Dynamic Games of Complete Information

Repeated Games

Infinitely Repeated Games

Infinitely Repeated Games

Let $G = (\{1,2\}, (S_1, S_2), (u_1, u_2))$ be a strategic-form game of two players.

An *infinitely repeated game G_{irep}* based on G proceeds in *stages* so that in each stage, say t, players choose a strategy profile $s^t = (s_1^t, s_2^t)$.

Recall that a *history of length $t \geq 0$* is a sequence $h = s^1 \cdots s^t \in S^t$ of t strategy profiles. Denote by $H(t)$ the set of all histories of length t.

A *pure strategy* for player i in the infinitely repeated game G_{irep} is a function

$$\tau_i : \bigcup_{t=0}^{\infty} H(t) \to S_i$$

which for every possible history chooses a next step for player i.

Every pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} induces a sequence of pure strategy profiles $w_\tau = s^1 s^2 \cdots$ in G so that $s_i^t = \tau_i(s^1 \cdots s^{t-1})$.
(Here for $t = 0$ we have that $s^1 \cdots s^{t-1} = \epsilon$.)

Infinitely Repeated Games & Discounted Payoff

Let $\tau = (\tau_1, \tau_2)$ be a pure strategy profile in G_{irep} such that $w_\tau = s^1 s^2 \cdots$

Given $0 < \delta < 1$, we define a *δ-discounted payoff* by

$$u_i^\delta(\tau) = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \cdot u_i(s^{t+1})$$

Given a strategic-form game G and $0 < \delta < 1$, we denote by G_{irep}^δ the infinitely repeated game based on G together with the δ-discounted payoffs.

Infinitely Repeated Games & Discounted Payoff

Definition 78
A strategy profile $\tau = (\tau_1, \tau_2)$ is a Nash equilibrium in G^{δ}_{irep} if for both $i \in \{1, 2\}$ and for every τ'_i we have that

$$u^{\delta}_i(\tau_i, \tau_{-i}) \quad \geq \quad u^{\delta}_i(\tau'_i, \tau_{-i})$$

Given a history $h = s^1 \cdots s^t$ and a strategy τ_i of player i, we define a strategy τ^h_i in the infinitely repeated game G_{irep} by

$$\tau^h_i(\bar{s}^1 \cdots \bar{s}^{\bar{t}}) = \tau_i(s^1 \cdots s^t \bar{s}^1 \cdots \bar{s}^{\bar{t}}) \quad \text{for every sequence } \bar{s}^1 \cdots \bar{s}^{\bar{t}}$$

(i.e. τ^h_i behaves as τ_i after h)

Now $\tau = (\tau_1, \tau_2)$ is a SPE in G^{δ}_{irep} if for every history h we have that (τ^h_1, τ^h_2) is a Nash equilibrium.
Note that (τ^h_1, τ^h_2) must be a NE also for all histories h that are *not* visited when the profile (τ_1, τ_2) is used.

Example

Consider the infinitely repeated game G_{irep} based on Prisoner's dilemma:

	C	S
C	$-5,-5$	$0,-20$
S	$-20,0$	$-1,-1$

What are the Nash equilibria and SPE in G_{irep}^{δ} for a given δ ?

Consider a pure strategy profile (τ_1, τ_2) where $\tau_i(s^1 \cdots s^T) = C$ for all $T \geq 1$ and $i \in \{1, 2\}$. Is it a NE? A SPE?

Consider a "grim trigger" profile (τ_1, τ_2) where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

Is it a NE? Is it a SPE?

One-Shot Deviation Principle

A pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} satisfies one-shot deviation property in G^{δ}_{irep} if for every $i \in \{1,2\}$ and every $\bar{\tau}_i$, *differing from* τ_i *just on a single history* h, we have $u_i^{\delta}(\bar{\tau}_1^h, \tau_2^h) \leq u_i^{\delta}(\tau_1^h, \tau_2^h)$.

