Agent logit equilibrium

Summer 26

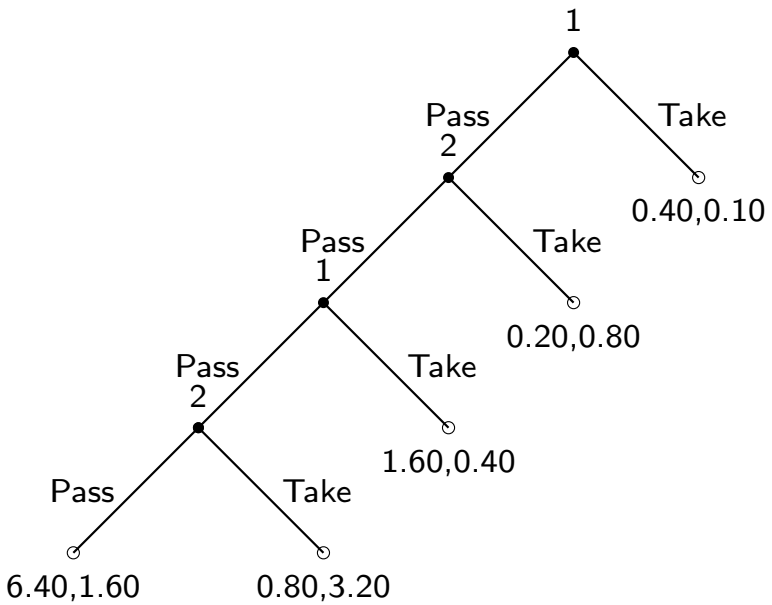
Niels Boissonnet, Bielefeld

Discussion

Centipede Game

Agent logit equilibrium

Centipede Game



By BI, P1 chooses Take and the play stops after the first node.

An Experimental Study of the Centipede Game

Agent logit equilibrium

The experiment

- Ten centipede games of either two rounds (four moves) or of three rounds (six moves).
- Before each game, each subject is matched with a random other subject

Motivation

- BI yields a solution that is intuitively unsatisfactory, so it is worth observing experimentally how subjects behave in this game

- Testing if deviations
  1. are due to existence of altruistic players,
  2. or the fact that some DMs make mistakes.

Results: descriptive statistics

- f_i is the proportion of games ending at the ith node
- p_i is the probability of stopping the game conditional on reaching the ith node

$$p_i = \frac{f_i}{\sum_{j=i}^{n+1} f_j}$$

TABLE IIA

PROPORTION OF OBSERVATIONS AT EACH TERMINAL NO

	Session		N	f_1	f_2	f_3	f_4
Four Move	1	(PCC)	100	.06	.26	.44	.20
	2	(PCC)	81	.10	.38	.40	.11
	3	(CIT)	100	.06	.43	.28	.14
	Total	1–3	281	.071	.356	.370	.153
High Payoff	4	(High-CIT)	100	.150	.370	.320	.110
Six Move	5	(CIT)	100	.02	.09	.39	.28
	6	(PCC)	81	.00	.02	.04	.46
	7	(PCC)	100	.00	.07	.14	.43
	Total	5–7	281	.007	.064	.199	.384

TABLE IIB[a]

IMPLIED TAKE PROBABILITIES FOR THE CENTIPEDE GAME

	Session	p_1	p_2	p_3	p_4	p_5	p_6
Four Move	1 (PCC)	.06 (100)	.28 (94)	.65 (68)	.83 (24)		
	2 (PCC)	.10 (81)	.42 (73)	.76 (42)	.90 (10)		
	3 (CIT)	.06 (100)	.46 (94)	.55 (51)	.61 (23)		
	Total 1–3	.07 (281)	.38 (261)	.65 (161)	.75 (57)		
High Payoff	4 (CIT)	.15 (100)	.44 (85)	.67 (48)	.69 (16)		
Six Move	5 (CIT)	.02 (100)	.09 (98)	.44 (89)	.56 (50)	.91 (22)	.50 (2)
	6 (PCC)	.00 (81)	.02 (81)	.04 (79)	.49 (76)	.72 (39)	.82 (11)
	7 (PCC)	.00 (100)	.07 (100)	.15 (93)	.54 (79)	.64 (36)	.92 (13)
	Total 5–7	.01 (281)	.06 (279)	.21 (261)	.53 (205)	.73 (97)	.85 (26)

