

15. Repeated Games

Duarte Gonçalves

University College London

MRes Microeconomics

Repeated Interaction

Repeated strategic interaction:

- (i) Firms setting prices.
- (ii) Workers exerting effort at their job.
- (iii) Students learning the materials.
- (iv) Politicians running for reelection.
- (v) Viruses mutating.
- (vi) Investors adjusting their portfolios.

What can we say about repeated strategic interactions?

How do predictions change when players interact multiple times?

Overview

1. Repeated Interaction
2. Repeated Games
3. Characterising Nash Equilibria
4. Characterising Subgame-Perfect Nash Equilibria
5. Applications
6. More

Overview

1. Repeated Interaction

2. Repeated Games

- Setup
- Interpreting the Discount Factor
- Average Discounted Payoffs

3. Characterising Nash Equilibria

4. Characterising Subgame-Perfect Nash Equilibria

5. Applications

6. More

Example

Firm pricing

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > 4/7 p_H > 0$.

Unit mass of consumers.

$1/8$ consumers are loyal to brand A and $1/8$ to B .

$3/4$ consumers buy from cheapest firm; ties broken uniformly at random.

$$\pi_i(p_i, p_j) = p_i \left(\frac{1}{8} + \mathbf{1}_{p_i < p_j} \frac{6}{8} + \mathbf{1}_{p_i = p_j} \frac{3}{8} \right).$$

Stage game:

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

One-shot interaction: unique NE (L, L) .

but... firms face repeated pricing problem.

Can firms collude in pricing higher?

Setup

Repeated game consists of players repeatedly facing a **stage game** Γ .

For simplicity, we take Γ as a normal-form game:

$\Gamma := \langle I, A, \pi \rangle$, where

I : set of players; A_i : player i 's set of available actions, and $A := \times_{i \in I} A_i$;

$\pi_i : A \rightarrow \mathbb{R}$ player i 's bounded stage-game vNM utility function, with $\pi = (\pi_i)_{i \in I}$.

$\pi_i : \times_{i \in I} \Delta(A_i) \rightarrow \mathbb{R}$, where $\alpha_i \in \Delta(A_i)$ is mixed action.

Setup

Repeated game (or 'supergame'): a stage game Γ , repeated for $T \in \mathbb{N} \cup \{\infty\}$ periods, and a discounting parameter δ .

- Focus on games of **perfect monitoring**: at the end of every period t , every player observes the actions chosen by all players in that period.

(When this is not the case, repeated game is of **imperfect monitoring** or unobservable actions.)

Histories:

At period t , each player observed **history** of actions $h^t = \{a^1, \dots, a^{t-1}\} \in A^t$, where $a^t = (a_i^t)_{i \in I}$.

H^t : **set of all possible histories at period t** , $H^t = A^{t-1}$, $H^1 = \{\emptyset\}$.

$H := \cup_{t=1}^T H^t$ is **set of all possible histories in the repeated game**.

NB: Here, slightly different definition of histories; allow for simultaneous actions.

Also, history = info set.

Setup

Repeated game (or ‘supergame’): a stage game Γ , repeated for $T \in \mathbb{N} \cup \{\infty\}$ periods, and a discounting parameter δ .

Strategies:

Pure strategy: s_i for player i determines an action at any possible history/info set, i.e., $s_i : H \rightarrow A_i$. (as before).

For convenience, we focus on **behavioural strategies** (instead of mixed), writing $\sigma_i : H \rightarrow \Delta(A_i)$

Payoffs:

Strategy profile induces distribution over actions in each period.

At each period t , an **(expected) stage-game payoff** $\pi_i(\sigma(h^t))$ given history h^t .

Aggregate payoffs for supergame via exponential discounting with δ .

(Expected) discounted payoff for the repeated game given σ (inducing a distribution over histories):

$$u_i(\sigma) := \sum_{t=1}^T \delta^{t-1} \mathbb{E}_{\sigma}[\pi_i(\sigma(h^t))] = \sum_{t=1}^T \delta^{t-1} \sum_{h^t \in H^t} \mathbb{P}_{\sigma}(h^t) \pi_i(\sigma(h^t)).$$

(looks more complicated than it is)

If $T = \infty$, need $\delta \in [0, 1)$ for payoffs to be well-defined.

Interpretation of the Discount Factor

Discounting parameter $\delta \in [0, 1)$ has two interpretations

- (i) Player's time preference.
(can have player-specific δ_i , just makes it more complicated)
- (ii) Prob. game continues at every given period (generating an exponential distrib.).
- (iii) Not mutually exclusive: if players exponentially discount future periods at rate ρ and game continues to next period wp λ , then $\delta = \rho\lambda$ (insofar as $\rho\lambda \in [0, 1)$).

Repeated Game Γ_δ^T : stage game Γ repeated T periods, with discount factor δ .

- For any history h which is subhistory of h' , $\Gamma_\delta^T(h)$ is a proper subgame of $\Gamma_\delta^T(h')$.
- Every period we have a new proper subgame.

Average Discounted Payoffs

Average discounted payoff

$$\tilde{u}_i(\sigma) := \frac{(1 - \delta)}{1 - \delta^T} u_i(\sigma) = \frac{(1 - \delta)}{1 - \delta^T} \sum_{t=1}^T \delta^{t-1} \mathbb{E}_{\sigma}[\pi_i(\sigma(h^t))].$$

Convenient normalisation in infinitely repeated games.