Theorem 79

Let $G = (\{1,2\}, (S_1, S_2), (u_1, u_2))$ be a two-player strategic-form game such that both u_1 and u_2 are bounded on $S = S_1 \times S_2$. Let $0 < \delta < 1$. A pure strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} is a SPE in G^{δ}_{irep} **iff** *it satisfies the one-shot deviation property in G^{δ}_{irep}.*

Before proving Theorem 79, let us note the following:

- The one shot deviation property is concerned with all strategies $\bar{\tau}_i$ that differ from τ_i on a single history. This means that we have to consider *all* histories h, even those that *can not* be visited using τ_i with any opponent.
- The one-shot deviation property immediately implies the following: If $\bar{\tau}_i$ *does not* differ from τ_i on any history of the form $h' = hh''$ where $h'' \neq \varepsilon$ (i.e., on any history obtained by prolonging h), then $u_i^{\delta}(\bar{\tau}_1^h, \tau_2^h) \leq u_i^{\delta}(\tau_1^h, \tau_2^h)$.

  Indeed, note that τ_i^h differs from $\bar{\tau}_i^h$ only on h.

One-Shot Deviation Principle

Proof. $\Rightarrow$: Trivial.
$\Leftarrow$: Assume that τ satisfies the one-shot deviation property but is not a SPE. That is, a deviation may increase payoff of one of the players in a subgame. Assume, w.l.o.g., that player 1 gains by deviation to a strategy $\bar{\tau}_1$ in a subgame starting with a h, i.e.,

$$u_1^{\delta}(\bar{\tau}_1^h, \tau_2^h) > u_1^{\delta}(\tau_1^h, \tau_2^h) \tag{29}$$

Since $\delta < 1$ and u_i are bounded on S, we may safely choose $\bar{\tau}_1$ so that $\bar{\tau}_1(h') = \tau_1(h')$ for all sufficiently long histories h'.

Indeed, since u_i is bounded on pure strategies of G, the sum $\sum_{t=\ell}^{\infty} \delta^t \cdot u_i(s^{t+1})$ goes to 0 as ℓ goes to ∞; hence the strict inequality (29) remains valid even if $\bar{\tau}_1$ is arbitrarily modified in a very distant future.

One-Shot Deviation Principle

Let h' be a history of *maximum length* such that h is a prefix of h' and $\bar{\tau}_1(h') \neq \tau_1(h')$. (Note that then $\bar{\tau}_1(h'h'') = \tau_1(h'h'')$ for all $h'' \neq \varepsilon$.)

Let $\bar{\tau}_{11}$ be a strategy of player 1 obtained from $\bar{\tau}_1$ by changing $\bar{\tau}_1(h')$ to $\tau_1(h')$. Now note that the one-shot deviation property implies, that

$$u_1^\delta(\bar{\tau}_{11}^{h'}, \tau_2^{h'}) = u_1^\delta(\tau_1^{h'}, \tau_2^{h'}) \geq u_1^\delta(\bar{\tau}_1^{h'}, \tau_2^{h'})$$

and thus $u_1^\delta(\bar{\tau}_{11}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_1^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$. Note that $\bar{\tau}_{11}^h$ has a strictly smaller number of deviations from τ_1^h than $\bar{\tau}_1^h$.

Repeating the same argument with $\bar{\tau}_{11}$ in place of $\bar{\tau}_1$ we obtain $\bar{\tau}_{12}$ such that $u_1^\delta(\bar{\tau}_{12}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_{11}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$. Here $\bar{\tau}_{12}^h$ has even less deviations from τ_1^h than $\bar{\tau}_{11}^h$.

Then repeating with $\bar{\tau}_{12}$ in place of $\bar{\tau}_1$ we obtain $\bar{\tau}_{13}$ such that $u_1^\delta(\bar{\tau}_{13}^h, \tau_2^h) \geq u_1^\delta(\bar{\tau}_{12}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$, etc., still decreasing the number of deviations from τ_1^h.

