Copyright (C) 2010 David K. Levine

This document is an open textbook it is released under your choice of

(a) the creative commons license: Attribution-ShareAlike License

(b) version 1 of the open text license amendment to version 2 of the GNU General Public License. The open text license amendment is published by Michele Boldrin et al at http://levine.sscnet.ucla.edu/general/gpl.htm; the GPL is published by the Free Software Foundation at http://www.gnu.org/copyleft/gpl.html.

Learning in Games

Introduction and Basic Concepts

David K. Levine

October 5, 2010

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)

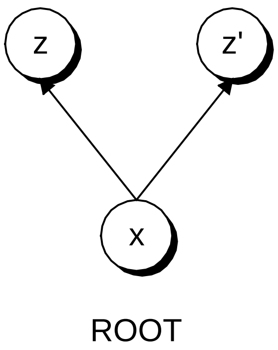

Players and Information Sets

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

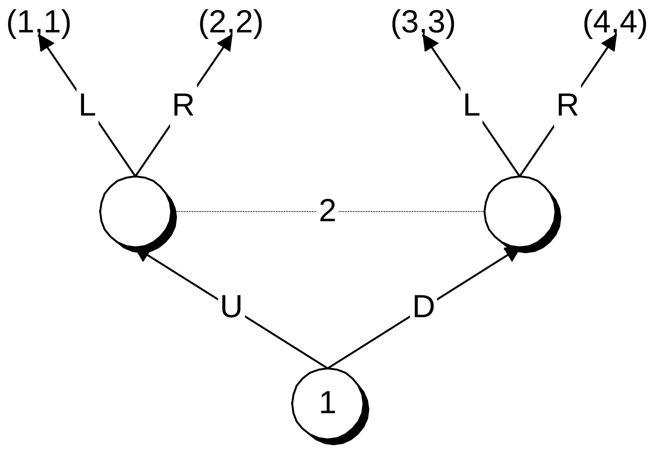

Behavior Strategies

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

S_i are the set of pure strategies

$\sigma_i \in \Sigma_i$ are mixed strategies, probability distributions over pure strategies

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

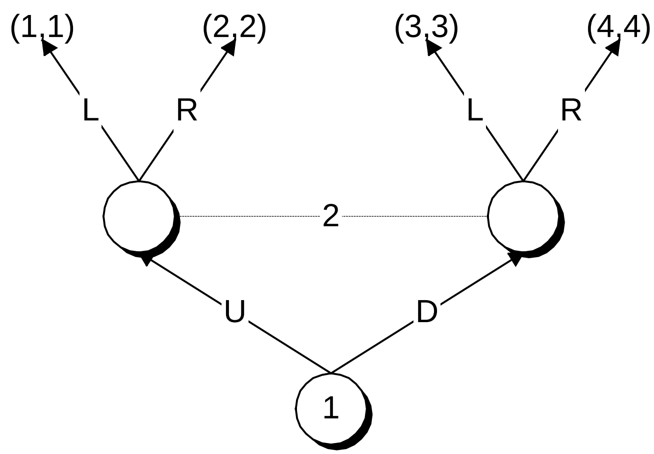

	L	R
U	1,1	2,2
D	3,3	4,4

Kuhn's Theorem

every mixed strategy gives rise to a unique behavior strategy

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

The converse is NOT true

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 $u(\sigma_i, s_{-i}) = u(\sigma'_i, s_{-i})$

Additional Notation

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

$\hat{\rho}(\pi)$, $\hat{\rho}(\sigma) \equiv \hat{\rho}(\hat{\pi}(\sigma))$ distribution over terminal nodes

μ_i a probability measure on Π_{-i}

$u_i(s_i \mid \mu_i)$ preferences

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

Nash Equilibrium

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

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

Why Might We Be At Nash Equilibrium?

The rush hour traffic game

Potential games

Dynamics versus statics: two different questions

- What sort of outcomes can arise from asymptotic of learning? Nash? Self-confirming?
- What does the adjustment path look like?

Focus on statics first

Active versus passive learning

Unitary Self-Confirming Equilibrium

What does learning tell us in extensive form games?

- $\mu_i(\Pi_{-i}(\sigma_{-i} \mid \bar{H}(\sigma))) = 1$

Theorem: *Path equivalent to Nash equilibrium when there are two players*

Why?