If s specifies every period same action profile played on-path $a_t = a \forall t$,
then $\tilde{u}_i(s) = \pi_i(a)$.

$\tilde{u}_i$ bounded when π_i is bounded, even if $T = \infty$:

$\tilde{u}_i(\sigma) \in \text{co}(\{\pi_i(a), a \in A\}) \implies$ Allows to take $\delta \rightarrow 1$ (limit of patient players).

Overview

1. Repeated Interaction
2. Repeated Games
3. Characterising Nash Equilibria
 - Feasible and Individually Rational Payoffs
 - Folk Theorems for Nash Equilibria
4. Characterising Subgame-Perfect Nash Equilibria
5. Applications
6. More

Example

Firm pricing

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > \frac{4}{7}p_H > 0$.

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

One-shot interaction: unique NE (L, L) .

but... firms face repeated pricing problem.

Can firms collude in pricing higher?

Characterising Possible Equilibria

Set of equilibria in repeated games may expand relative to the set of equilibria in stage games (played repeatedly).

Players can now make use of future gameplay to punish deviations from a target behaviour.

What sort of equilibrium play can be sustained in repeated games?

Characterise equilibria

1st: Nash equilibria; later: subgame-perfect Nash equilibria.

Nash Equilibria

Characterising Nash equilibria: no player has an incentive to deviate from the *equilibrium path of play*.

Only need to check for deviations at histories that are reached with positive probability.

Still: lots of deviations and complicated to characterise all equilibria

Idea: look at equilibrium payoffs.

Definition

The set of **stage-game payoffs** are given by $\tilde{V} := \{(\pi_i(a))_{i \in I} \mid a \in A\}$.

The set of **feasible payoffs** are given by $V^F := \text{co}(\tilde{V})$.

Feasible payoffs: $v \in V^F \iff \exists p \in \Delta(A) : \forall i, v_i = \mathbb{E}_p[\pi_i(a)]. \quad \forall \sigma, u_i(\sigma) \in V.$


Interpretation:

- Randomisation with a public device: Players coordinate with an observable random numbers generator that produces the distribution p .
Many natural public randomisation devices: rain outcome, first n characters of a newspaper, etc.
- Alternating actions: order that players agreed upon before play starts, yielding average payoff very close to v for large T and δ close to 1.

Write $\pi(\alpha) := (\pi_i(\alpha))_{i \in I}$.

More Definitions: Feasible Payoffs

		Col	
		H	L
Row	H	4, 2	0, 3
	L	3, 0	2, 4



Definition

Minmax payoff of player i 's for the stage game Γ is

$$\underline{v}_i := \min_{\alpha_{-i}} \max_{\alpha_i} \pi_i(\alpha_i, \alpha_{-i}).$$

Minmax payoff = highest payoff i can get if everyone else coordinates on minimising i 's payoff.

- Minmax actions can be used to punish player who deviates from equilibrium path.
- Minmax payoffs set bounds for avg payoffs that we can expect:
player i can always secure at least $\underline{v}_i$.


Definition

$v \in \mathbb{R}^{|I|}$ is **strictly individually rational** if $\forall i, v_i > \underline{v}_i$.

The set of **strictly individually rational payoffs** are given by $V^{SIR} := \{v_i > \underline{v}_i \forall i\}$.

More Definitions: Feasible and IR Payoffs

		Col	
		H	L
Row	H	4, 2	0, 3
	L	3, 0	2, 4



Folk Theorem for NE in Finitely Repeated Games

Theorem

Let Γ be a stage game with a Nash equilibrium α s.t. $\pi(\alpha) \in V^{SIR}$. Let $v \in V^F \cap V^{SIR}$. For every $\varepsilon > 0$, $\exists \delta^* \in (0, 1)$ s.t., $\forall \delta \in (\delta^*, 1)$, Γ_δ^T has a Nash equilibrium σ for which $|\tilde{u}_i(\sigma) - v_i| < \varepsilon$, for any finite $T > T^*(\delta)$.

Translation:

If (1) $\exists$ SIR NE, (2) stage game is repeated for long enough, and (3) players are sufficiently patient, then every feasible and SIR payoff vector can be approx. attained as the eqm avg. discounted payoff by some NE of Γ_δ^T .

Folk Theorem for NE in Finitely Repeated Games

Theorem

Let Γ be a stage game with a Nash equilibrium α with strictly individually rational equilibrium payoffs, $\pi(\alpha) \in V^{SIR}$. Let $v \in V^F \cap V^{SIR}$.

For every $\varepsilon > 0$, $\exists \delta^* \in (0, 1)$ s.t., $\forall \delta \in (\delta^*, 1)$, Γ_δ^T has a Nash equilibrium σ for which $|\tilde{u}_i(\sigma) - v_i| < \varepsilon$, for any finite $T > T^*(\delta)$.

Proof idea

Intuition: coordinate on prob. p that generates v for T_0 periods and switch to playing α at periods $t > T_0$; if anyone deviates, everyone else 'minmaxes' that player forever.