Eventually, as $\bar{\tau}_1^h$ has only finitely many deviations from τ_1^h, we get $\bar{\tau}_{1k}^h = \tau_1^h$ for some k and thus $u_1^\delta(\tau_1^h, \tau_2^h) = u_1^\delta(\bar{\tau}_{1k}^h, \tau_2^h) > u_1^\delta(\tau_1^h, \tau_2^h)$, a contradiction. $\square$

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

The grim trigger profile (τ_1, τ_2) where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

is a SPE.

A Simple Version of Folk Theorem

Let $G = (\{1,2\}, (S_1, S_2), (u_1, u_2))$ be a two-player strategic-form game where u_1, u_2 are bounded on $S = S_1 \times S_2$ (but S may be infinite) and let s^* be a Nash equilibrium in G.

Let s be a strategy profile in G satisfying $u_i(s) > u_i(s^*)$ for all $i \in N$.

Consider the following *grim trigger for s using s** strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} s_i & T = 0 \\ s_i & s^\ell = s \text{ for all } 1 \leq \ell \leq T \\ s_i^* & \text{otherwise} \end{cases}$$

Then for

$$\delta \quad \geq \quad \max_{i \in \{1,2\}} \frac{\max_{s_i' \in S_i} u_i(s_i', s_{-i}) - u_i(s)}{\max_{s_i' \in S_i} u_i(s_i', s_{-i}) - u_i(s^*)}$$

we have that (τ_1, τ_2) is a SPE in G_{irep}^δ and $u_i^\delta(\tau) = u_i(s)$.

Proof: Consider a possible one-shot deviation $\bar{\tau}_1$ of player 1, i.e., there is exactly one h such that $\bar{\tau}_1(h) \neq \tau_1(h)$. We distinguish two cases depending on h.

Proof of Simple Folk Theorem (Cont.)

Case 1: $h \neq s \cdots s$. Then there is a deviation from s in h and thus according to (τ_1^h, τ_2^h) both players play s^* forever :

$$u_1^\delta(\tau_1^h, \tau_2^h) = (1-\delta)\sum_{k=0}^{\infty}\delta^k u_1(s^*) = u_1(s^*)(1-\delta)\sum_{k=0}^{\infty}\delta^k = u_1(s^*)$$

Now $(\bar{\tau}_1^h, \tau_2^h)$ gives a sequence $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s_1', s_2^*)s^*s^* \cdots$ where s_1' is a strategy of player 1 to which he deviates after h.

Here player 2 plays s_2^* all the time after h because one of the players has already deviated in h.

We obtain

$$\begin{aligned} u_1(\bar{\tau}_1^h, \tau_2^h) &= (1-\delta)\left(u_1(s_1', s_2^*) + \sum_{k=1}^{\infty}\delta^k u_1(s^*)\right) \\ &\leq (1-\delta)\left(u_1(s_1^*, s_2^*) + \sum_{k=1}^{\infty}\delta^k u_1(s^*)\right) \\ &= u_1(s^*) \end{aligned}$$

So this deviation cannot be beneficial no matter what δ is.

Proof of Simple Folk Theorem (Cont.)

Case 2: $h = s \cdots s$. Clearly, $u_1(\tau_1^h, \tau_2^h) = u_1(s)$.

Now $(\bar{\tau}_1^h, \tau_2^h)$ gives a sequence $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s_1', s_2)s^* s^* \cdots$ where s_1' is a strategy of player 1 to which he deviates after h.
As opposed to the previous case, here player 2 first plays s_2 (since the deviation of player 1 to s_1' is the first deviation in the history) and then both players react by playing s^* forever.

If $u_1(s_1', s_2) < u_1(s)$, then

$$
\begin{aligned}
u_1^\delta(\bar{\tau}_1^h, \tau_2^h) &= (1-\delta)\left(u_1(s_1', s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*)\right) \\
&< (1-\delta)\left(u_1(s_1, s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*)\right) \\
&< (1-\delta)\left(u_1(s) + \sum_{k=1}^{\infty} \delta^k u_1(s)\right) = u_1(s) = u_1^\delta(\tau_1^h, \tau_2^h)
\end{aligned}
$$

and thus this deviation is also not beneficial no matter what δ is.

