Reputation

Extensive Form Examples

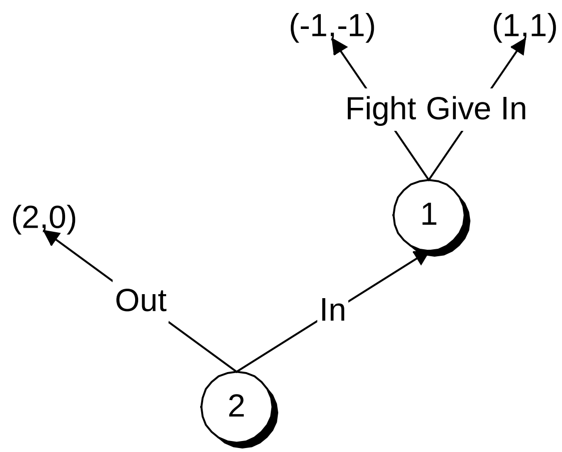

Chain Store Game

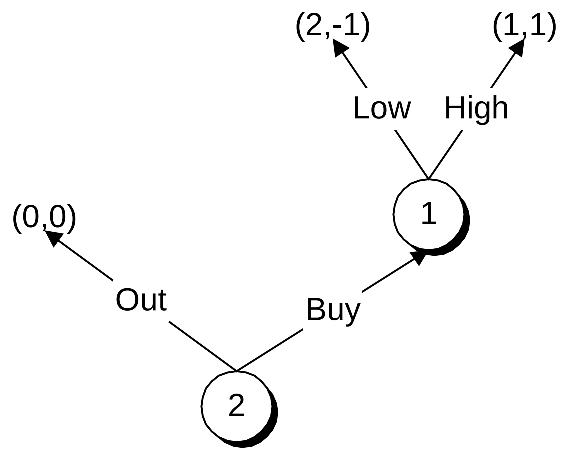

Quality Game

Simultaneous Move Examples

Modified Chain Store

	out	in
fight	2-ε, 0	-1,-1
give in	2,0	1,1**

Inflation Game

	Low	High
Low	0,0	-2,-1
High	1,-1	-1,0

Inflation Game: LR=government, SR=consumers

consumer preferences are whether or not they guess right

	Low	High
Low	0,0	0,-1
High	-1,-1	-1,0

with a hard-nosed government

The Model

multiple types of long-run player $\omega \in \Omega$

Ω is a countable set of types

type is fixed forever (does not change from period to period)

$u^1(a, \omega)$ utility depends on type

strategy $\sigma^1(h, \omega)$ depends on type

types are privately known to long-run player, not known to short run player

strategy $\sigma^2(h)$ does not depend on type

μ probability distribution over Ω commonly known short-run player prior over types

Truly Committed Types

type $\omega(a^1)$ has a dominant strategy to play a^1 in the repeated game:

$$u^1(\tilde{a}^1, a^2, \omega(a^1)) = \begin{cases} 1 & \tilde{a}^1 = a^1 \\ 0 & \tilde{a}^1 \neq a^1 \end{cases}$$

for example

A truly committed type “can't be bargained with…can't be reasoned with…doesn't feel pity, or remorse, or fear.”

Let $n(\omega)$ be the least utility received by a type ω in any Nash equilibrium

Fix a type ω_0: let a^1* be a pure strategy Stackelberg strategy for that type with corresponding committed type $\omega* \equiv \omega(a^1*)$ and utility

$u^1* = \max_{\alpha^1} \min_{\alpha^2 \in BR(\alpha^1)} u^1(\alpha^1, \alpha^2, \omega_0)$

Note the use of min rather than max in the second spot

let $\underline{u}^1$ be the least possible utility for type ω_0

Theorem: Suppose $\mu(\omega_0) > 0$ and $\mu* \equiv \mu(\omega*) > 0$. Then there is a constant $k(\mu*)$ otherwise independent of μ, Ω and δ such that

$$n(\omega) \geq \delta^{k(\mu*)} u^1* + (1 - \delta^{k(\mu*)}) \underline{u}^1$$

Proof

define $\pi_t *$ to be the probability at the beginning of period t by the short-run player that the long-run player will play $a^1 *$

Let $N(\pi_t * \leq \overline{\pi})$ be the number of times $\pi_t * \leq \overline{\pi}$

Lemma 1: There is $\overline{\pi} < 1$ such that if $\pi_t * > \overline{\pi}$ the SR player plays a best response to $a^1 *$

- Why?

Lemma 2: Suppose that LR plays $a^1 *$ always. Then for any history h that has positive probability

$$pr(N(\pi_t * \leq \overline{\pi}) > \log \mu * / \log \overline{\pi} \mid h) = 0$$

- Why do the Lemmas imply the Theorem?

