Extensive Form Games

Definition of Extensive Form Game

a finite game tree X with nodes $x \in X$

nodes are partially ordered and have a single root (minimal element)

terminal nodes are $z \in Z$ (maximal elements)

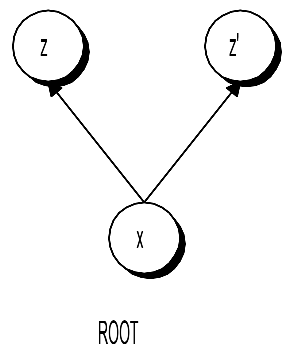

Players and Information Sets

player 0 is nature

information sets $h \in H$ are a partition of $X \setminus Z$

each node in an information set must have exactly the same number of immediate followers

each information set is associated with a unique player who “has the move” at that information set

$H_i \subset H$ information sets where *i* has the move

More Extensive Form Notation

information sets belonging to nature $h \in H_0$ are singletons

$A(h)$ feasible actions at $h \in H$

each action and node $a \in A(h), x \in h$ is associated with a unique node that immediately follows x on the tree

each terminal node has a payoff $r_i(z)$ for each player

by convention we designate terminal nodes in the diagram by their payoffs

Example: a simple simultaneous move game

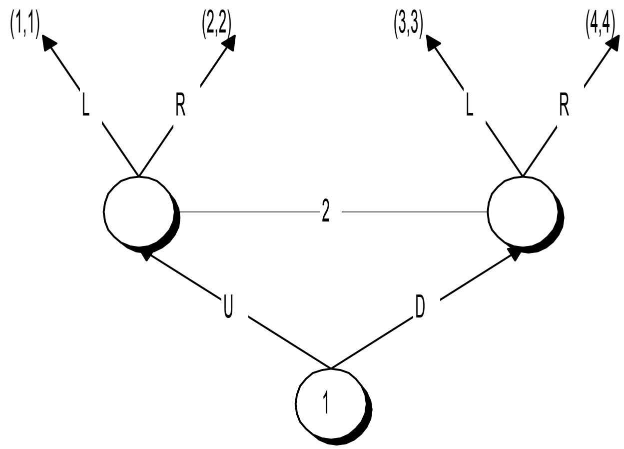

Behavior Strategies

a *pure strategy* is a map from information sets to feasible actions $s_i(h_i) \in A(h_i)$

a *behavior strategy* is a map from information sets to probability distributions over feasible actions $\pi_i(h_i) \in P(A(h_i))$

Nature's move is a behavior strategy for Nature and is a fixed part of the description of the game

We may now define $u_i(\pi)$

normal form are the payoffs $u_i(s)$ derived from the game tree

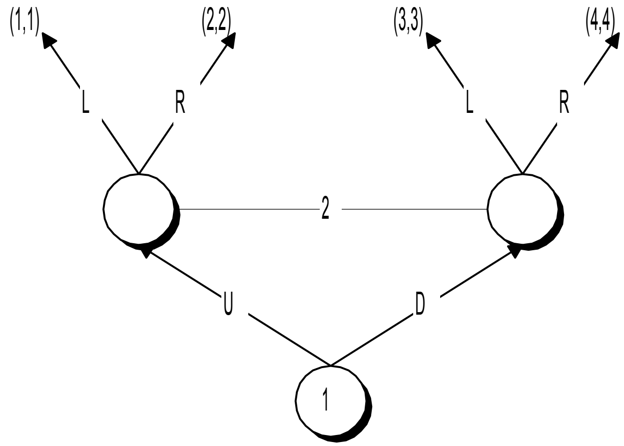

	L	R
U	1,1	2,2
D	3,3	4,4

Kuhn's Theorem:

