

Long and short-run Players

Example (Chain store model)

Long-run player, Chain store, can either fight with or accommodate with smaller stores who decide whether to enter or not each period. In a one-shot game the long-run player can achieve ~~a~~ higher payoff by accommodating. However, in a repeated-game it may be profitable for the incumbent to fight some entrants in order to build a reputation for fighting and thus to deter entry.

The general framework (2 players, 1 long-run)

Let the one-shot game be $g: S_1 \times S_2 \rightarrow \mathbb{R}^2$ with player 1 being the long-run player and player 2 being the short-run player.

Let $V = \{(v_1, v_2) \mid v_i = g_i(s_1, s_2) \text{ for some } (s_1, s_2) \text{ and } s_2 \text{ is a BR to } s_1\}$
 $V^* = \text{convex hull of } V$

best response } different to usual definition

Note: We will consider only pure-strategy equilibria here.
 See Mailith and Samuelson for mixed-strategy extensions.

Further define:

$$v_1 = \min_{\substack{s_2 \in BR(s_1) \\ \text{for some } s_1}} \max_{s_1'} g_1(s_1', s_2)$$

Note that this is a different min max concept than before.

Consider

	L	M	R
U	0, 2	2, 3	-2, -1
D	-1, 2	1, 1	-2, -1

With 2 long-run players playing "R" minimaxes player 1. However a short-run player cannot be expected to play "R". In the case where player 2 is short run "L" is the strategy which minimaxes player 1.

Theorem (Folk-like theorem)

If $(v_1, v_2) \in V^*$, $v_1 > v_2$, then there exists δ^* s.t. for all $\delta > \delta^*$ there exists SPE of the repeated game with δ in which (v_1, v_2) are the average equilibrium payoffs.

Proof

Normal phase

choose (s_1, s_2) to obtain (v_1, v_2)

Punishment

↳
in the
previous
example they
would play $(0, L)$


if player 1 deviates minmax him large enough
to wipe out gain and 1 plays the action that
makes 2 best respond. Then return to normal phase.

↳ if no punishment
is restarted

See Mailath-Samuelson
or Fudenberg-Kreps-Maskin

Application (chain store game)

extensive form



Normal Form

	Out	In	
Fight	2, 0	-1, -1	2, 0 ∈ V
Accomodate	2, 0	1, 1	1, 1 ∈ V

By the theorem above any convex combination of
 $(2, 0)$ and $(1, 1)$ is achievable as SPE for δ big enough

INTRODUCING UNCERTAINTY TO THIS GAME

Entrant believes that with probability ϵ , the incumbent likes to fight.

	Out	In
Fight	2, 0	-1, -1
Accomodate	2, 0	1, 1

	Out	In
Fight	2, 0	2, 0
Accomodate	1, 0	1, 1

Prob. $1 - \epsilon$

prob ϵ

This uncertainty makes for sharper predictions:

Claim In any sequential equilibrium if δ is near 1, then
player 1's average payoff in equilibrium is close to 2.

Proof

N.B. This holds for pure- and mixed-strategy equilibria,
but we shall only prove it for pure-strategy equilibria.

↳ Let t be the first period in which player 1 plays "accomodate".

If $t = \infty$, player 1 gets payoff = 2.

If $t < \infty$, suppose 1 deviates and plays F and continues to
play F in all subsequent periods.

Consider seq
equilibrium
 (σ_1, σ_2)
↑
for rational
player 1

Eric Maskin

Proof (continued)

Player 1 was getting 2 before period t , gets -1 in period t and gets 2 after that. If player 1 can deviate and get a payoff close to 2, then in equilibrium (σ_1, σ_2) he must be getting at least that much. ■

Idea: What if player 1 was crazy and accommodated instead of fought?

Claim: Suppose that there are finitely many types of "crazy" player 1, but that the type that always likes to fight has positive prob. For δ near 1, in any sequential equilibrium (σ_1, σ_2) player 1's payoff is nearly 2. — again a sharp prediction

Proof Because there are only finitely many types, $\exists T^*$ s.t. if $T > T^*$, $\Pr(1's \text{ next action is F} \mid \text{in the first } T \text{ periods } 1 \text{ played } \& 1 \text{ is crazy}) = 1$

player 1 can disprove that he's some type that will at some point play "a".