Results: descriptive statistics

- very few subjects stop in round 1 $\rightsquigarrow$ backward induction or rationality not satisfied
- probability of stopping is increasing as game is played $\rightsquigarrow$ not simply mistakes

TABLE IIA

PROPORTION OF OBSERVATIONS AT EACH TERMINAL NO

	Session		N	f_1	f_2	f_3	f_4
Four Move	1	(PCC)	100	.06	.26	.44	.20
	2	(PCC)	81	.10	.38	.40	.11
	3	(CIT)	100	.06	.43	.28	.14
	Total	1–3	281	.071	.356	.370	.153
High Payoff	4	(High-CIT)	100	.150	.370	.320	.110
Six Move	5	(CIT)	100	.02	.09	.39	.28
	6	(PCC)	81	.00	.02	.04	.46
	7	(PCC)	100	.00	.07	.14	.43
	Total	5–7	281	.007	.064	.199	.384

TABLE IIB[a]

IMPLIED TAKE PROBABILITIES FOR THE CENTIPEDE GAME

	Session	p_1	p_2	p_3	p_4	p_5	p_6
Four Move	1 (PCC)	.06 (100)	.28 (94)	.65 (68)	.83 (24)		
	2 (PCC)	.10 (81)	.42 (73)	.76 (42)	.90 (10)		
	3 (CIT)	.06 (100)	.46 (94)	.55 (51)	.61 (23)		
	Total 1–3	.07 (281)	.38 (261)	.65 (161)	.75 (57)		
High Payoff	4 (CIT)	.15 (100)	.44 (85)	.67 (48)	.69 (16)		
Six Move	5 (CIT)	.02 (100)	.09 (98)	.44 (89)	.56 (50)	.91 (22)	.50 (2)
	6 (PCC)	.00 (81)	.02 (81)	.04 (79)	.49 (76)	.72 (39)	.82 (11)
	7 (PCC)	.00 (100)	.07 (100)	.15 (93)	.54 (79)	.64 (36)	.92 (13)
	Total 5–7	.01 (281)	.06 (279)	.21 (261)	.53 (205)	.73 (97)	.85 (26)

Results: descriptive statistics

- Subjects stop later in first five games than in last five: learning?
- Number of subjects that adopt the dominated strategy of passing on the last move drops from 25% to 18% and from 19% to 10% – less altruism?

TABLE IIIB

IMPLIED TAKE PROBABILITIES
COMPARISON OF EARLY VERSUS LATE PLAYS IN THE LOW PAYOFF CENTIPEDE GAMES

Treatment	Game	p_1	p_2	p_3	p_4	p_5	p_6
Four Move	1–5	.06	.32	.57	.75		
		(145)	(136)	(92)	(40)		
	6–10	.08	.49	.75	.82		
		(136)	(125)	(69)	(17)		
Four Move	1–5	.00	.06	.18	.43	.75	.81
		(145)	(145)	(137)	(112)	(64)	(16)
	6–10	.01	.07	.25	.65	.70	.90
		(136)	(134)	(124)	(93)	(33)	(10)

Results: descriptive statistics

- Some individuals (5%) passed at every opportunity they have;
- Some subjects' behavior does not seem to make sense indicating they make mistakes.