Fudenberg-Kreps Example

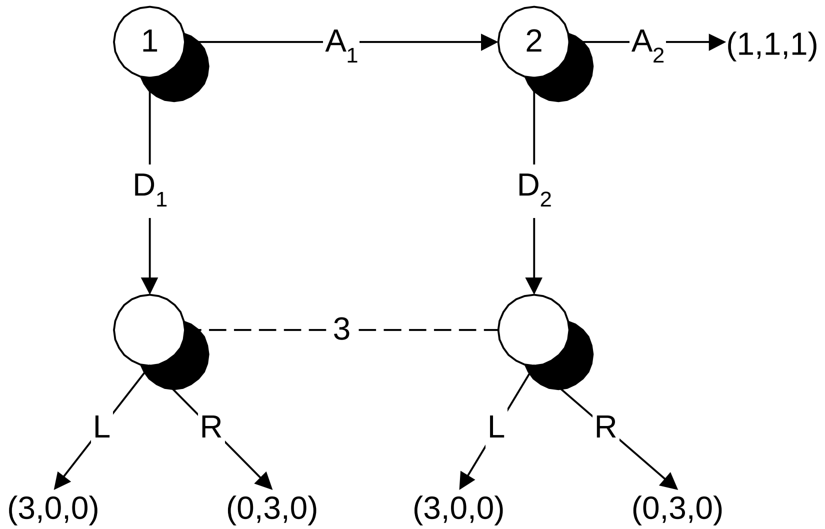

A_1, A_2 is self-confirming, but not Nash

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

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

The “observation function”

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

Public Randomization

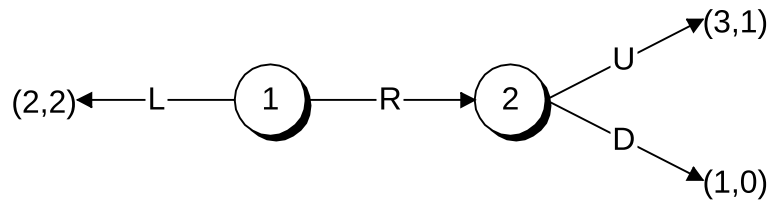

Remark: In games with perfect information, the set of heterogeneous self-confirming equilibrium payoffs (and the probability distributions over outcomes) are convex

Example Without Public Randomization

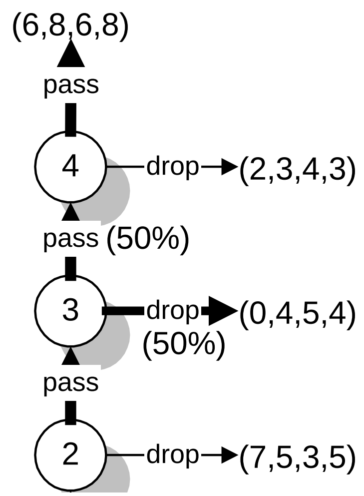

Knowing and Unknowing Losses

The relative importance of learning

Ultimatum Bargaining Results

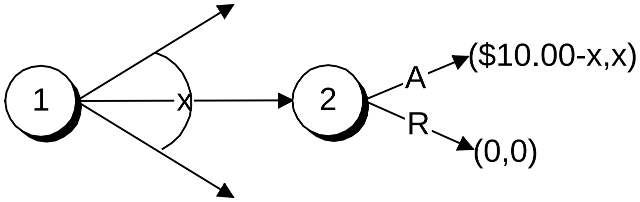

Raw US Data for Ultimatum

x	*Offers*	*Rejection Probability*
$2.00	1	100%
$3.25	2	50%
$4.00	7	14%
$4.25	1	0%
$4.50	2	100%
$4.75	1	0%
$5.00	13	0%
	27	

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; but really a matter of social preference
- tripling the stakes increases the size of losses a bit less than proportionally (losses roughly double)
- key fact: unknowing losses considerably larger than knowing losses – relative importance of learning

Centipede Game: Palfrey and McKelvey

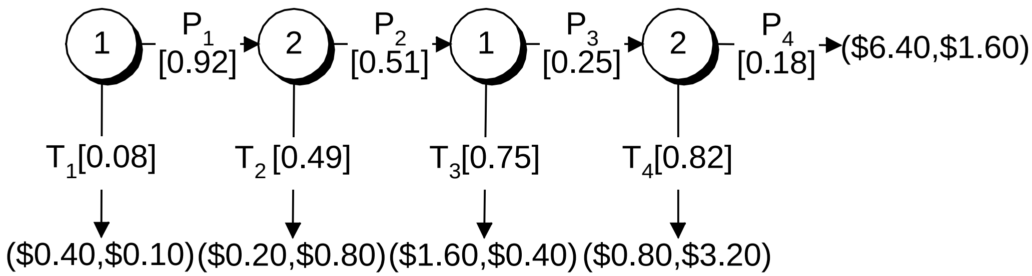

Numbers in square brackets correspond to the observed conditional probabilities of play corresponding to rounds 6-10, stakes 1x below.

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

Summary of Experimental Results

Trials Rnd	Rnds	Stake		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

*The data on which from which this case is computed is reported above.

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" indicative however of social preference. 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.