every mixed strategy gives rise to a unique behavior strategy

The converse is NOT true

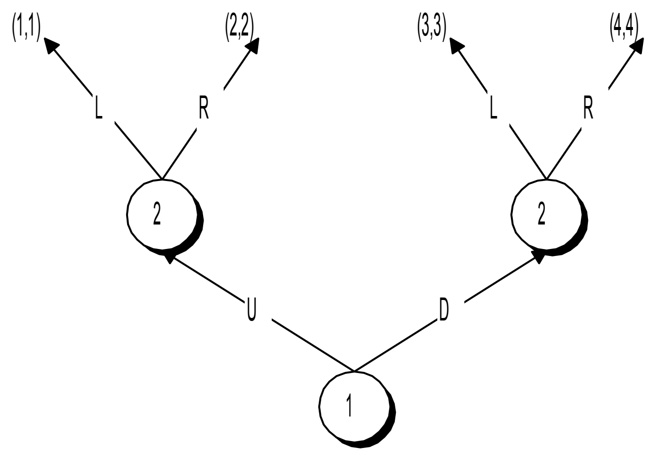

1 plays .5 U

behavior: 2 plays .5L at U; .5L at R

mixed: 2 plays .5(LL),.5(RR)

2 plays .25(LL),.25(RL),.25(LR),.25(RR)

however: if two mixed strategies give rise to the same behavior strategy, they are *equivalent*, that is they yield the same payoff vector for each opponents profile $u(\sigma_i, s_{-i}) = u(\sigma'_i, s_{-i})$

Subgame Perfection

some games seem to have too many Nash equilibria

Ultimatum Bargaining

player 1 proposes how to divide $10 in pennies

player 2 may accept or reject

Nash: any proposal by player 1 with all poorer proposals rejected and equal or better proposals accepted

Chain Store

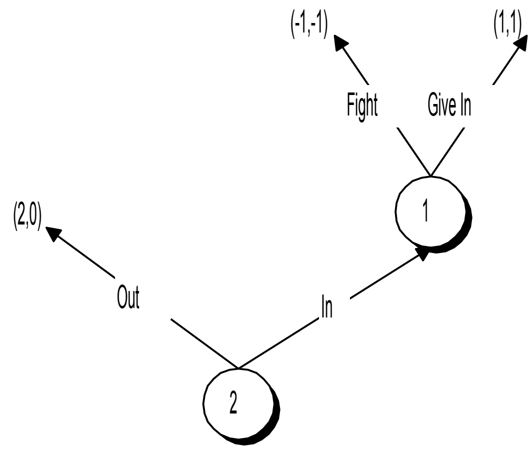

	out	in
fight	2*,0*	-1,-1
give in	2,0	1*,1*

Subgame Perfection

A *subgame perfect Nash Equilibrium is a Nash equilibrium in every subgame*

A subgame starts at a singleton information set

Selten Game

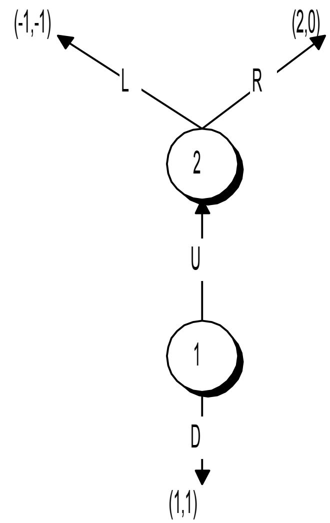

	L	R
U	-1,-1	2*,0*
D	1*,1*	1,1

equilibria:

UR is subgame perfect

D and .5 or more L is Nash but not subgame perfect

Application to Rubinstein Bargaining

the pie division game: there is one unit of pie; player 1 demands p_1

player 2 accepts or rejects

if player 2 rejects one period elapses, then the roles are reversed, with player 2 demanding p_2

common discount factor $0 < \delta < 1$

Nash equilibrium: player 1 gets all pie, rejects all positive demands by player 2; player 2 indifferent and demands nothing

conversely: player 2 gets all the pie

wait 13 periods then split the pie 50-50; if anyone makes a positive offer during this waiting period, reject then revert to the equilibrium where the waiting player gets all the pie

subgame perfection: one player getting all pie is not an equilibrium: if your opponent must wait a period to collect all pie, he will necessarily accept demand of $1-\delta-\varepsilon$ today, since this give him $\delta+\varepsilon$ in present value, rather than δ the present value of waiting a period