Proof of Lemma 1

Bayes Law

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})\pi(h_t \mid \omega^*, h_{t-1})}{\pi(h_t \mid h_{t-1})}$$

given h_{t-1} player 1 and 2 play independently

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})\pi(h_t \mid \omega^*, h_{t-1})}{\pi(h_t^1 \mid h_{t-1})\pi(h_t^2 \mid h_{t-1})}$$

repeating from the previous page

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})\pi(h_t \mid \omega^*, h_{t-1})}{\pi(h_t^1 \mid h_{t-1})\pi(h_t^2 \mid h_{t-1})}$$

since player 1's type isn't known to player 2 rewrite denominator

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})\pi(h_t \mid \omega^*, h_{t-1})}{\pi(h_t^1 \mid h_{t-1})\pi(h_t^2 \mid \omega^*, h_{t-1})}$$

player 1's strategy is to always play $a^1 *$ so $\pi(h_t^1 \mid \omega^*, h_{t-1}) = 1$

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})\pi(h_t^2 \mid \omega^*, h_{t-1})}{\pi(h_t^1 \mid h_{t-1})\pi(h_t^2 \mid \omega^*, h_{t-1})}$$

$$= \frac{\pi(\omega^* \mid h_{t-1})}{\pi(h_t^1 \mid h_{t-1})} = \frac{\pi(\omega^* \mid h_{t-1})}{\pi_t *}$$

the conclusion reiterated:

$$\pi(\omega^* \mid h_t) = \frac{\pi(\omega^* \mid h_{t-1})}{\pi_t \ ^*}$$

- what does this say?

the Lemma now derives from the fact that $\pi(\omega^* \mid h_t) \leq 1$

Observational Equivalence

$\rho(y \mid \alpha)$ outcome function

$\alpha^2 \in W(\alpha^1)$ if there exists $\tilde{\alpha}^1$ such that $\rho(\cdot \mid \tilde{\alpha}^1, \alpha^2) = \rho(\cdot \mid \alpha^1, \alpha^2)$ and $\alpha^2 \in BR(\tilde{\alpha}^1)$

$$u^{1*} = \max_{\alpha^1} \min_{\alpha^2 \in W(\alpha^1)} u^1(\alpha^1, \alpha^2, \omega_0)$$

here the min can have real bite

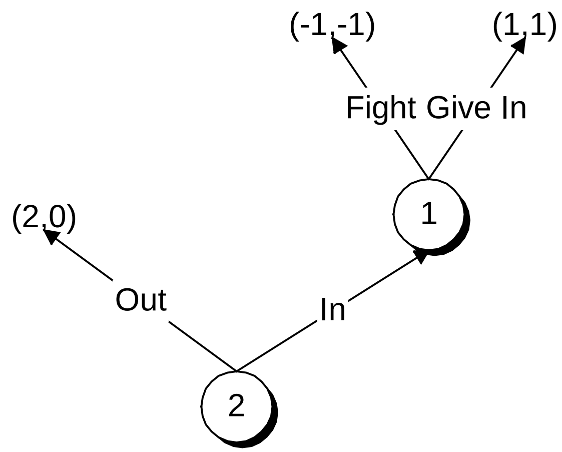

Chain Store Game

strategies that are observationally equivalent

	out	in	mixed
fight	all	fight	fight
give	all	give	give
mixed	all	mixed	mixed

weak best responses

fight: out

give: in, out

mixed: in, out?

Best case fight:out so $u^{1*} = 2$

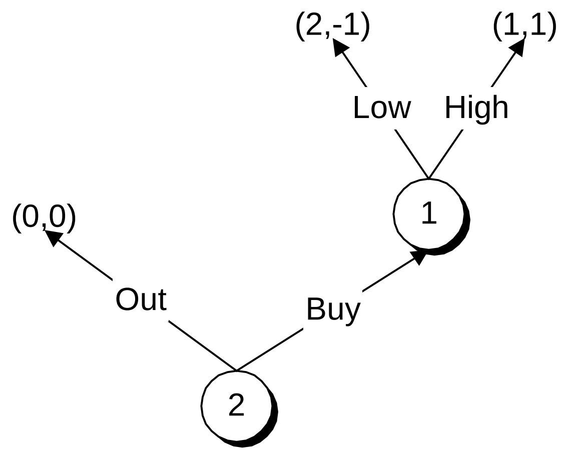

Quality Game

strategies that are observationally equivalent

	out	buy	mixed
hi	all	hi	hi
lo	all	lo	lo
mixed	all	mixed	mixed

weak best responses

hi: in, out

lo: out

mixed: in?, out

in every case out is a weak best response so $u^{1*} = 0$

Moral Hazard and Mixed Commitments

$\rho(y \mid \alpha)$ outcome function

expand space of types to include types committed to mixed strategies: leads to technical complications because it requires a continuum of types

