:

Consider deviation "always defect"

Player 1: D, D, D, …

Player 2: C, D, D, …

Payoff 1: $3 + \delta + \delta^2 + \delta^3 \ldots$

Tit for tat

Consider Deviation Always defect.

Payoff 1: $3 + (\delta + \delta^2 + \delta^3 \ldots)$

$3 + \delta(1 + \delta + \delta^2 \ldots)$

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

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

Tit for tat

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

- Start with C
- Play C if other player played C in previous period
- Play D if other player played D in previous period

$$U_1(TFT, TFT) = \frac{2}{1-\delta}$$

Step 4: this is a rational deviation iff:

$$3 + \frac{\delta}{(1-\delta)} > \frac{2}{(1-\delta)}$$

$$(3 - 3\delta) + \delta > 2$$

$$3 - 2\delta > 2$$

$$\delta < 1/2$$

Conclude: If the players value the future moderately, δ > ½, (TFT,TFT) is a N.E.

Conclude: If the players don't value the future moderately, δ < ½, cooperation **cannot** be sustained.

Tit for tat

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

U_1(TFT,TFT) = $\frac{2}{(1-\delta)}$

- Start with C
- Play C if other player played C in previous period
- Play D if other player played D in previous period

Why is looking at deviation in the first round sufficient for the case of TFT against TFT?

Cooperation in infinitely repeated PD

- Cooperation along the equilibrium path of play can be supported by several different strategy profiles.

- Cooperation is supported by the threat of punishment and a sufficient level of patience.

  - Note: (all C, all C) is not an equilibrium strategy. Even a nice strategy must be able to punish.

- The level of patience required is smaller if punishment is more severe (e.g., grim trigger requires less patience, TFT requires more patience).

Alternate equilibrium path

- Instead of (C,C) in every period, is there a NE where the players alternate between (D,C) and (C,D)?
- Consider an alternating grim trigger set of strategies (*AltGT*):
  - Play (D,C) in odd number periods, play (C,D) in even number periods
  - If either player deviates from this path of play, they will play D forever.

Stage game

	C	D
C	2, 2	0, 3
D	3, 0	1, 1

$$U_1(AltGT, AltGT) = 3 + 0\delta + 3\delta^2 + 0\delta^3 + \ldots$$

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

$$U_2(AltGT, AltGT) = 0 + 3\delta + 0\delta^2 + 3\delta^3 + \ldots$$

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

Alternate equilibrium path

Since Player 1 gets highest payoff in period 1, consider deviation to D in period 2

$$\begin{aligned} U_1(Dev, AltGT) &= 3 + \delta + \delta^2 + \delta^3 + \ldots \\ &= 3 + \frac{\delta}{(1-\delta)} \end{aligned}$$

Player 1 has no incentive to deviate if

$$U_1(AltGT, AltGT) \geq U_1(Dev, AltGT)$$

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

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

Alternate equilibrium path

Since Player 2's best deviation is to start playing D in period 1

$$U_2(Dev, AltGT) = 1 + \delta + \delta^2 + \delta^3 + \ldots$$
$$= \frac{1}{1-\delta}$$

Player 2 has no incentive to deviate if

$$U_2(AltGT, AltGT) \geq U_2(Dev, AltGT)$$
$$\frac{3\delta}{1-\delta^2} \geq \frac{1}{1-\delta}$$
$$\delta \geq \frac{1}{2}$$

Alternate equilibrium paths

- Thus, for the stage game with the payoffs given, there is a Nash equilibrium where players alternative between (D,C) and (C,D) along the equilibrium path.
  - Note: in this case they are playing different strategies.
- This suggests that outcomes other than full cooperation or full defection can be supported in equilibrium as well.

Remarks

The folk theorem (which we did not introduce) tells us that in infinitely repeated games there is a *multiplicity of equilibria* – we cannot make sharp empirical predictions.

In the PD, cooperation is sustainable in equilibrium—but it is *not the only possible outcome*. All defect is in equilibrium against all defect as well, and mixed cooperation and defection strategies are also in equilibrium.

The folk theorem tells us which *payoffs* are supportable in *some* Nash equilibrium. It does *not* tell us anything the actual strategy profiles that might be used.

Therefore, repeated play might engender cooperation, but it might not.

Application: Green House Gasses

- Why is it so difficult to get countries to agree to reduce green house gasses?
  - No country has incentive to reduce its own emissions, because doing so alone requires a significant cost with little benefit in terms of climate change.

- Consider a PD game between us and them.

		THEM	
		Cut emissions	Don’t
US	Cut emissions	-1, -1	-20, 0
	Don’t	0, -20	-12, -12

  - The British government suggests that coordinated effort may come at a cost of 1% of GDP / nation, whereas inaction could cost roughly 12% of GDP / nation. Dixit, Skeath, and Reiley extrapolate the remaining payoffs.

Application: Green House Gasses

Michael Liebriech, argues that despite the Kyoto agreement, this is not a one shot interaction. Countries repeatedly interact and negotiate additional amendments to the existing agreement, making it a repeated game.

1. Do you think international behavior towards Global Warming can be explained by a repeated PD?

2. Does Kyoto and other accords provide for the type of retaliatory behavior that is required for cooperation to be an equilibrium in a repeated game?

3. What types of discount factors should we expect in such a game? Note, countries may interact roughly infinitely, but term limits suggest that politicians have shorter time horizons.

4. Even though repeated play can lead to cooperative behavior, we have seen that it can also lead to non-cooperative behavior as well? What guarantees that we will get to a cooperative equilibrium?

5. How can we help countries coordinate on the cooperative equilibria?