Proof of Simple Folk Theorem (Cont.)

Finally, if $u_1(s_1', s_2) \geq u_1(s)$, then

$$u_1^\delta(\bar{\tau}_1^h, \tau_2^h) = (1-\delta)\left(u_1(s_1', s_2) + \sum_{k=1}^{\infty} \delta^k u_1(s^*)\right)$$

$$= (1-\delta)u_1(s_1', s_2) + (1-\delta)u_1(s^*) \cdot \delta \sum_{k=0}^{\infty} \delta^k$$

$$= u_1(s_1', s_2) - \delta \cdot u_1(s_1', s_2) + \delta \cdot u_1(s^*)$$

Thus

$u_1^\delta(\bar{\tau}_1^h, \tau_2^h) \leq u_1^\delta(\tau_1^h, \tau_2^h) = u_1(s)$ **iff**

$u_1(s_1', s_2) - \delta \cdot u_1(s_1', s_2) + \delta \cdot u_1(s^*) \leq u_1(s)$ **iff**

$u_1(s_1', s_2) - u_1(s) \leq \delta \cdot (u_1(s_1', s_2) - u_1(s^*))$ **iff**

$$\delta \geq \frac{u_1(s_1', s_2) - u_1(s)}{u_1(s_1', s_2) - u_1(s^*)}$$

Proof of Simple Folk Theorem (Cont.)

Thus (τ_1, τ_2) satisfies the one-shot deviation property in G^{δ}_{irep} w.r.t. player 1 if

$$\delta \geq \frac{u_1(s'_1, s_2) - u_1(s)}{u_1(s'_1, s_2) - u_1(s^*)} \text{ for all } s'_1 \in S_1 \text{ satisfying } u_1(s'_1, s_2) \geq u_1(s)$$

Note that the right-hand-side expression is maximized when $u_1(s'_1, s_2)$ is maximized and thus we get

$$\delta \geq \frac{\max_{s'_1 \in S_1} u_1(s'_1, s_2) - u_1(s)}{\max_{s'_1 \in S_1} u_1(s'_1, s_2) - u_1(s^*)}$$

Proving the same for player 2 and putting the results together, we obtain that (τ_1, τ_2) satisfies the one-shot deviation property in G^{δ}_{irep} if

$$\delta \geq \max_{i \in \{1,2\}} \frac{\max_{s'_i \in S_i} u_i(s'_i, s_{-i}) - u_i(s)}{\max_{s'_i \in S_i} u_i(s'_i, s_{-i}) - u_i(s^*)} \tag{30}$$

Thus by Theorem 79, (τ_1, τ_2) is a SPE in G^{δ}_{irep} if δ satisfies ineq. (30).

Simple Folk Theorem – Example

Consider the infinitely repeated game G_{irep} based on the following game G:

	m	f	r
M	$4,4$	$-1,5$	$3,0$
F	$5,-1$	$1,1$	$0,0$
R	$0,3$	$0,0$	$2,2$

NE in G : (F, f)

Consider the grim trigger for (M, m) using (F, f), i.e., the profile (τ_1, τ_2) in G_{irep} where

- τ_1 : Plays M in a given stage if (M, m) was played in all previous stages, and plays F otherwise.
- τ_2 : Plays m in a given stage if (M, m) was played in all previous stages, and plays f otherwise.

This is a SPE in G^{δ}_{irep} for all $\delta \geq \frac{1}{4}$. Also, $u_i(\tau_1, \tau_2) = 4$ for $i \in \{1, 2\}$.

Are there other SPE? Yes, a grim trigger for (R, r) using (F, f). This is a SPE in G^{δ}_{irep} for $\delta \geq \frac{1}{2}$.