(the proof continues as before)

Let $t =$ the first period where σ_1 does not play F.

If $t < \infty$, consider a deviation in which 1 plays F in period t and in all subsequent periods. By the above 1 gets an average payoff of close to 2. Thus proving the claim. ■

Evolution

Multiplicity of equilibria in repeated games is both a curse and blessing — predictions aren't sharp, but lots of behaviours are consistent with equilibrium.

Evolutionary considerations, aside from renegotiation and reputation/incomplete information, can be used to cut down the set of equilibria.

Players can be identified by strategies and we will consider equilibria which are not susceptible to invasions by mutants with non-equilibrium strategies.

→ In economics consider cultural evolution as opposed to genetic evolution in biology.

→ Firms may imitate successful mutant strategies if these prove to be successful by delivering higher payoffs.

Definition

(2 player game, repeated)
infinitely repeated game

A strategy σ is evolutionarily stable if the majority of the population (fraction $1-q$) plays σ , but there is a proportion of mutants q who play σ' , and we have that:

$$\forall \sigma', (1-q)g_i^*(\sigma, \sigma) + qg_i^*(\sigma, \sigma') > (1-q)g_i^*(\sigma', \sigma) + qg_i^*(\sigma', \sigma')$$

for all q sufficiently small.

if $q=0$ then this is the definition of Nash equilibrium

Before the more formal definitions, consider some examples.

Example (Repeated Prisoners Dilemma)

	C	D
c	2,2	4,3
d	3,1	0,0

The strategy "always D" can be invaded by the "trigger strategy".

The invaders and incumbents have the following payoffs:

	Always D	Trigger
Always D	0	$0+\epsilon$
Trigger	$0-\epsilon$	2

For any given q , T will get higher expected payoff for δ big enough.

Remark: While always D is eliminated by this equilibrium concept, some quite bad equilibrium strategies are not.

Consider strategy "Alt" Go back and forth between C&D as long as alternating pattern continues, but if pattern is broken then play D.

	^{1-q} ALT	^q σ'
ALT	1	0
σ'	0	2

for any σ'

best case scenario

still for small enough q no strategy can invade.

Similarly strategy "Alt 2" where C, D, D is played repeatedly is not ruled out by evolution.

Furthermore any strategy "Alt n" can't be ruled out by evolutionary stability. Seems like it's bad news for this concept.

However, this strategy is quite fragile. In the real world mistakes happen and the punishment for such a mistake is too drastic in this case.

Introduce a small probability p of a mistake in play. In this case "alt" is not evolutionarily stable.

Consider strategy " σ " where the player goes back and forth between C & D as long as pattern continues, but if it is broken then play C. If the other player has also played C, then play C thereafter, otherwise play D thereafter.

If there are no mistakes:

	ALT	σ'
ALT	1	1
σ'	1	1

If there is 1 mistake, continuation payoffs are:

	ALT	σ'
ALT	0	$0-\epsilon$
σ'	$0-\epsilon$	2

Since $p > 0$, for $q > 0$ σ does better than alt and thus alt is not evolutionarily stable.

Thus we would not expect alt to survive evolutionarily.

Evolution and probability of mistakes reduces the set of equilibrium payoffs for the prisoners' dilemma to payoffs close to 2.

This doesn't happen in every game. For example the battle of the sexes.

	B	S
B	0,0	2,1
S	1,2	0,0

Formal Treatment

Consider finite, 2 player symmetric game, i.e. $g_1(s_1, s_2) = g_2(s_2, s_1)$

could be easily generalised to n players

this is important to ensure only 1 population is evolving, else it would be complicated as the rate at which respective populations are evolving would be important.

Normalize minmax payoffs to zero.

Let $V^* :=$ convex hull of feasible payoff set

Definition (strong Efficiency)

$(v_1, v_2) \in V^*$ is strongly efficient if $v_1 + v_2 = \max_{(v'_1, v'_2) \in V^*} \{v'_1 + v'_2\}$

Definition (history)

$h^{t+1} := ((s_1^1, s_2^1), \dots, (s_1^t, s_2^t))$, $\pi(h^{t+1}) := ((s_2^1, s_1^1), \dots, (s_2^t, s_1^t))$
(set of actions used in past) \hookrightarrow the flip of h^{t+1} .