$p(h_{t-1})$ probability distribution over outcomes conditional on the history (a vector)

$p^+(h_{t-1})$ probability distribution over outcomes conditional on the history and the type being in Ω^+

Theorem: for every $\varepsilon > 0, \Delta_0 > 0$ and set of types Ω^+ with $\mu(\Omega^+) > 0$ there is a K such that if Ω^+ is true there is probability less than ε that there are more than K periods with

$$\| p^+(h_{t-1}) - p(h_{t-1}) \| > \Delta_0$$

look for tight bounds

let $\underline{n}, \overline{n}$ be best and worst Nash payoffs to LR

try to get

$$\liminf_{\delta\to 1} \underline{n}(\omega) = \limsup_{\delta\to 1} \overline{n}(\omega) = \max_{\alpha^2 \in BR(\alpha^1)} u^1(\alpha)$$

game is *non-degenerate* if there is no undominated pure action a^2 such that for some $\alpha^2 \neq a^2$

$$u^i(\cdot, a^2) = u^i(\cdot, \alpha^2)$$

counterexample: player 2 gets zero always, player 1 gets either zero or one depending only on player 2's action game is *identified* if for all α^2 that are not weakly dominated $\rho(\cdot \mid \alpha^1,\alpha^2) = \rho(\cdot \mid \tilde{\alpha}^1,\alpha^2)$ implies $\alpha^1 = \tilde{\alpha}^1$

$$\rho(\cdot \mid \alpha^1,\alpha^2) = \alpha^1 R(\alpha^2)$$

condition for identification $R(\alpha^2)$ has full row rank for all α^2

Patient Short Run Players: Schmidt

short run preferences $\begin{bmatrix} 10 & 0 \\ 0 & 1 \end{bmatrix}$

long run preferences

$\mu^0 = 0.1$ $\begin{bmatrix} 10 & 0 \\ 0 & 1 \end{bmatrix}$ pure coordination

$\mu^* = 0.01$ $\begin{bmatrix} 10 & 10 \\ 0 & 1 \end{bmatrix}$ commitment type

$\mu^i = 0.89$ $\begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}$ indifferent type strategies:

normal: play U except if you previously did D, then switch to D

commitment: always play U

indifferent type: U until deviation then D

SR: play L then alternate between R and L (on path)

if 1 deviated to D switch to R forever

if 2 deviated play L; if 1 reacts with U continue with L

reacts with D continue with R

$\delta_1 \geq .15, \delta_2 \geq .75$ then this is a subgame perfect equilibrium

- interesting deviation for SR when supposed to do R deviate to L; but then indifferent type switches to D forever
- for the normal type to prove he's not type "i" he must play D revealing he is not the commitment type

Suppose that LR can minmax SR in a pure strategy $\underline{a}^1$

Theorem: LR gets at least $\min_{\alpha^2 \in BR^2(\underline{a}^1)} u^1(\underline{a}^1, \alpha^2)$

let $\underline{u}^2$ be SR minmax

let $\tilde{u}^2$ be second best against $\underline{a}^1$

$$N = \frac{\ln(1-\delta_2) + \ln(\underline{u}^2 - \tilde{u}^2) - \ln(\overline{u}^2 - \tilde{u}^2)}{\ln \delta_2}$$

$$\varepsilon = \frac{(1-\delta_2)^2(\underline{u}^2 - \tilde{u}^2)}{(\overline{u}^2 - \tilde{u}^2)} - {\delta_2}^N(1-\delta_2)$$

commit to $\underline{a}^1$

Lemma: suppose $a_2^{t+1} \notin BR(\underline{a}^1)$ with positive probability, then SR must believe that in $t+1,\ldots,t+N$ there is a probability of at least ε of not having $\underline{a}^1$

- why is this sufficient?

Proof of Lemma:

2 can get at least $\underline{u}^2$ so

$$(1-\delta_2)u(\alpha^1, a_2^{t+1}) + \delta_2 V \geq \underline{u}^2$$

if $pr(\underline{a}^1) > 1-\varepsilon$ in $t+1, \ldots, t+N$

suppose the optimum is not a br; consider deviating to a br

gain at least $\underline{u}^2 - \tilde{u}^2$ at $t+1$; starting at $t+2$ earn no less than $\underline{u}^2$

if you hadn't deviated, starting at $t+2$ would have earned at most

$$(1-\delta_2)\sum_{t=1}^{N} \delta_2^t (\underline{u}^2(1-\varepsilon) + \varepsilon \overline{u}^2) + \delta_2^{N+1}\overline{u}^2$$

but we chose N and ε so that the loss exceeds the gain