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

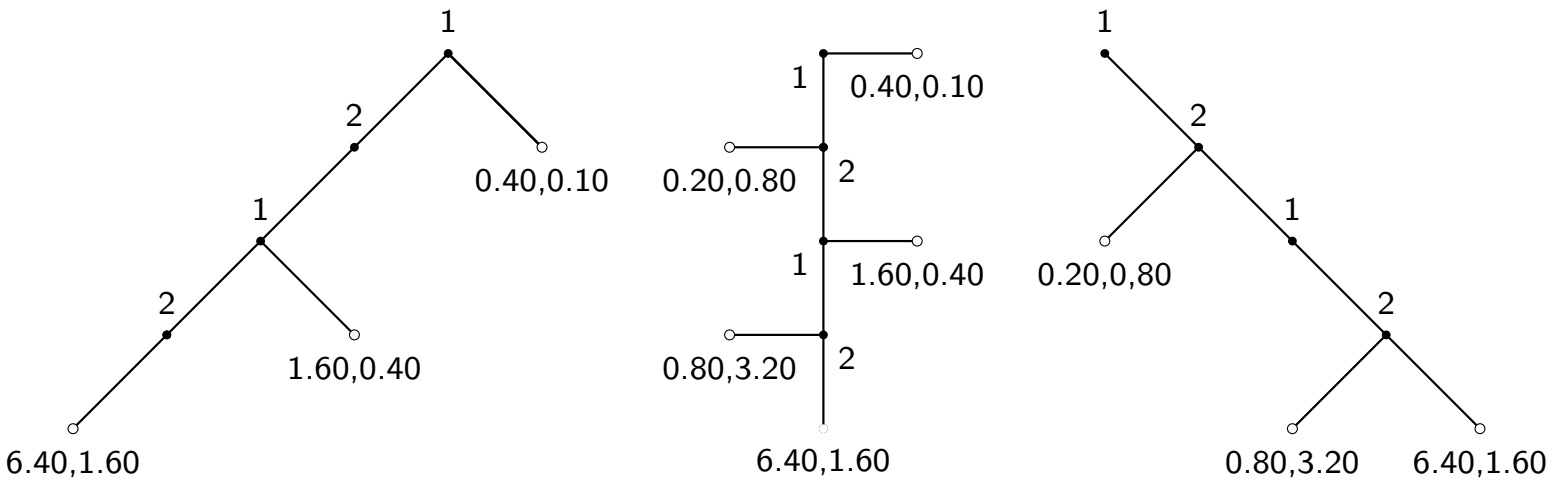


Figure: Extensive form representation

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

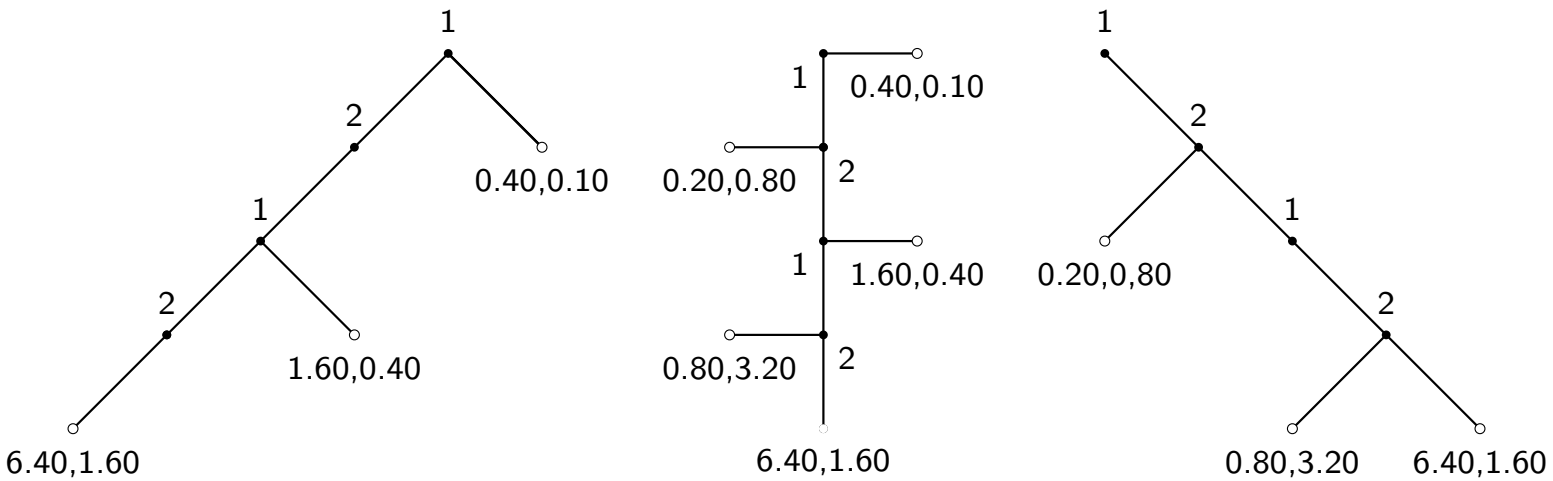


Figure: Extensive form representation

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

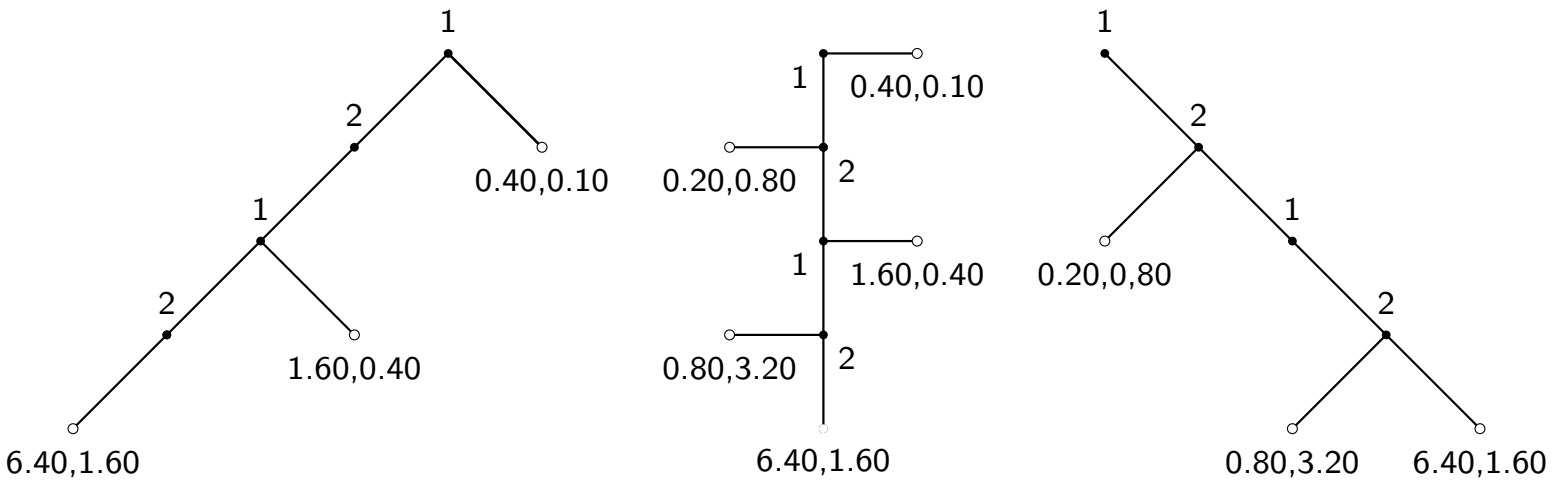


Figure: Extensive form representation