Rubinstein's Theorem:

there is a unique subgame perfect equilibrium

players always make the same demands, and if they demand no more than the equilibrium level their demands are accepted

to compute the unique equilibrium observe that a player may reject an offer, wait a period, make the equilibrium demand of *p* and have it accepted, thus getting δp today; this means the opposing player may demand up to $1-\delta p$ and have the demand accepted; the equilibrium condition is

$p = 1 - \delta p$ or $p = \frac{1}{1+\delta}$

notice that the player moving second gets

$\frac{\delta}{1+\delta}$ and that as $\delta \to 1$ the equilibrium converges to a 50-50 split

Uniqueness

a problem: if offers are in pennies, subgame perfect equilibrium is not unique

How to prove the equilibrium is unique:

let p be such that any higher demand will be rejected in every equilibrium

let q be such that any lower demand will be accepted in every equilibrium

if you accept p you get $1-p$ versus at least δq by rejecting, so $1-\delta q$ or less will be rejected in any equilibrium and $p \leq 1-\delta q$

if you accept q you get $1-q$ versus at most δp by rejecting so $(1-q)/\delta$ will be accepted in any equilibrium and $p \geq (1-q)/\delta$

moreover $p \geq q$

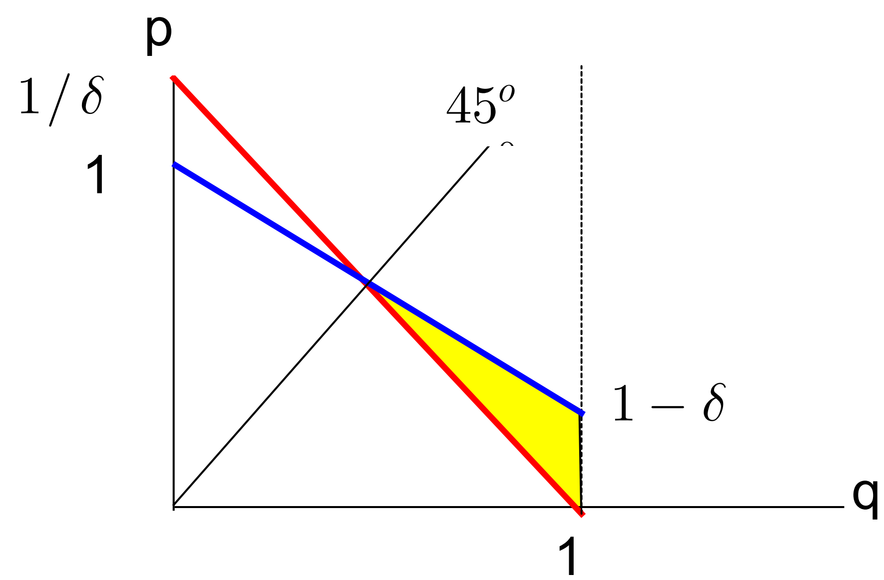
p
$1/\delta$
1
45^o
$1-\delta$
q
1

Trembling Hand Perfection in the Selten Game

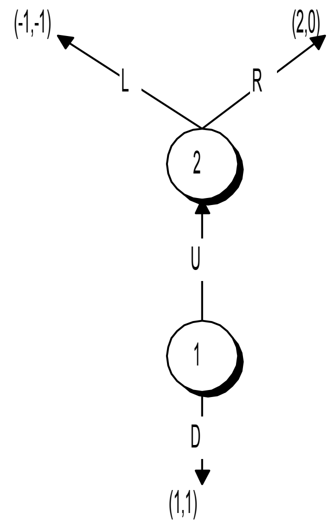

	L	R
U	-1,-1	2,0
D	1,1	1,1

Subgame Perfect Equilibria

UR is subgame perfect

D and .5 or more L is Nash but not subgame perfect

can also solve by weak dominance

or by trembling hand perfection

Example of Trembling Hand not Subgame Perfect

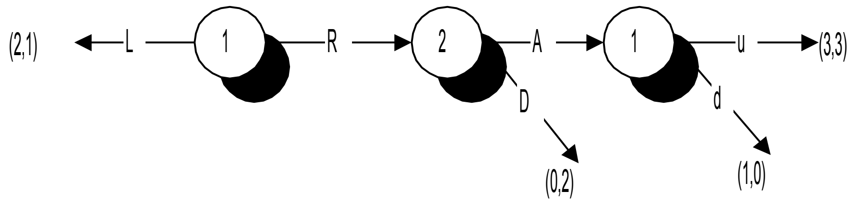

	A	D	
Lu=Ld	2,1	2,1	(n-2)/n
Ru	3,3	0,2	1/n
Fd	1,0	0,2	1/n
	1/n	(n-1)/2	

Here Ld,D is trembling hand perfect but not subgame perfect

Definition of the Agent Normal Form

each information set is treated as a different player, e.g. 1a, 1b if player 1 has two information sets; players 1a and 1b have the same payoffs as player 1

extensive form trembling hand perfection is trembling hand perfection in the agent normal form

what is sequentiality??

Robustness – The Selten Game

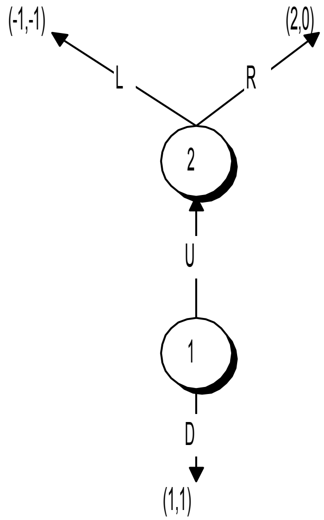

genericity in normal form

	L	R
U	-1,-1	2**,0**
D	1**,1*($\pm\epsilon$)	1,1

Learning

the traffic game

focus on longer term learning and implications for "steady states"

active vs. passive learning

for active learning players must be patient so willing to undertake investment

The Mystery in Human Learning