Tacit Collusion

Consider the Cournot duopoly game model $G = (N, (S_i)_{i\in N}, (u_i)_{i\in N})$

- $N = \{1, 2\}$
- $S_i = [0, \kappa]$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

If the firms sign *a binding contract* to produce only $\theta/4$, their profit would be $\theta^2/8$ which is higher than the profit $\theta^2/9$ for playing the NE $(\theta/3, \theta/3)$.

However, such contracts are forbidden in many countries (including US).

Is it still possible that the firms will behave selfishly (i.e. only maximizing their profits) and still obtain such payoffs?

In other words, is there a SPE in the infinitely repeated game based on G (with a discount factor δ) which gives the payoffs $\theta^2/8$?

Tacit Collusion

Consider the Cournot duopoly game model $G = (N, (S_i)_{i \in N}, (u_i)_{i \in N})$

- $N = \{1, 2\}$
- $S_i = [0, \infty)$
- $u_1(q_1, q_2) = q_1(\kappa - q_1 - q_2) - q_1 c_1 = (\kappa - c_1)q_1 - q_1^2 - q_1 q_2$
  $u_2(q_1, q_2) = q_2(\kappa - q_2 - q_1) - q_2 c_2 = (\kappa - c_2)q_2 - q_2^2 - q_2 q_1$

Assume for simplicity that $c_1 = c_2 = c$ and denote $\theta = \kappa - c$.

Consider the grim trigger profile for $(\theta/4, \theta/4)$ using $(\theta/3, \theta/3)$:
Player i will

- produce $q_i = \theta/4$ whenever all profiles in the history are $(\theta/4, \theta/4)$,
- whenever one of the players deviates, produce $\theta/3$ from that moment on.

Assuming that $\kappa = 100$ and $c = 10$ (which gives $\theta = 90$), this is a SPE G_{irep}^{δ} for $\delta \geq 0.5294 \cdots$. It results in $(\theta/4, \theta/4)(\theta/4, \theta/4) \cdots$ with the discounted payoffs $\theta^2/8$.

Dynamic Games of Complete Information

Repeated Games

Infinitely Repeated Games

Long-Run Average Payoff and Folk Theorems

Infinitely Repeated Games & Average Payoff

In what follows we assume that all payoffs in the game G are positive and that S is finite!

Let $\tau = (\tau_1, \tau_2)$ be a strategy profile in the infinitely repeated game G_{irep} such that $w_\tau = s^1 s^2 \cdots$.

Definition 80

We define a *long-run average payoff* for player i by

$$u_i^{avg}(\tau) = \limsup_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} u_i(s^t)$$

(Here $\limsup$ is necessary because τ_i may cause non-existence of the limit.)

The lon-run average payoff $u_i^{avg}(\tau)$ is *well-defined* if the limit $u_i^{avg}(\tau) = \lim_{T\to\infty} \frac{1}{T} \sum_{t=1}^{T} u_i(s^t)$ exists.

Given a strategic-form game G, we denote by G_{irep}^{avg} the infinitely repeated game based on G together with the long-run average payoff.

Definition 81

A strategy profile τ is a Nash equilibrium if $u_i^{avg}(\tau)$ is well-defined for all $i \in N$, and for every i and every τ'_i we have that

$$u_i^{avg}(\tau_i, \tau_{-i}) \quad \geq \quad u_i^{avg}(\tau'_i, \tau_{-i})$$

(Note that we demand existence of the defining limit of $u_i^{avg}(\tau_i, \tau_{-i})$ but the limit does not have to exist for $u_i^{avg}(\tau'_i, \tau_{-i})$.)

Moreover, $\tau = (\tau_1, \tau_2)$ is a SPE in G_{irep}^{avg} if for every history h we have that (τ_1^h, τ_2^h) is a Nash equilibrium.