A strategy is $\sigma: H \rightarrow S$.
 \uparrow
all histories

Focus on finitely complex strategies.

\hookrightarrow a strategy that could be played by a finite computer with finite memory.

\hookrightarrow trigger is ok, alternating is ok, tit-for-tat ok
 \hookrightarrow rule out strategies with infinite states.

Definition (mistake)

Let p be the probability that σ makes a mistake in any period, (so that $\sigma(h) \neq s$), and make all mistakes equally likely to fix ideas.

Definition (discount rate) $\delta = \frac{1}{1+r}$, r is the discount rate, δ is discount factor.

Let $U_i^{r,p}(\sigma_i, \sigma_{-i}^p) = \frac{r}{1+r} \mathbb{E} \left[\sum_{t=1}^{\infty} \left(\frac{1}{1+r}\right)^{t-1} g_i(\sigma_i(h^t), \sigma_{-i}^p(h^t)) \mid \sigma_i, \sigma_{-i}, p, h \right]$
 \nearrow continuation payoff

Definition (evolutionary stability)

σ is evolutionarily stable wrt g, r, p if $\forall \sigma' \neq \sigma \quad \forall q' < q$

$$(1-q')U_i^{r,p}(\sigma, \sigma) + q'U_i^{r,p}(\sigma, \sigma') > (1-q')U_i^{r,p}(\sigma', \sigma) + q'U_i^{r,p}(\sigma', \sigma')$$

"No two strategies can deliver the same payoff"

\leftarrow Strict inequality since every ~~part~~ part of the game tree is reached with positive probability, not just the equilibrium path.
because of mistakes

References: Maskin-Fudenberg (1989)

"Evolution and Cooperation in Noisy repeated games"

Homework problems: Will be on Blackboard early next week
due 2 weeks after posting (send by email)

Evolution and repeated games (continued)

Recall: Definition (evolutionary stability)

σ is evolutionarily stable wrt q, r, p if $\forall \sigma' \neq \sigma, q' < q$ ^{Proportion of mutants}
 $(1-q')u_i^{r,p}(\sigma, \sigma) + q'u_i^{r,p}(\sigma, \sigma') > (1-q')u_i^{r,p}(\sigma', \sigma) + q'u_i^{r,p}(\sigma', \sigma')$

Notation:

$\underline{v} := \min\{v \geq 0 \mid \text{there exists } v' \text{ with } (v, v') \text{ strongly efficient}\}$

Theorem (Lower bound on payoffs in a repeated game).

Given $\epsilon > 0, q > 0, \exists r^*, p^*$ s.t $\forall r < r^*, p < p^*$ if σ is ES wrt q, r, p then
 $u_i^{r,p}(\sigma, \sigma | h) \geq \underline{v} - \epsilon \quad \forall h.$

example (Prisoner's dilemma)

	C	D
c	2,2	-1,3
d	3,-1	0,0

here $\underline{v} = 2$
 So the above theorem implies almost efficient outcomes.

example (Battle of the sexes)

	B	S
B	0,0	1,2
S	2,1	0,0

$\underline{v} = 1$

Proof: Suppose $u_i^{r,p}(\sigma, \sigma | h) < \underline{v} - \epsilon$ for a sequence $(r, p) \rightarrow (0, 0)$.

To Do: Find σ' which invades for r, p sufficiently small.

choose $h^* \in \arg \min_h u_i^{r,p}(\sigma, \sigma | h)$. If multiple such histories, choose h^* where the payoff of the other player is maximized.

Finite complexity $\Rightarrow h^*$ exists.

Define $\sigma'(h) = \sigma(h)$ for any h which is not a continuation of h^* or $\pi(h^*)$. Note $\pi(h^*) \neq h^*$ in general.

e.g. For "ALT" h^* is any history where the alternating pattern is broken

After h^* , σ' "signals" willingness to cooperate by playing $\sigma'(h^*) \neq \sigma(h)$.
 If any other player "responds" (by playing differently from σ) in the period after h^*+1 , then σ' goes to a strongly efficient point (v^*, v^{**}) , where

$$w^* = v^* + v^{**}$$

↑
max possible sum of payoffs

If other player does not respond ~~to~~ (i.e. the other player follows σ in period after the signal), then σ' plays like σ .