- not why people learn so badly – why they learn so well.
- behavioral economists, psychologists, economists and computer scientists model human learning using naïve and primitive models.
- models designed by computer scientists to make the best possible decisions cannot come close to the learning ability of the average human child, chimpanzee or even rat.
- equilibrium models and rational expectations: if we have to choose between best models of learning and perfect learning – for most situations of interest to economists perfect learning fits the facts better

Global Convergence?

- "grail" of learning research: global convergence theorem for convincing learning processes
- easy to construct examples of learning processes that don't converge
- non-convergence looks like cob-web; people repeat the same mistakes over and over; not terrifically plausible
- we seem to see much "equilibriumness" around us: traffic, refugee camps

Overview

- stochastic procedures can be globally stable: fishing for Nash equilibrium
- there are procedures of this type that satisfy sensible criteria for being “good”

Worst-case or Universal analysis vs. Bayesian analysis

- opponents may be smarter than you
- their process of optimization may result in play not in the support of your prior
- probability 1 with respect to your own beliefs is not meaningful in the setting of a game
- example: everyone believing that they face a stationary process (a common statistical assumption) implies that no one will actually behave in a stationary way
- these deficiencies in the robustness of Bayes learning are why there is no satisfactory global convergence theorem for Bayesian learning procedures

The Two Armed Bandit

- gambling machine with two arms: $a \in A, B$
- each arms either pays either 0 or 1 and pay 1 with probability π_a
- you only see what happens for the arm you choose!!
- discount factor is δ
- prior is the beta distribution

$$\frac{\pi_a^{\alpha_a - 1}(1-\pi_a)^{\beta_a - 1}}{B(\alpha_a, \beta_a)}$$

where α_a, β_a are parameters and B is the beta function

The Conjugate Prior

•

this is a “conjugate prior” a success results in the posterior

$$\frac{\pi_a^{\alpha_a}(1-\pi_a)^{\beta_a-1}}{B(\alpha_a,\beta_a)}$$

and a failure in

$$\frac{\pi_a^{\alpha_a-1}(1-\pi_a)^{\beta_a}}{B(\alpha_a,\beta_a)}$$

- we can think of α_a, β_a as the number of successes and failures
- (but note that the only requirement of these parameters is that the be strictly positive)

the posterior mean is $\alpha_a/(\alpha_a+\beta_a)$

the fraction of successes...