(i) $T_0, T \implies$ approximate v well enough and make it SIR.

$$|(1 - \delta^{T_0})v_i + (\delta^{T_0} - \delta^T)\pi_i(\alpha) - (1 - \delta^T)v_i| < \varepsilon \iff (\delta^{T_0} - \delta^T)|\pi_i(\alpha) - v_i| < \varepsilon$$

Folk Theorem for NE in Finitely Repeated Games

Theorem

Let Γ be a stage game with a Nash equilibrium α with strictly individually rational equilibrium payoffs, $\pi(\alpha) \in V^{SIR}$. Let $v \in V^F \cap V^{SIR}$.

For every $\varepsilon > 0$, $\exists \delta^* \in (0, 1)$ s.t., $\forall \delta \in (\delta^*, 1)$, Γ_δ^T has a Nash equilibrium σ for which $|\tilde{u}_i(\sigma) - v_i| < \varepsilon$, for any finite $T > T^*(\delta)$.

Proof idea

Intuition: coordinate on prob. p that generates v for T_0 periods and switch to playing α at periods $t > T_0$; if anyone deviates, everyone else 'minmaxes' that player forever.

(i) $T_0, T \implies$ approximate v well enough and make it SIR.

(ii) δ large $\implies$ players patient $\implies$ punishment effective in deterring deviations.

Payoff from adhering to strategy: $v_i - \varepsilon$.

Payoff from deviating at $t < T_0$: $(1 - \delta^t)v_i + \delta^t(1 - \delta) \max_{a_i} \pi_i(a_i, p_{-i}) + \delta^t \delta(1 - \delta^{T-t}) \underline{v}_i$.

NE if $(1 - \delta) \max_{a_i} \pi_i(a_i, p_{-i}) + \delta(1 - \delta^{T-t}) \underline{v}_i < v_i - \delta^{-t} \varepsilon$. $\square$

Infinitely Repeated Games

What if the game has no final period?

Often the game doesn't have a fixed terminal horizon.

Intuition suggests: more repetitions $\implies$ more opportunities for punishment
 $\implies$ can do more than in finitely repeated games.

Recall: key part of folk theorem for finitely repeated games was T large enough.

Folk Theorem for NE in Infinitely Repeated Games

Theorem

Let Γ be a stage game with a Nash equilibrium α with strictly individually rational equilibrium payoffs, $\pi(\alpha) \in V^{SIR}$. Let $v \in V^F \cap V^{SIR}$.

$\exists \delta^* \in (0, 1)$ s.t., $\forall \delta \in (\delta^*, 1)$, Γ_δ^∞ has a Nash equilibrium σ that yields average discounted payoffs $\tilde{u}(\sigma) = v$.

Translation:

If (1) $\exists$ SIR NE, (2) stage game is repeated for long enough, and (3) players are sufficiently patient, then every feasible and SIR payoff vector can be approx. attained as the eqm avg. discounted payoff by some NE of Γ_δ^∞ .

Folk Theorem for NE in Infinitely Repeated Games

Theorem

Let $v \in V^F \cap V^{SIR}$. $\exists \delta^* \in (0, 1)$ s.t., $\forall \delta \in (\delta^*, 1)$, Γ_δ^∞ has a Nash equilibrium σ that yields average discounted payoffs $\tilde{u}(\sigma) = v$.

Proof idea

Intuition: coordinate on prob. p that generates v ; if anyone deviates, everyone else 'minmaxes' that player forever.

If there are multiple players deviating, then fix one to be punished.

Only need to check on-path deviations (NE not SPNE).

Payoff from adhering to strategy: v_i .

Payoff from deviating at h^t : $(1 - \delta) \max_{a_i} \pi_i(a_i, p_{-i}) + \delta \underline{v}_i$.

NE if $\delta \geq \frac{\max_{a_i} \pi_i(a_i, p_{-i}) - v_i}{\max_{a_i} \pi_i(a_i, p_{-i}) - \underline{v}_i}$.

□

Example

Firm pricing

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > 4/7 p_H > 0$.

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

One-shot interaction: unique NE (L, L) .

but... firms face repeated pricing problem.

Can firms collude in pricing higher?

Payoffs $(p_H/2, p_H/2)$ attainable as NE of infinitely repeated game if players sufficiently patient, and (H, H) as on-path eqm play.

Overview

1. Repeated Interaction
2. Repeated Games
3. Characterising Nash Equilibria
4. Characterising Subgame-Perfect Nash Equilibria
 - One-Shot Deviation Principle
 - SPNE in Repeated Games
 - Folk Theorems for SPNE
5. Applications
6. More

Subgame-Perfect Nash Equilibrium

NE: no player wants to deviate.

SPNE: no player wants to deviate *in any subgame* (i.e., following any history).

Rule out non-credible threats. At every period, following any history, players can re-optimize strategies.

A lot of deviations to check!

Even if A is finite, dimensionality of player i 's *pure* strategy space can be very large!

$$|A_i|^{\sum_{t=1}^T (\prod_{j \in I} |A_j|)^{t-1}}.$$

With two players – two actions stage game, the size of the pure strategy space of each player is 2^{2^T} ; grows very very fast with time horizon.

Questions:

How can one characterise equilibrium strategies?

What can be achieved in equilibrium?

One-Shot Deviation Principle

Intuition: need but consider one deviation at a single information set. (In repeated games, following a single history, since new subgame follows every non-terminal history.)

Continuation play $\sigma_i |_{h^t}$: strategy σ_i restricted to subgame that ensues history h^t .