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

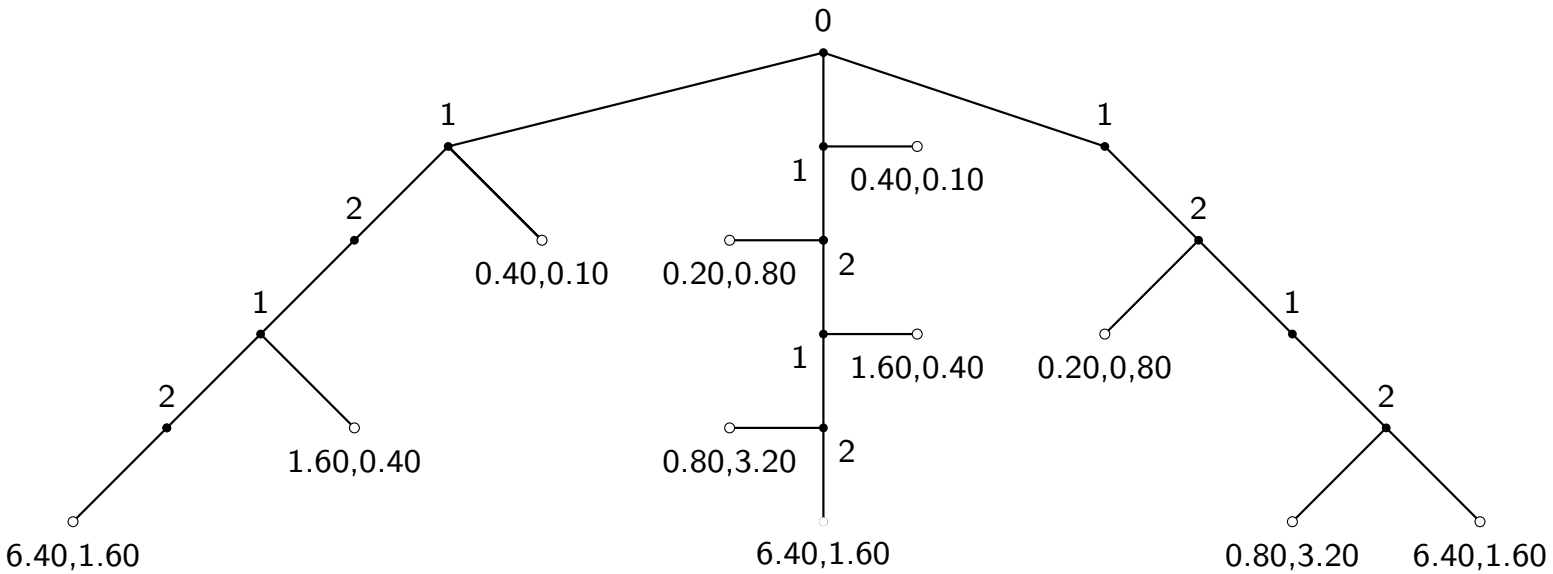


Figure: Extensive form representation

The model

Incomplete information: The opponent is

- selfish with probability q
- altruist with probability $1 - q$

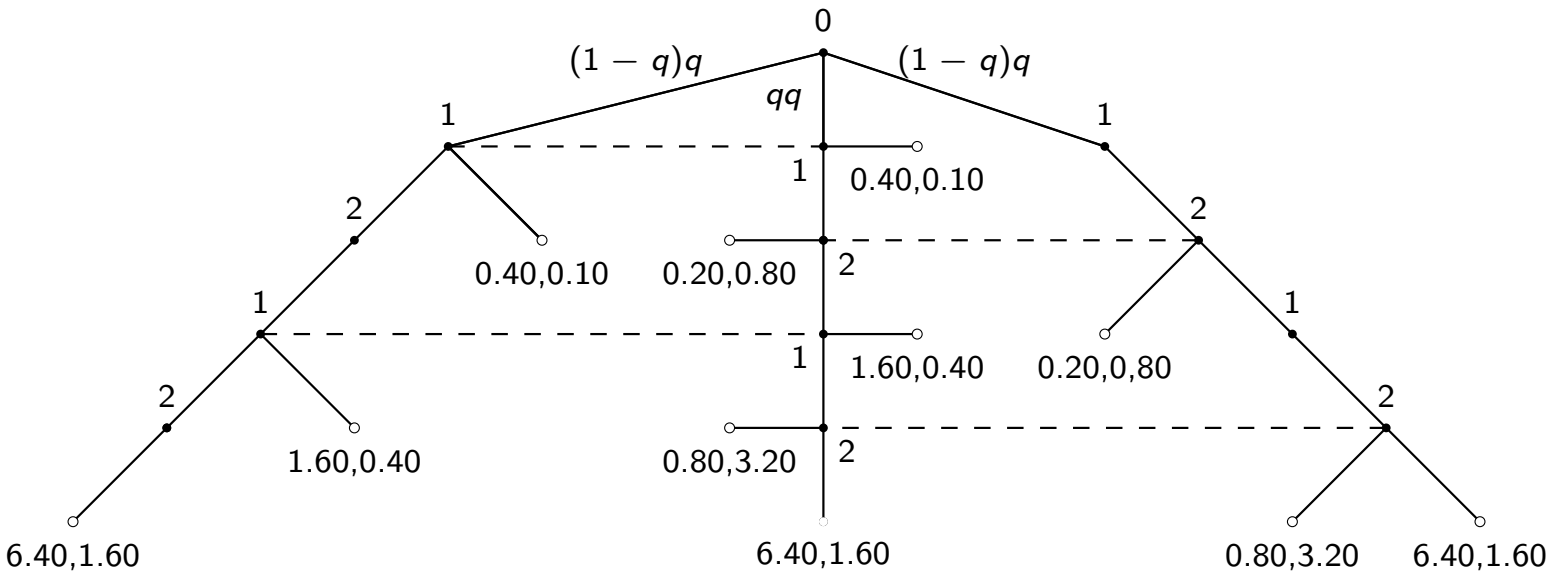


Figure: Extensive form representation

Proposition

Let q be the common belief that a selected player is selfish. For a generic q there exists a sequential equilibrium denoted $p(q)$ with the following properties.

1. *For any q, Player 2 chooses* **Take** *with probability 1 on its last move.*
2. *if $q = 1$ they choose* **Take**
3. *If $1 - q > 1/7$, both Players always choose* **Pass**, *except on the last move*
4. *If $1 - q < 1/7$ the equilibrium involves mixed strategies.*

The form of the imitation in our setting is obvious: selfish players sometimes pass, to mimic an altruist.