The Gittins Index

for each arm given the current belief compute the certain amount g_a

that is indifferent to pulling the arm

this is called the Gittins index

optimal strategy:

- pull the arm with the highest Gittins index and recompute
- optimal experimentation
- the Gittins index tells the option value of learning
- if you pull a bad arm and find out you are wrong you can take advantage of that forever at the cost of just a one period trial

actually pretty hard to compute the Gittins index and there are good heuristics known to computer scientists

What Happens?

- With probability one after some finite period of time you pull the same arm forever after
- it isn't necessarily the correct arm
- if δ is close to one the probability of the wrong arm forever is low

consider:

- you wrongly think arm A is a great arm

  you pull it and learn you are wrong

  "on the equilibrium path error"
- you wrongly think arm B is a terrible arm

  you never pull it and never learn you are wrong

  "off the equilibrium path error"

Self Confirming Equilibrium

$s_i \in S_i$ pure strategies for *i;* $\sigma_i \in \Sigma_i$ mixed

H_i information sets for *i*

$\overline{H}(\sigma)$ reached with positive probability under σ

$\pi_i \in \Pi_i$ behavior strategies

$\hat{\pi}(h_i|\sigma_i)$ map from mixed to behavior strategies

$\hat{\rho}(\pi)$, $\hat{\rho}(\sigma) \equiv \hat{\rho}(\hat{\pi}(\sigma))$ distribution over terminal nodes μ_i a probability measure on Π_{-i}

$u_i(s_i|\mu_i)$ preferences

$$\Pi_{-i}(\sigma_{-i}\Big|J) \equiv \{\pi_{-i}\Big|\pi_i(h_i) = \hat{\pi}(h_i|\sigma_i), \forall h_i \in H_{-i} \cap J\}$$

Notions of Equilibrium

Nash equilibrium

a mixed profile σ such that for each $s_i \in \text{supp}(\sigma_i)$ there exist beliefs μ_i such that

- s_i maximizes $u_i(\cdot|\mu_i)$
- $\mu_i(\Pi_{-i}(\sigma_{-i}|H)) = 1$

Unitary Self-Confirming Equilibrium

- $\mu_i(\Pi_{-i}(\sigma_{-i}|\overline{H}(\sigma))) = 1$

(=Nash with two players)

Fudenberg-Kreps Example

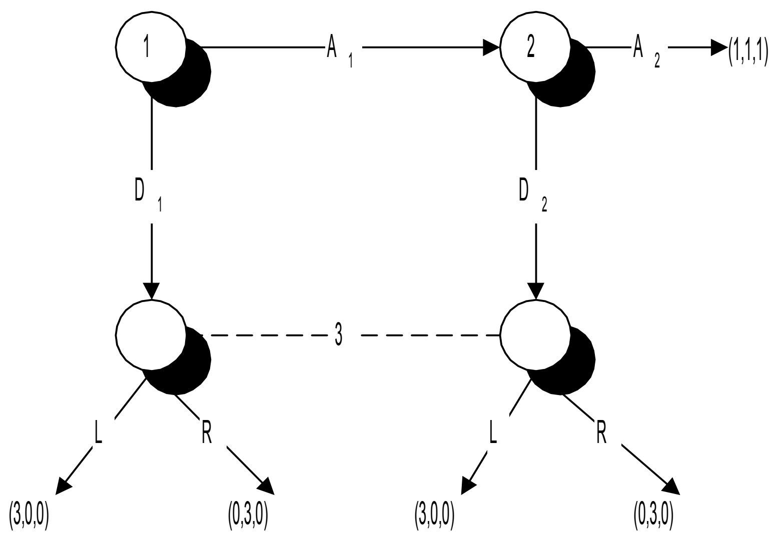

A_1, A_2 is self-confirming, but not Nash

any strategy for 3 makes it optimal for either 1 or 2 to play down

but in self-confirming, 1 can believe 3 plays R; 2 that he plays L

Heterogeneous Self-Confirming Equilibrium

- $\mu_i(\Pi_{-i}(\sigma_{-i} | \overline{H}(s_i, \sigma))) = 1$

Can summarize by means of “observation function”