After $\pi(h^*)$, σ' plays $\sigma(\pi(h^*))$ the first period if the other player has played $\neq \sigma(h^*)$ that period, then σ' goes to strongly efficient play, (i.e. goes to v^{**}); otherwise σ' plays like σ .

Since σ' performs exactly like σ in all histories except h^* and $\pi(h^*)$. Consider play after h^* :

	σ	σ'
σ	x	x
σ'	$\approx x$	v^{**}

Consider play after $\pi(h^*)$:

	σ	σ'
σ	y	$\approx y$
σ'	y	v^{**}

σ' facing history h^* .

Note $x+y < w^*$
since $u_i^{r,p}(\sigma, \sigma | h) < v - \epsilon //$

Averaging (or adding) the above boxes we get:

	σ	σ'
σ	$< w^*$	$< w^*$
σ'	$< w^*$	w^*

Thus, σ^* can invade successfully. ■

Theorem (Existence of evolutionarily stable strategies)

Let $(v, v) \in V^*$ and $v > \underline{v}$. For all $\epsilon > 0$, $\exists q > 0$ and $r^* > 0$ such that $\forall r < r^* \exists p^*$ such that $\forall p < p^* \exists$ an evolutionarily stable σ wrt q, r, p for which

$$|u_i^{r,p}(\sigma, \sigma) - v| < \epsilon.$$

given r , need p to be relatively small.

Proof Construction: In normal phase, play whatever gets you (v, v) .

If player i deviates, minimax him long enough to wipe out gain, then go to strongly efficient payoff which is worse for the deviator than (v, v) . - make punishment phase as short as possible.

Proof: (continued)

Let us see this concept by examining a number of games.

Consider the prisoners' dilemma:

* σ plays c in first period.

Then play c if previously (c,c) or (d,d);

play d if (c,d) or (d,c) was played.

↳ minimizes the punishment time.

	C	D
c	2, 2	-1, 3
d	3, -1	0, 0

Consider the possible invasion by a mutant.

Such a mutant can only invade when the above strategy plays (d,D) on purpose (discard mistakes).

Let σ' be the mutant strategy.

behaves like σ until first deviation from (c,c) after that play c for 1 period. If other player has done the same, then play c forever, else σ' follows σ .

Now consider the meeting of (σ, σ'). After a mistake:

→ There is a loss of 1 ^{on avg.} in the period of a mistake

→ There is a loss of 2 in the next period

Total loss 3. occurs w.p. p.

Thus expected value of losses is: $3p \left(\frac{r}{1+r} \right)$

Consider the meeting of (σ, σ'). After the first mistake:

→ There will be no response

consider a meeting of (σ', σ').

→ ~~no~~ response and then cooperate.

	σ	σ'
σ	$3p \left(\frac{r}{1+r} \right)$	$3p \left(\frac{r}{1+r} \right)$
σ'	$3p \left(\frac{r}{1+r} \right)$	$1p \left(\frac{r}{1+r} \right)$

σ' cannot get a toe-hold.

Consider a modified Battle of the sexes:

σ plays C as long as no one deviates from c.

If a mistake occurs go to (b,a) if 1 made the mistake and (a,b) if 2 made the mistake.

Ignore deviations in punishment phase.

All punishments are strongly efficient.

Even though (c,c) is not strongly efficient, it is evolutionarily stable.

	A	B	C
a	0, 0	4, 1	0, 0
b	1, 4	0, 0	0, 0
c	0, 0	0, 0	2, 2

Note: (c,c) is not strongly efficient.

Private Monitoring

Consider a very simplified game of 2 periods:

	W	S
W	2, 2	-1, 3
S	3, -1	0, 0

	G	B
g	3, 3	0, 0
b	0, 0	1, 1

Instead of knowing payoffs, players will get a signal.