Continuation payoffs $u_i |_{h^t}$: payoffs from stage games following history h^t .

$$u_i |_{h^t} (\sigma |_{h^t}) := \sum_{\ell=1}^{T-t} \delta^{\ell-1} \mathbb{E}_{\sigma |_{h^t}} [\pi(\sigma |_{h^t} (h^{t+\ell}))].$$

One-Shot Deviation Principle

Definition

Fix a strategy profile σ . A **profitable one-shot deviation** for player i is strategy $\sigma'_i \neq \sigma_i$ for which $\exists!$ information set h^t s.t.

- (i) $\forall h \in \mathcal{H}_i \setminus \{h^t\}, \sigma'_i(h) = \sigma_i(h)$; and
- (ii) $u_i |_{h^t} (\sigma'_i |_{h^t}, \sigma_{-i} |_{h^t}) > u_i |_{h^t} (\sigma |_{h^t})$.

Remark

- (1) Deviation at a single info set (a one-shot deviation) can induce completely different *path of play* from then on; behaviour at histories subsequent to h^t can depend on what was played at h^t .
- (2) Information set h^t may never be reached given σ ; deviation is *conditional* on h^t .

One-Shot Deviation Principle

Intuition: need but consider one deviation at a single information set. (In repeated games, following a single history, since new subgame follows every non-terminal history.)

Continuation play $\sigma_i |_{h^t}$: strategy σ_i restricted to subgame that ensues history h^t .

Continuation payoffs $u_i |_{h^t}$: payoffs from stage games following history h^t .

$$u_i |_{h^t} (\sigma |_{h^t}) := \sum_{h \in [T-t]} \delta^{h-1} \mathbb{E}_{\sigma |_{h^t}} [\pi(\sigma |_{h^t} (h^{t+h}))].$$

Proposition (One-Shot Deviation Principle)

σ is SPNE if and only if there are no profitable one-shot deviations.

One-Shot Deviation Principle

Proof

Profitable one-shot deviations $\implies$ Not SPNE. Immediate.

No profitable one-shot deviations $\implies$ SPNE. Less so.

- (1) σ not SPNE $\implies \exists$ profitable deviation for player i , $\hat{\sigma}_i$, in some subgame following history h^t $u_i|_{h^t}(\hat{\sigma}_i, \sigma_{-i}) - u_i|_{h^t}(\sigma_i, \sigma_{-i}) =: \varepsilon > 0$. (drop $|_{h^t}$ for simplicity)
- (2) Create auxillary strategy $\tilde{\sigma}_i$ s.t.
 - (a) $\tilde{\sigma}_i(h^{t+\ell}) = \hat{\sigma}_i(h^{t+\ell}) \quad \forall h^{t+\ell}$ following h^t up to some large period T ($\ell = 0, 1, \dots, T$);
 - (b) $\tilde{\sigma}_i(h^{t-\ell}) = \hat{\sigma}_i(h^{t-\ell}) = \sigma_i(h^{t-\ell}) \quad \forall$ proper subhistory $h^{t-\ell}$ of h^t ($\ell = 1, 2, \dots, t$); and
 - (c) $\tilde{\sigma}_i(h^{T+\ell}) = \sigma_i(h^{T+\ell}) \quad \forall$ history $h^{T+\ell}$ from period $T+1$ onward ($\ell = 1, 2, \dots$).

If the game is finite, T can be the total number of repetitions.

If game is infinitely repeated, choose T large enough s.t. $u_i(\hat{\sigma}_i, \sigma_{-i}) - u_i(\tilde{\sigma}_i, \sigma_{-i}) < \varepsilon/2$.

One-Shot Deviation Principle

Proof

- (1) Create auxillary strategy $\tilde{\sigma}_i$ s.t.
- (a) $\tilde{\sigma}_i(h^{t+\ell}) = \hat{\sigma}_i(h^{t+\ell}) \quad \forall h^{t+\ell}$ following h^t up to some large period T ($\ell = 0, 1, \dots, T$);
 - (b) $\tilde{\sigma}_i(h^{t-\ell}) = \hat{\sigma}_i(h^{t-\ell}) = \sigma_i(h^{t-\ell}) \quad \forall$ proper subhistory $h^{t-\ell}$ of h^t ($\ell = 1, 2, \dots, t$); and
 - (c) $\tilde{\sigma}_i(h^{T+\ell}) = \sigma_i(h^{T+\ell}) \quad \forall$ history $h^{T+\ell}$ from period $T+1$ onward ($\ell = 1, 2, \dots$).
- (2) If game is infinitely repeated, choose T large enough so that $u_i(\hat{\sigma}_i, \sigma_{-i}) - u_i(\tilde{\sigma}_i, \sigma_{-i}) < \varepsilon/2$.