$$J(s_i, \sigma) = H, \overline{H}(\sigma), \overline{H}(s_i, \sigma)$$

Public Randomization

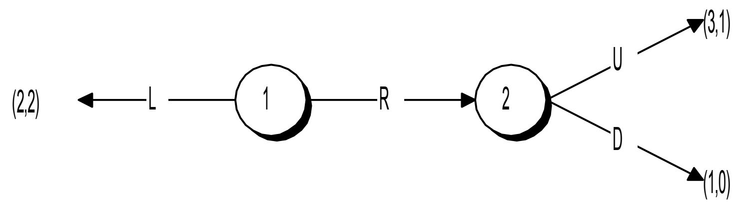

Remark: In games with perfect information, the set of heterogeneous self-confirming equilibrium payoffs (and the probability distributions over outcomes) are convex

Example Without Public Randomization

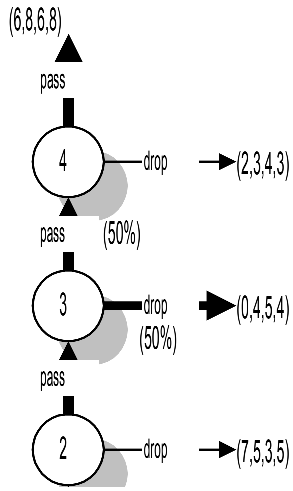

Subgame Perfection and Best Shot

Prasnikar and Roth

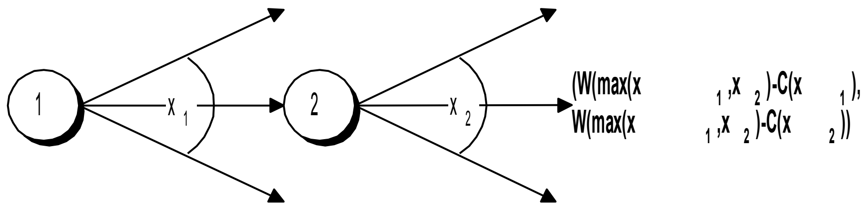

x	*W(x)*	*C(x)*
0	$0.00	$0.00
1	$1.00	$0.82
2	$1.95	$1.64
3	$2.85	$2.46
4	$3.70	$3.28
5	$4.50	$4.10
6	$5.25	$4.92
7	$5.95	$5.74
8	$6.60	$6.50

Discussion of Best Shot

if the other player makes any contribution at all, it is optimal to contribute nothing

unique subgame perfect equilibrium player 1 contributes nothing

another Nash equilibrium player 2 to contributes nothing regardless of player 1's play

Best-Shot Results

Hirshleifer-Harrison partial information, but alternating roles

Prasnikar-Roth fixed roles, both partial and full information

In the full information case and partial information heterogeneous case player 2 occasionally contributes less than 4 when player 1 has contributed nothing; Note that the player who contributes nothing gets $3.70 against $0.42 for the opponent who contributes 4

- full information case: player 1 never contributed anything
- partial information case: sometimes roles reverse