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

The grim trigger profile (τ_1, τ_2) where

$$\tau_i(s^1 \cdots s^T) = \begin{cases} S & T = 0 \\ S & s^\ell = (S, S) \text{ for all } 1 \leq \ell \leq T \\ C & \text{otherwise} \end{cases}$$

is a SPE which gives the long-run average payoff -1 to each player.

The intuition behind the grim trigger works as for the discounted payoff: Whenever a player i deviates, the player $-i$ starts playing C for which the best response of player i is also C. So we obtain $(S, S) \cdots (S, S)(X, Y)(C, C)(C, C) \cdots$ (here (X, Y) is either (C, S) or (S, C) depending on who deviates). Apparently, the long-run average payoff is -5 for both players, which is worse than -1.

Example

Consider the infinitely repeated game based on Prisoner's dilemma:

	C	S
C	$-5, -5$	$0, -20$
S	$-20, 0$	$-1, -1$

However, other payoffs can be supported by NE. Consider e.g. a strategy profile (τ_1, τ_2) such that

- Both players **cyclically** play as follows:
  - 9 times (S, S)
  - once (S, C)
- If one of the players deviates, then, from that moment on, both play (C, C) forever.

Then (τ_1, τ_2) is also SPE.

Apparently, $u_1^{avg}(\tau_1, \tau_2) = \frac{9}{10} \cdot (-1) + (-20)/10 = -29/10$ and $u_1^{avg}(\tau_1, \tau_2) = \frac{9}{10}(-1) = -9/10$.

Player 2 gets better payoff than from the Pareto optimal profile (S, S)!

Outline of the Folk Theorems

The previous examples suggest that other (possibly all?) convex combinations of payoffs may be obtained by means of Nash equilibria.

This observation forms a basis for a bunch of theorems, collectively called Folk Theorems.
No author is listed since these theorems had been known in games community long before they were formalized.

In what follows we prove several versions of Folk Theorem concerning achievable payoffs for repeated games.

Ordered by increasing technical and conceptual difficulty, we consider the following variants:

- Long-run average payoffs & SPE
- Discounted payoffs & SPE
- Long-run average payoffs & Nash equilibria

Folk Theorems – Feasible Payoffs

Definition 82
We say that a vector of payoffs $v = (v_1, v_2) \in \mathbb{R}^2$ is *feasible* if it is a convex combination of payoffs for pure strategy profiles in G with rational coefficients, i.e., if there are rational numbers β_s, here $s \in S$, satisfying $\beta_s \geq 0$ and $\sum_{s \in S} \beta_s = 1$ such that for both $i \in \{1, 2\}$ holds

$$v_i = \sum_{s \in S} \beta_s \cdot u_i(s)$$

We assume that there is $m \in \mathbb{N}$ such that each β_s can be written in the form $\beta_s = \gamma_s / m$.

The following theorems can be extended to a notion of feasible payoffs using *arbitrary, possibly irrational,* coefficients β_s in the convex combination. Roughly speaking, this follows from the fact that each real number can be approximated with rational numbers up to an arbitrary error. However, the proofs are technically more involved.

Folk Theorems – Long-Run Average & SPE

Theorem 83
Let s^* be a pure strategy Nash equilibrium in G and let $v = (v_1, v_2)$ be a feasible vector of payoffs satisfying $v_i \geq u_i(s^*)$ for both $i \in \{1, 2\}$. Then there is a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} such that

- *τ is a SPE in G_{irep}^{avg}*
- *$u_i^{avg}(\tau) = v_i$ for $i \in \{1, 2\}$*

Proof: Consider a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} which gives the following behavior:

1. Unless one of the players deviates, the players play **cyclically** all profiles $s \in S$ so that each s is always played for γ_s rounds.
2. Whenever one of the players deviates, then, from that moment on, each player i plays s_i^*.

It is easy to see that $u_i^{avg}(\tau) = v_i$.

We verify that τ is SPE.

Folk Theorems – Long-Run Average & SPE

Fix a history h, we show that $\tau^h = (\tau_1^h, \tau_2^h)$ is a NE in G_{irep}^{avg}.