Limitations

For any value of q, there is at least one outcome with a 0 or close to 0 probability of occurrence.

- So, the model will fit poorly data in which all of the possible outcomes occur.
- Nor can it account for any consistent patterns of learning in the data, for substantial variations across individuals, or for some of the irregularities described earlier.

To account for these features of the data, they introduce two additional elements:

- *The subject explore the experiment to see what happens;*
- *Subjects have heterogeneous beliefs about their opponents' type.*

Subjects explore the experiment

In game t, let p^* be the equilibrium probability of TAKE

- Noisy play: Assume the player actually TAKE with probability

$$(1 - \epsilon_t)p^*,$$

  and makes a random move (i.e. TAKE or PASS with probability 0.5) with probability ϵ_t.

- Learning: Individuals follow an exponential learning curve with ϵ as the initial error rate.

$$\epsilon_t = \epsilon e^{-\delta(t-1)}$$

Subjects have heterogeneous beliefs about the type

- Q is the true probability that individuals are selfish, but each player i has a belief q_i that her opponent is selfish, where $q_i \neq Q$ is possible;
- Beliefs of the individuals are independently drawn from a Beta distribution with parameters (α, β), where the mean of the distribution is simply equal to some $q = \frac{\alpha}{\alpha+\beta}$;
- Egocentric model hypothesis: Each player plays the game as if it were common knowledge that the opponent had the same belief, i.e., while the econometrician assumes there is heterogeneity in beliefs, players do have heterogeneous beliefs, but believe that everyone's beliefs are alike;
- Rational expectations hypothesis: The average belief $\frac{\alpha}{\alpha+\beta}$ equals the true proportion of selfish players.

$$\mathbb{E}[q_i] = \frac{\alpha}{\alpha+\beta} = Q$$

Estimation

- If Q is the true population parameter for the fraction of selfish players, then the likelihood of observing i choosing D_i, conditional on her belief q, is

$$\pi_i(D_i|Q,\epsilon,\delta,q) = QPr_i^{selfish}(D_i|\epsilon,\delta,q) + (1-Q)Pr_i^{altruist}(D_i|\epsilon,\delta,q)$$

- Since q is drawn from the Beta distribution with parameters (α,β), and density $B(q|\alpha,\beta)$, then the likelihood of observing D_i without conditioning on q is given by:

$$s_i(D_i|Q,\epsilon,\delta,\alpha,\beta) = \int_0^1 \pi_i(D_i|Q,\epsilon,\delta,q)\cdot B(q|\alpha,\beta)\ dq$$

- Parameters are estimated by maximizing the likelihood

$$L(D|Q,\epsilon,\delta,\alpha,\beta) = \sum_{i=1}^{I} log\big(s_i(D_i|Q,\epsilon,\delta,\alpha,\beta)\big)$$

Tested hypotheses

- Errors in action: Is ϵ different from 0?
- Heterogeneity of beliefs: Is the variance of the estimated distribution of priors significantly different from 0 ?
- Rational Expectations: Is the estimated value of Q equal to the mean of the estimated distribution of priors, $\frac{\alpha}{\alpha+\beta}$?
- Learning: Is the estimated value of δ positive and significantly different from 0?

Estimation Results and Tests

- $\hat{\mu} = \frac{\hat{\alpha}}{\hat{\alpha}+\hat{\beta}}$
- $\sigma =$ variance of the distributed priors beliefs

TABLE IV

RESULTS FROM MAXIMUM LIKELIHOOD ESTIMATION[a]

Treatment	$\hat{\alpha}$	$\hat{\beta}$	$\hat{\mu}$	$\hat{Q}$	$\hat{\varepsilon}$	$\hat{\delta}$	$-\ln L$
unconstrained	42	2.75	.939	.956	.18	.045	327.35
$\mu = Q$	44	2.75	.941	.941	.18	.045	327.41
$\delta = 0$	68	2.50	.965	.950	.21	.000	345.08
$\sigma = 0$	—	—	.972	.850	.23	.020	371.04

TABLE V

LIKELIHOOD RATIO TESTS

	Hypothesis	Treatment	d.f.	-2 log likelihood ratio
Theoretical Hypotheses	Heterogeneity ($\sigma = 0$)	4-move	1	87.38*
		6-move	1	181.78*
	Rational expectations ($\mu = Q$)	4-move	1	.12
		6-move	1	1.38
	Learning ($\delta = 0$)	4-move	1	35.46*
		6-move	1	37.88*

Quantal Response Equilibria for Extensive Form Games

Agent logit equilibrium

Agent quantal response equilibrium (AQRE)

Generalization of QRE to extensive form games (McKelvey and Palfrey, 1998).