To see this, denote $\hat{\sigma} = (\hat{\sigma}_i, \sigma_{-i})$ and $\tilde{\sigma} = (\tilde{\sigma}_i, \sigma_{-i})$. Then,

$$\begin{aligned} |u_i(\hat{\sigma}_i, \sigma_{-i}) - u_i(\tilde{\sigma}_i, \sigma_{-i})| &= \left| \sum_{\ell=1}^{\infty} \delta^{T+\ell-1} \sum_{h^{T+\ell} \in H^{T+\ell}} \mathbb{P}_{\hat{\sigma}}(h^{T+\ell}) \pi_i(\hat{\sigma}(h^{T+\ell})) - \mathbb{P}_{\tilde{\sigma}}(h^{T+\ell}) \pi_i(\tilde{\sigma}(h^{T+\ell})) \right| \\ &\leq \sum_{\ell=1}^{\infty} \delta^{T+\ell-1} \left| \sum_{h^{T+\ell} \in H^{T+\ell}} \mathbb{P}_{\hat{\sigma}}(h^{T+\ell}) \pi_i(\hat{\sigma}(h^{T+\ell})) - \mathbb{P}_{\tilde{\sigma}}(h^{T+\ell}) \pi_i(\tilde{\sigma}(h^{T+\ell})) \right| \\ &\leq \sum_{\ell=1}^{\infty} \delta^{T+\ell-1} 2 \max_{\alpha} |\pi_i(\alpha)| = \frac{\delta^T}{1-\delta} 2 \max_{\alpha} |\pi_i(\alpha)| < \varepsilon/2 \\ &\iff \frac{\ln(4/((1-\delta)\varepsilon) \max_{\alpha} \pi_i(\alpha))}{-\ln(\delta)} < T. \end{aligned}$$

One-Shot Deviation Principle

Proof

No profitable one-shot deviations $\implies$ SPNE. Immediate.

No profitable one-shot deviations $\longleftarrow$ SPNE:

- (1) σ not SPNE $\implies \exists$ profitable deviation for player i , $\hat{\sigma}_i$, in some subgame following history h^t $u_i |_{h^t}(\hat{\sigma}_i, \sigma_{-i}) - u_i |_{h^t}(\sigma_i, \sigma_{-i}) =: \varepsilon > 0$. (drop $|_{h^t}$ for simplicity)
- (2) Create auxillary strategy $\tilde{\sigma}_i$ s.t.
 - (a) $\tilde{\sigma}_i(h^{t+\ell}) = \hat{\sigma}_i(h^{t+\ell}) \quad \forall h^{t+\ell}$ following h^t up to some large period T ($\ell = 0, 1, \dots, T$);
 - (b) $\tilde{\sigma}_i(h^{t-\ell}) = \hat{\sigma}_i(h^{t-\ell}) = \sigma_i(h^{t-\ell}) \quad \forall$ proper subhistory $h^{t-\ell}$ of h^t ($\ell = 1, 2, \dots, t$); and
 - (c) $\tilde{\sigma}_i(h^{T+\ell}) = \sigma_i(h^{T+\ell}) \quad \forall$ history $h^{T+\ell}$ from period $T+1$ onward ($\ell = 1, 2, \dots$).

If the game is finite, T can be the total number of repetitions.

If game is infinitely repeated, choose T large enough s.t. $u_i(\hat{\sigma}_i, \sigma_{-i}) - u_i(\tilde{\sigma}_i, \sigma_{-i}) < \varepsilon/2$.

- (3) $\tilde{\sigma}_i$ is also a profitable deviation, differing from σ_i in finitely many histories.

One-Shot Deviation Principle

Proof

No profitable one-shot deviations $\implies$ SPNE. Immediate.

No profitable one-shot deviations $\longleftarrow$ SPNE:

(4) $\tilde{\sigma}_i$ is also a profitable deviation, differing from σ_i in finitely many histories.

Find a one-shot profitable deviation:

(5) For $\ell = 0, 1, \dots, T - t$:

If, for every $h \in H^{T-\ell}$, $u_{i|h}(\tilde{\sigma}_{i|h}, \sigma_{-i|h}) \leq u_{i|h}(\sigma_{i|h}, \sigma_{-i|h})$, redefine $\tilde{\sigma}_i(h) = \sigma_i(h)$.

Redefined $\tilde{\sigma}_i$ is still a profitable deviation.

Else, pick some $h^* \in H^{T-\ell}$ and redefine

$\tilde{\sigma}_i(h') = \sigma_i(h')$ for $h' \in H \setminus \{h^*\}$ and $\tilde{\sigma}_i(h^*) = \hat{\sigma}_i(h^*)$.

(6) h^* is our one-shot profitable deviation.

(one such h^* has to exist as $\tilde{\sigma}_i$ is a profitable deviation)

One-Shot Deviation Principle

Proposition (One-Shot Deviation Principle)

σ is SPNE if and only if there are no profitable one-shot deviations.

Proof extends immediately to perfect information games.

Version one-shot deviation principle works for extensive-form games in general:

There we define one-shot deviations to be a player deviating at any information set that belongs to a subgame but not to any of its proper subgames.

Implications of the One-Shot Deviation Principle in Repeated Games

Definition

σ is history-independent iff for any non-terminal histories h, h' $\sigma(h) = \sigma(h') = \alpha$.

Proposition

A history-independent σ is SPNE if and only if $\forall t, \sigma(h^t) = \alpha$ is a Nash equilibrium of the stage game.

Proof

If $\sigma(h)$ is NE of stage game h , then it is an SPNE.

If σ is history-independent, then deviation at h does not change continuation play.

Hence, $\exists$ profitable one-shot deviation at any given history h

if and only if $\exists i$ s.t. $\alpha_i = \sigma_i(h)$ not BR to $\alpha_{-i} = \sigma_{-i}(h)$

if and only if $\alpha = \sigma(h)$ not NE of the stage game.

Implications of the One-Shot Deviation Principle in Repeated Games

Definition

σ is history-independent iff for any non-terminal histories h, h' $\sigma(h) = \sigma(h') = \alpha$.