- If h does not contain any deviation from the cyclic behavior 1., then τ^h continues according to 1., thus $u_i^{avg}(\tau^h) = v_i$.
- If h contains a deviation from 1., then

  $$w_{\tau^h} = s^* s^* \cdots$$

  and thus $u_i^{avg}(\tau^h) = u_i(s^*)$.
- Now if a player i deviates to $\bar{\tau}_i^h$ from τ_i^h in G_{irep}^{avg}, then

  $$w_{(\bar{\tau}_i^h, \tau_{-i}^h)} = (s_i^1, s'_{-i})(s_i^2, s_{-i}^*)(s_i^3, s_{-i}^*) \cdots$$

  where $s_i^1, s_i^2, \ldots$ are strategies of S_i and s'_{-i} is a strat. of S_{-i}. However, then $u_i^{avg}(\bar{\tau}_i^h, \tau_{-i}^h) \leq u_i(s^*) \leq v_i$ since s^* is a Nash equilibrium and thus $u_i(s_i^k, s_{-i}^*) \leq u_i(s^*)$ for all $k \geq 1$.

  Intuitively, player $-i$ punishes player i by playing s_{-i}^*. $\square$

Folk Theorems – Discounted Payoffs & SPE

Theorem 84

Let s^* be a pure strategy Nash equilibrium in G and let $v = (v_1, v_2)$ be a feasible payoff satisfying $v_i > u_i(s^*)$ for both $i \in \{1,2\}$. Then there is a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} and $\delta < 1$ such that

- *τ is a SPE in $G_{irep}^{\delta'}$ for every $\delta' \in [\delta, 1)$ and*
- *$\lim_{\delta' \to 1} u_i^{\delta'}(\tau) = v_i$.*

Proof: The following claim allows us to reduce the discounted payoff to the long-run-average.

Claim 5

Let τ be a well-defined strategy profile. Then

$$\lim_{\delta \to 1^-} u_i^{\delta}(\tau) = u_i^{avg}(\tau)$$

Now to prove Theorem 84, consider the strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} from the proof of Theorem 83.

We check the one-shot deviation property in G_{irep}^{δ} for δ close to 1.

Folk Theorems – Discounted Payoffs & SPE

Fix a history h and consider $\tau^h = (\tau_1^h, \tau_2^h)$.

- If h does not contain any deviation from 1., then both players follow 1., and $u_i^\delta(\tau^h)$ is close to $u_i^{avg}(\tau^h) = v_i$ for δ close to 1.
- If h contains any deviation from 1., then $w_{\tau^h} = s^* s^* \cdots$ and $u_i^\delta(\tau^h) = u_i(s^*)$.
- Now assume, w.l.o.g., that player 1 deviates exactly after h, which gives a strategy $\bar{\tau}_1^h$ differing from τ_1^h only on h. Thus $w_{(\bar{\tau}_1^h, \tau_2^h)} = (s_1', s_2') s^* s^* \cdots$ where s_1' is a strategy of S_1 and s_2' is either the next step in the cyclic behavior described by 1. (if h follows 1.), or equal to s_2^* (h does not follow 1.)

  Note that for δ close to 1, we have that $u_i^\delta(\bar{\tau}_i^h, \tau_{-i}^h)$ is close to $u_i^{avg}(\bar{\tau}_i^h, \tau_{-i}^h) = u_i(s^*)$.
  - If h follows 1., then $u_1^\delta(\tau^h)$ is close to v_1 which is greater than $u_1(s^*)$ to which $u_1^\delta(\bar{\tau}_1^h, \tau_2^h)$ is close.
  - If h does not follow 1., then $s_2' = s_2^*$ (players punish due to a deviation in h), and thus $u_1^\delta(\bar{\tau}_1^h, \tau_2^h) \le u_1(s^*) = u_1^\delta(\tau^h)$. $\square$

Folk Theorems – Individually Rational Payoffs

Definition 85

$v = (v_1, v_2) \in \mathbb{R}^2$ is *individually rational* if for both $i \in \{1,2\}$ holds

$$v_i \geq \min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$$

That is, v_i is at least as large as the value that player i may secure by playing best responses to the most hostile behavior of player $-i$.