Ultimatum Bargaining

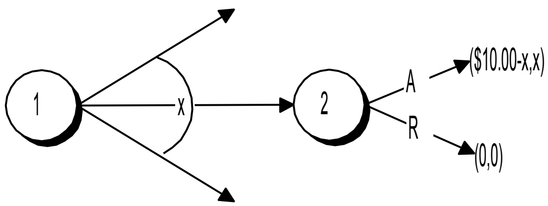

US Data for Ultimatum

x	*Offers*	*Rejection Probability*
$2.60	3	33%
$4.25	13	18%
$5.00	13	0%
	29	

US $10.00 stake games, round 10

Trials	Rnd	Cntry Stake	Case	Expected Loss			Max Gain	Ratio
				Pl 1	Pl 2	Both		
27	10	US	H	$0.00	$0.67	$0.34	$10.00	3.4%
27	10	US	U	$1.30	$0.67	$0.99	$10.00	9.9%
10	10	USx3	H	$0.00	$1.28	$0.64	$30.00	2.1%
10	10	USx3	U	$6.45	$1.28	$3.86	$30.00	12.9%
30	10	Yugo	H	$0.00	$0.99	$0.50	$10?	5.0%
30	10	Yugo	U	$1.57	$0.99	$1.28	$10?	12.8%
29	10	Jpn	H	$0.00	$0.53	$0.27	$10?	2.7%
29	10	Jpn	U	$1.85	$0.53	$1.19	$10?	11.9%
30	10	Isrl	H	$0.00	$0.38	$0.19	$10?	1.9%
30	10	Isrl	U	$3.16	$0.38	$1.77	$10?	17.7%
	WC		H			$5.00	$10.00	50.0%

Rnds=Rounds, WC=Worst Case, H=Heterogeneous, U=Unitary

Comments on Ultimatum

- every offer by player 1 is a best response to beliefs that all other offers will be rejected so player 1's heterogeneous losses are always zero.
- big player 1 losses in the unitary case
- player 2 losses all knowing losses from rejected offers; magnitudes indicate that subgame perfection does quite badly
- as in centipede, tripling the stakes increases the size of losses a bit less than proportionally (losses roughly double).

Centipede Game: Palfrey and McKelvey

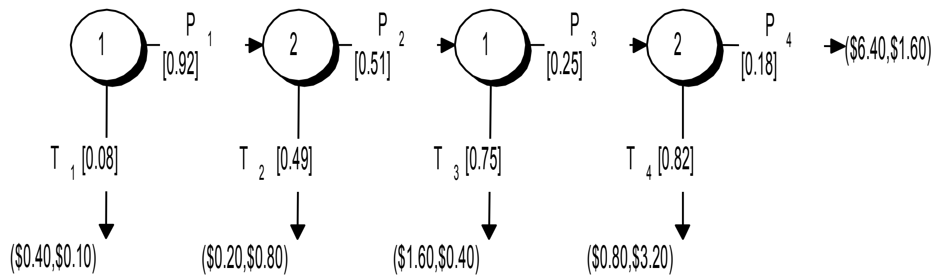

Numbers in square brackets correspond to the observed conditional probabilities of play corresponding to rounds 6-10, stakes 1x below.

This game has a unique self-confirming equilibrium; in it player 1 with probability 1 plays T_1

Summary of Results

Trials / Rnd	Rnds	Stake	Case	Expected Loss			Max Gain	Ratio
				Pl 1	Pl 2	Both		
29*	6-10	1x	H	$0.00	$0.03	$0.02	$4.00	0.4%
29*	6-10	1x	U	$0.26	$0.17	$0.22	$4.00	5.4%
	WC	1x	H			$0.80	$4.00	20.0 %
29	1-10	1x	H	$0.00	$0.08	$0.04	$4.00	1.0%
10	1-10	4x	H	$0.00	$0.28	$0.14	$16.00	0.9%

Rnds=Rounds, WC=Worst Case, H=Heterogeneous, U=Unitary

Comments on Experimental Results

- heterogeneous loss per player is small; because payoffs are doubling in each stage, equilibrium is very sensitive to a small number of player 2's giving money away at the end of the game.
- unknowing losses far greater than knowing losses
- quadrupling the stakes very nearly causes $\bar{\varepsilon}$ to quadruple
- theory has substantial predictive power: see WC
- losses conditional on reaching the final stage are quite large--inconsistent with subgame perfection. McKelvey and Palfrey estimated an incomplete information model where some "types" of player 2 liked to pass in the final stage. This cannot explain many players dropping out early so their estimated model fits poorly.