Proposition

A history-independent σ is SPNE if and only if $\forall t, \sigma(h^t) = \alpha$ is a Nash equilibrium of the stage game.

Playing same stage-game NE every period is SPNE

Playing NE (not necessarily the same) every period is SPNE (insofar as it doesn't depend on what was played before).

Holds for extensive-move stage games as well, but considering stage-game SPNE.

But are there other SPNE?

Finitely Repeated Games

Proposition

In any finitely repeated game s.t. stage game has unique Nash equilibrium α , the unique SPNE of the repeated game is the history-independent strategy profile $\sigma : \sigma(h^t) = \alpha \forall h^t \in H^t, \forall t = 1, \dots, T$.

Proof

Follows immediately by backward induction.

- (1) In period T , must be playing NE of the stage game. (Example)
- (2) When there are multiple stage-game NE, there may be SPNE in which there is non-stage-game-NE play at periods $t = 1, \dots, T - 1$. (Example)

SPNE in Infinitely Repeated Games

SPNE: need to check deviations following every possible history.

(on equilibrium path and off equilibrium path)

Punishments need to be credible: need to be playing a NE at all subgames.

Minmaxing is not going to be a credible punishment in general:

Punishers need to be playing a Nash equilibrium of the subgame!

Idea: punish using “bad” Nash equilibria.

Player i 's **Nash-threat payoff** is given by

$$\underline{v}_i := \inf \{ v_i = \pi_i(\alpha) \mid \alpha \text{ is a (possibly mixed) stage-game Nash equilibrium} \}.$$

Nash-reversion (trigger) strategy:

play a_i ; if no player deviates from prescribed path of play, keep playing a_i ;

if player j deviates, revert to playing Nash equilibrium that yields $\underline{v}_j$ (“bad” NE for j).

Define $V^T := \{v \in \mathbb{R}^I \mid v_i > \underline{v}_i \forall i\}$

Example

Firm pricing

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > \frac{4}{7}p_H > 0$.

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

One-shot interaction: unique NE (L, L) .

but... firms face repeated pricing problem.

Can firms collude in pricing higher?

Example

Firm pricing

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > 4/7 p_H > 0$.

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

$$\pi_i(L, L) = p_L/2 = \underline{v}_i.$$

Sustain high prices with Nash reversion:

- Play H ; if $h^t = (H, H) =_{\ell=1}^{t-1}$, play H at t ; otherwise, player L .
- Relevant histories $h^t = (H, H) =_{\ell=1}^{t-1}$, and all other.
- Payoff from sticking to strategy: $p_H/2$.
- Payoff from deviating: if $h^t = \emptyset$ or $(H, H) \neq_{\ell=1}^{t-1}$, OSD yields $p_L 7/8 > p_H/2$ today and $p_L/2$ forever after.

Avg discounted payoff: $(1 - \delta)p_L 7/8 + \delta p_L/2$.

- Payoff from deviating: if any other h^t , not profitable to deviate (playing stage NE).
- Need $(1 - \delta)p_L 7/8 + \delta p_L/2 \leq p_H/2 \iff \delta \geq \frac{p_L - 4/7 p_H}{3/7 p_L} \in (0, 1)$

Nash-threats Folk Theorem for SPNE in Infinitely Repeated Games

Theorem (Friedman 1971)

Let $v \in V^F \cap V^T$. $\exists \delta^* \in (0,1)$ s.t. $\forall \delta > \delta^*$, Γ_δ^∞ has a SPNE σ that yields average discounted equilibrium payoffs of v .

Intuition: coordinate on prob. p that generates v ; if player j deviates, revert to j -worst stage-game NE forever.

If there are multiple players deviating, then choose one to be punished.

OSDP: only need to check for one-shot profitable deviations.

Nash-threats Folk Theorem for SPNE in Infinitely Repeated Games

Theorem (Friedman, 1971 RES)

Let $v \in V^F \cap V^T$. $\exists \delta^* \in (0, 1)$ s.t. $\forall \delta > \delta^*$, Γ_δ^∞ has a SPNE σ that yields average discounted equilibrium payoffs of v .

Proof idea

Intuition: coordinate on prob. p that generates v ; if player j deviates first, revert to j -worst stage-game NE α^j , forever (no matter what happens).

If there are multiple players deviating first, then choose one to be punished.

OSDP: Only need to check one-shot deviations.

At h s.t. someone already deviated: playing stage NE $\implies$ SPNE forever after.

At h s.t. no one deviated:

Payoff from adhering to strategy: v_j .

Payoff from deviating at h^t : $(1 - \delta) \max_{a_i} \pi_i(a_i, p_{-i}) + \delta \underline{v}_j$.

Better not to deviate if $\delta \geq \frac{\max_{a_i} \pi_i(a_i, p_{-i}) - v_j}{\max_{a_i} \pi_i(a_i, p_{-i}) - \underline{v}_j}$.

□

More SPNE in Repeated Games

Can we do even more?

Yes!

Idea: punish using minmaxing, but... it's a bit more complicated because, again, punishments need to be credible, and minmaxing forever is not in general credible.

Folk Theorem for SPNE in Infinitely Repeated Games

Theorem (Fudenberg & Maskin, 1986 Ecta)

Let (i) $v \in V^F \cap V^{SIR}$, and (ii) $\text{int}(V^F \cap V^{SIR}) \neq \emptyset$.