Example:

	m	f	r
M	$4,4$	$-1,5$	$3,0$
F	$5,-1$	$1,1$	$0,0$
R	$0,3$	$0,0$	$2,2$

Here any (v_1, v_2) such that $v_1 \geq 2$ and $v_2 \geq 1$ is individually rational.

Folk Theorems – Long-Run Average & NE

Theorem 86

Let $v = (v_1, v_2)$ be a feasible and individually rational vector of payoffs. Then there is a strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} such that

- *τ is a Nash equilibrium in G_{irep}^{avg}*
- *$u_i^{avg}(\tau) = v_i$ for $i \in \{1, 2\}$*

Proof: It suffices to use a slightly modified strategy profile $\tau = (\tau_1, \tau_2)$ in G_{irep} from Theorem 83:

- Unless one of the players deviates, the players play **cyclically** all profiles $s \in S$ so that each s is always played for γ_s rounds.
- Whenever a player i deviates, the opponent $-i$ plays a strategy $s_{-i}^{\min} \in \operatorname{argmin}_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i})$.

It is easy to see that $u_i^{avg}(\tau) = v_i$.

If a player i deviates, then his long-run average payoff cannot be higher than $\min_{s_{-i} \in S_{-i}} \max_{s_i \in S_i} u_i(s_i, s_{-i}) \leq v_i$, so τ is a NE. $\square$

Folk Theorems – Long-Run Average & NE

Theorem 87

If a strategy profile $\tau = (\tau_1, \tau_2)$ is a NE in G_{irep}^{avg}, then $\left(u_1^{avg}(\tau), u_2^{avg}(\tau)\right)$ is individually rational.

Proof: Suppose that $\left(u_1^{avg}(\tau), u_2^{avg}(\tau)\right)$ is not individually rational. W.l.o.g. assume that $u_1^{avg}(\tau) < \min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2)$.

Now let us consider a new strategy $\bar{\tau}_1$ such that for an arbitrary history h the pure strategy $\bar{\tau}_1(h)$ is a best response to $\tau_2(h)$.

But then, for every history h, we have

$$u_1(\bar{\tau}_1(h), \tau_2(h)) \geq \min_{s_2 \in S_2} \max_{s_1 \in S_1} u_1(s_1, s_2) > u_1^{avg}(\tau)$$

So clearly $u_1^{avg}(\bar{\tau}_1, \tau_2) > u_1^{avg}(\tau)$ which contradicts the fact that (τ_1, τ_2) is a NE. □

Note that if irrational convex combinations are allowed in the definition of feasibility, then vectors of payoffs for Nash equilibria in G_{irep}^{avg} are exactly feasible and individually rational vectors of payoffs. Indeed, the coefficients β_s in the definition of feasibility are exactly frequencies with which the individual profiles of S are played in the NE.

Folk Theorems – Summary

- We have proved that "any reasonable" (i.e. feasible and individually rational) vector of payoffs can be justified as payoffs for a Nash equilibrium in G_{irep}^{avg} (where the future has "an infinite weight").
- Concerning SPE, we have proved that any feasible vector of payoffs dominating a Nash equilibrium in G can be justified as payoffs for SPE in G_{irep}^{avg}.

  This result can be generalized to arbitrary feasible and *strictly* individually rational payoffs by means of a more demanding construction.
- For discounted payoffs, we have proved that an arbitrary feasible vector of payoffs strictly dominating a Nash equilibrium in G can be approximated using payoffs for SPE in G_{irep}^{δ} as δ goes to 1.

  Even this result can be extended to feasible and strictly individually rational payoffs.

For a very detailed discussion of Folk Theorems see "A Course in Game Theory" by M. J. Osborne and A. Rubinstein.