If there was public monitoring and two signals \bar{y}, \underline{y} then

$$P(\bar{y} | s_1, s_2) = \begin{cases} p & \text{if } (w, w) \\ q & \text{if } (w, s) \text{ or } (s, w) \\ r & \text{if } (s, s) \end{cases}$$

play (g, G)
in second period

play (b, B)
in second period.

$$p > q > r.$$

To get an equilibrium where players work, need:

$$2 + 3p + (1-p) \geq 3 + 3q + 1 - q \Rightarrow \underline{2(p - q) \geq 1}$$

Keep this assumption.

With private monitoring — hard to say something for general monitoring structures.
players get private signals — either \bar{z} or \underline{z} . ↓
consider specific cases

Interested in characterizing:

	\bar{z}	\underline{z}
\bar{z}	$\eta(1-\epsilon)$	$\epsilon/2$
\underline{z}	$\epsilon/2$	$(1-\eta)(1-\epsilon)$

If $z_1 \in (\bar{z}, \underline{z})$ is highly correlated with $z_2 \in (\bar{z}, \underline{z})$ then we have

ALMOST PUBLIC MONITORING

To relate this to public monitoring:

$$\eta = \begin{cases} p & \text{if } (w, w) \\ q & \text{if } (w, s) \text{ or } (s, w) \\ r & \text{if } (s, s) \end{cases}$$

CONDITIONAL INDEPENDENCE

— the signal doesn't tell you anything about the opponent's actions.

$$P(z_1, z_2 | s_1, s_2) = P(z_1 | s_2) P(z_2 | s_1)$$

if (w, w)

	\bar{z}	\underline{z}		\bar{z}	\underline{z}
\bar{z}	$(1-\epsilon)^2$	$(1-\epsilon)\epsilon$	$1-\epsilon$	$(1-\epsilon)\epsilon$	$(1-\epsilon)^2$
\underline{z}	$(1-\epsilon)\epsilon$	ϵ^2	ϵ	ϵ^2	$(1-\epsilon)\epsilon$
	$1-\epsilon$	ϵ		$1-\epsilon$	ϵ

as $\epsilon \rightarrow 0$, we get perfect monitoring

A Flavour of the results:

Almost Public Monitoring

Goal: Construct an equilibrium in which player i plays:

- * W in 1st period
- * G if \bar{z}
- * B if \underline{z}

Working backwards: If you are player 1, you want to compute, for stage 2 the probability, $P(z_2 = \bar{z} | z_1 = \bar{z})$.

If $P(z_2 = \bar{z} | z_1 = \bar{z}) > \frac{1}{4}$ you would play G.

$\frac{P(1-\epsilon)}{P(1-\epsilon) + \epsilon/2}$ satisfied if ϵ is small.

In stage 1, to get players to work need:

$$2 + 3P(\bar{z}, \bar{z} | (w, w)) + 1P(\underline{z}, \underline{z} | (w, w)) \geq 3 + 3P(\bar{z}, \bar{z} | (s, w)) + P(\underline{z}, \underline{z} | (s, w))$$

~~$\Rightarrow 2 + 3P(\bar{z}, \bar{z} | (w, w)) + 1P(\underline{z}, \underline{z} | (w, w)) \geq 3 + 3P(\bar{z}, \bar{z} | (s, w)) + P(\underline{z}, \underline{z} | (s, w))$~~

$\Rightarrow 2(P-\epsilon)(1-\epsilon) \geq 1$

If public monitoring inequality holds strictly, then for ϵ small enough, we can replicate the public monitoring case (at least in the two-period version).

↓
In the infinitely repeated games need to confine ourselves to strategies which only look at finite histories.

Conditional Independence

Claim: If $\epsilon > 0$, \nexists equilibrium in pure strategies in which (w, w) is played in the first period.

ABWOC (w, w) is played in the first period. Want player i to play:

- * G if \bar{z}
- * B if \underline{z}

However if ϵ is small i will want to play G, regardless of the signal received.

↳ signals are useless in dictating second period play.

However, consider the case where the players are mixing in the first period.

Consider the case where players randomize between w and S with probabilities α and $(1-\alpha)$

Then consider the strategy for i :

- played w in 1st period and observed \bar{z} , play G
- played w in 1st period and observed \underline{z} , play B .
- played S in 1st period, play B in 2nd period.

As $\epsilon \rightarrow 0$, we can get very close to w, W in the 1st period.