$\exists \delta^* \in (0, 1) : \delta > \delta^*, \Gamma_\delta^\infty$ has a SPNE σ that yields average discounted equilibrium payoffs of v .

Proof intuition

“Cooperative phase”: Play p to get v , as long as no player deviates.

“Punishment phase for player j ”: If player j deviates, minmax j for T^j periods.

“Reward phase, after punishing j ”: After minmaxing of j is over, other players reward themselves forever by choosing outcome that is relatively good for them.

(And could yield more than v_{-j} or not, but yields less than v_j for player j .)

If a player i deviates during any phase, then punishment for player i starts.

Reward ensures that after j 's deviation, it's SPNE of subgame to punish & reward.

Players don't need to revert to the static stage-game NE as punishment: they minmax the deviating player.

Folk Theorem for SPNE in Infinitely Repeated Games

Theorem (Fudenberg & Maskin, 1986 Ecta)

Let (i) $v \in V^F \cap V^{SIR}$, and (ii) $\text{int}(V^F \cap V^{SIR}) \neq \emptyset$.

$\exists \delta^* \in (0, 1) : \delta > \delta^*, \Gamma_\delta^\infty$ has a SPNE σ that yields average discounted equilibrium payoffs of v .

Remarks

Can do away with (ii) for 2-player normal-form stage games.

Can have the stage game being an extensive-form game.

But beware: need some version of (ii) if game is extensive-form, even with 2 players (Wen, 2002 RES; Sorin 1995 GEB; and Rubinstein & Wolinsky 1995, GEB).

Folk Theorem for SPNE in Infinitely Repeated Games

Theorem (Fudenberg & Maskin, 1986 Ecta)

Let (i) $v \in V^F \cap V^{SIR}$, and (ii) $\text{int}(V^F \cap V^{SIR}) \neq \emptyset$.

$\exists \delta^* \in (0, 1) : \delta > \delta^*, \Gamma_\delta^\infty$ has a SPNE σ that yields average discounted equilibrium payoffs of v .

Remarks

Note discontinuity at $T = \infty$:

for $T < \infty$, Γ with unique NE $\implies$! SPNE of Γ_δ^T ;

for $T = \infty$, can get any payoff in $V^F \cap V^{SIR}$.

Folk Theorem for SPNE in Finitely Repeated Games

What about SPNE for finitely repeated games?

Theorem (Benoît & Krishna, 1985 Ecta)

Let $T < \infty$. Assume $\forall i, \exists$ stage-game NE $\alpha^i : \pi_i(\alpha^i) > \underline{v}_i$ and that $\text{int}(V^F \cap V^{SIR}) \neq \emptyset$.
 $\forall \epsilon > 0$ and all $v \in V^F \cap V^{SIR}$, $\exists T^*$ s.t. $\forall T > T^*, \exists$ SPNE of Γ_1^T such that the average payoff is within ϵ of v .

Note: we need there to be, for each player, two Nash equilibria with different payoffs.

Folk Theorem for SPNE in Infinitely Repeated Games

Extending the Folk Theorem:

Can do away with (ii) for 2-player normal-form stage games.

Can do without public randomisation (Fudenberg & Maskin, 1991 JET).

Can have strategies with limited memory/bounded recall (Barlo, Carmona, & Sabourian, 2016 JET).

Can have anonymous random matching: private monitoring, no public randomisation (Deb, Sugaya, & Wolitzky 2020 Ecta).

Can have the stage game being an extensive-form game.

But beware: need some version of (ii) if game is extensive-form, even with 2 players (Wen, 2002 RES; Sorin 1995 GEB; and Rubinstein & Wolinsky 1995, GEB).

Many, many generalisations (Mailath & Samuelson, 2006 Book)...

Characterising SPNE Payoffs: Abreu, Pearce, & Stachetti (1990 Ecta)

APS characterises set of SPNE payoffs for fixed δ (used in repeated games and applications, e.g., IO).

Also provide characterisation of supporting equilibrium strategies.

(Also: APS for stochastic games, Abreu, Brooks, Sannikov (2020 Ecta))

Anything goes?

Understanding of what kinds of strategy we need to support particular payoffs.
Restrictions on strategy space (e.g., Markov-perfect).

Overview

1. Repeated Interaction
2. Repeated Games
3. Characterising Nash Equilibria
4. Characterising Subgame-Perfect Nash Equilibria
5. Applications
 - Efficiency Wages
6. More

Stylised model of employment

- Firm makes wage offer $w \geq 0$ to worker.
- Worker chooses whether to accept and, if so, whether to “work hard” or “shirk”;
 $a \in \{W, S, NA\}$.
 - Working hard has cost $c > 0$.
 - Outside option $\bar{u} > 0$ if worker doesn't accept the offer. $u(w, NA) = \bar{u}, \quad u(w, W) = w - c, \quad u(w, S) = w.$
- Firm gets nothing if the worker turns down the offer;
 - $v - w$ if the worker accepts and works hard;
 - $-w$ if the worker accepts and shirks.Assume $v > \bar{u} + c$.

Efficiency Wages

One-Shot Interaction (SPNE)

- Worker rejects offers $w < \bar{u}$; accepts $w \geq \bar{u}$ and shirks.
- Firm gets $-w$ if worker accepts and shirks, and 0 if worker rejects offer.
- At any SPNE, the worker isn't hired.
- It would be efficient for firm to hire the worker if the worker will actually work; but, once hired, worker would like to slack off.

Efficiency Wages

Infinitely Repeated Interaction Γ_δ^∞ (SPNE).

Claim: If $v \geq \bar{u} + c/\delta$, $\exists$ SPNE in which the firm offers $w \in [\bar{u} + c/\delta, v]$ and the worker works hard in every period.

- At $t = 1$ or after any history where there has been no deviation: firm offers $w \in [\bar{u} + \frac{c}{\delta}, v]$ and worker works hard.
- After any deviation by the worker, the firm offers $w = 0$ and the worker does rejects.
- After any deviation by the firm offering some other w , the worker rejects if $w < \bar{u}$, and accepts and shirks if $w \geq \bar{u}$.
- Following a deviation, playing a stage-game SPNE; need to check if there is incentive to deviate when no deviation has yet occurred.

Efficiency Wages

Infinitely Repeated Interaction Γ_δ^∞ (SPNE).

Claim: If $v \geq \bar{u} + c/\delta$, $\exists$ SPNE in which the firm offers $w \in [\bar{u} + c/\delta, v]$ and the worker works hard in every period.

- At $t = 1$ or after any history where there has been no deviation:
firm offers $w \in [\bar{u} + \frac{c}{\delta}, v]$ and worker works hard.
- After any deviation by the worker, the firm offers $w = 0$ and the worker does rejects.
- After any deviation by the firm offering some other w , the worker rejects if $w < \bar{u}$,
and accepts and shirks if $w \geq \bar{u}$.
- Firm's best deviation payoff is zero: if it offers any other wage, worker either
rejects, or accepts and shirks.
Firm has no incentive to deviate since $v - w \geq 0$. (i.e. $v \geq w$)
- Worker has no incentive to reject as long as $w \geq \bar{u} + c$;
has an incentive to work rather than shirk if

$$\underbrace{w - c}_{\text{payoff today if W}} \geq \underbrace{(1 - \delta)w}_{\text{payoff today if S}} + \underbrace{\delta \bar{u}}_{\text{future payoff if fired}}$$
$$\Rightarrow w \geq \bar{u} + \frac{c}{\delta}$$

Overview

1. Repeated Interaction
2. Repeated Games
3. Characterising Nash Equilibria
4. Characterising Subgame-Perfect Nash Equilibria
5. Applications
6. More

Folk Theorems: Mailath & Samuelson (2006 Book, ch. 3), Abreu, Brooks & Sannikov (2020 Ecta), Abreu, Pearce & Stacchetti (1990 Ecta)

Imperfect Monitoring and Private Monitoring: Mailath & Samuelson (2006 Book, parts 2-3)

Price Wars: Mailath & Samuelson (2006 Book, ch. 6.1)

Markov-Perfect Equilibrium: Fudenberg & Tirole (1991 Book, ch. 13), Maskin & Tirole (2001 JET), Horst (2005 GEB)

Axiomatizing Play in Repeated Games: Mathevet (2018 GEB)

Limited Foresight and Complexity: Jehiel (2001 REStud), Eliaz (2003 GEB)

Experiments: Dal Bó (2005 AER), Dal Bó & Fréchette (2011 AER), Aoyagi, Bhaskar, Fréchette (2019 AEJMicro), Agranov & Elliott (2021 JEEA).

Appendix.

Example

Firm pricing

(Back)

Two firms, A, B ; can price p_H or p_L . $p_H > p_L > 4/7 p_H > 0$.

		Firm B	
		H	L
Firm A	H	$p_H/2, p_H/2$	$p_H/8, p_L 7/8$
	L	$p_L 7/8, p_H/8$	$p_L/2, p_L/2$

One-shot interaction: unique NE (L, L) .

but... firms face repeated pricing problem.

Can firms collude in pricing higher?

Avg. disc. payoffs $(p_H/2, p_H/2)$ approx. attainable as NE of repeated game if players sufficiently patient and game is repeated long enough.

Avg. disc. payoffs $(p_L/2, p_L/2)$ unique SPNE payoff insofar as game is finitely repeated (i.e., $T < \infty$).

Example

(Back)

		Col		
		A	B	C
Row	A	4,4	0,0	0,5
	B	0,0	1,1	0,0
	C	5,0	0,0	3,3

3 stage NE: (B,B), (C,C), and $(\frac{3}{4}B + \frac{1}{4}C, \frac{3}{4}B + \frac{1}{4}C)$.

Game repeated twice, Γ_1^2 .

Playing stage-game NE in each period is SPNE.

Can we have players not playing stage-game NE?

Idea: condition which stage NE is played at $t = 2$ depending on what is played at $t = 1$.

Example

(Back)

		Col		
		A	B	C
Row	A	4,4	0,0	0,5
	B	0,0	1,1	0,0
	C	5,0	0,0	3,3

Strategy:

Play (A,A) in 1st period.

If first period play is (A,A), play (C,C) in 2nd period.

Otherwise, play (B,B) in 2nd period.

Use OSDP; possible histories

partition histories: $\{\emptyset\}$, $\{(A,A)\}$, and $\{(A,B),(B,A),(B,B),(A,C),(C,A),(B,C),(C,B),(C,C)\}$.

$h = \emptyset$: deviating to C: $5+1$ vs $4+3$.

Other histories: $t = T = 2$ playing stage NE, no profitable deviation.