- including games with incomplete information

Agent quantal response equilibrium (AQRE)

Generalization of QRE to extensive form games (McKelvey and Palfrey, 1998).

- including games with incomplete information

Basic idea

- I_i denotes an information set of player i, σ denotes a completely mixed strategy profile, $\pi_i(s_i|\sigma)$ denotes the expected payoff of action $s_i \in S_i(I_i)$
- Assume the DM's perceived payoff is $\hat{\pi}_i(s_i|\sigma) = \pi_i(s_i|\sigma) + \epsilon_{s_i}$ where ϵ_{s_i} is a random variable, privately observed by the DM, and independently distributed among agents.
- A strategy is σ^* is an *agent quantal response equilibrium* if for all i, and all s_i

$$\sigma^*(s_i) = \mathbb{P}\left(\hat{\pi}_i(s_i|\sigma^*) \geq \hat{\pi}_i(s_i'|\sigma^*), \forall s_i' \in S_i(I_i)\right) =$$

$$\mathbb{P}\left(\pi_i(s_i|\sigma^*) - \pi_i(s_i'|\sigma^*) \geq \epsilon_{s_i'} - \epsilon_{s_i}, \forall s_i' \in S_i(I_i)\right)$$

Logit AQRE

- σ is an Agent logit equilibrium if ϵ_{s_i} is a Weibull Distribution ("extreme value Type I") with cumulative distribution function $F(\epsilon_{s_i}) = e^{-e^{\lambda \epsilon_{s_i}}}$
- This implies (see eg., McFadden, 1975) that for all actions in all information sets

$$\sigma(s_i) = \frac{\exp\{\lambda \cdot \pi_i(s_i|\sigma)\}}{\sum_{s_i' \in S_i(I_i)} \exp\{\lambda \cdot \pi_i(s_i'|\sigma)\}}$$

Logit AQRE

- σ is an Agent logit equilibrium if ϵ_{s_i} is a Weibull Distribution ("extreme value Type I") with cumulative distribution function $F(\epsilon_{s_i}) = e^{-e^{\lambda \epsilon_{s_i}}}$
- This implies (see eg., McFadden, 1975) that for all actions in all information sets

$$\sigma(s_i) = \frac{\exp\{\lambda \cdot \pi_i(s_i|\sigma)\}}{\sum_{s_i' \in S_i(I_i)} \exp\{\lambda \cdot \pi_i(s_i'|\sigma)\}}$$

Properties

- For $\lambda < \infty$, all actions are chosen with positive probability, thus there are no information sets off the path, equilibrium refinement is implied
- As λ approaches infinity, convergence to a sequential equilibrium

Can be combined with social preferences, too, e.g. in ultimatum games

Reinterpreting the centipede game date with AQRE

When $S_i(I_i) = \{s_i, s_{i'}\}$ is

$$\sigma(s_i) = \frac{1}{1 + exp(\lambda(\pi_i(s_i'|\sigma) - \pi_i(s_i|\sigma))}$$

So,

$$p_4 = \frac{1}{1 + exp(\lambda(1.6 - 3.2))}$$

$$p_3 = \frac{1}{1 + exp(\lambda(0.8p_4 + 6.4(1 - p_4) - 1.6)))}$$

$$p_2 = \frac{1}{1 + exp(\lambda(0.4p_3 + (1 - p_3)[3.2p_4 + 1.6(1 - p_4)) - 0.8))}$$

$$p_1 = \frac{1}{1 + exp(\lambda(0.2p_2 + (1 - p_2)(1.6p_3 + (1 - p_3)(0.8p_4 + 6.4(1 - p_4))$$

An Experimental Study of Constant-sum Centipede Games

Agent logit equilibrium

Motivation:

In the previous paper there were two very different ideas:

1. With the existence of altruistic subjects, there is an incentive for selfish players to choose PASS,
2. The subjects make errors throughout the game

To disentangle this idea they study a constant-sum centipede game in which there are no Pareto improvements over the course of the game, so that no selfish player has an interest to choose PASS.

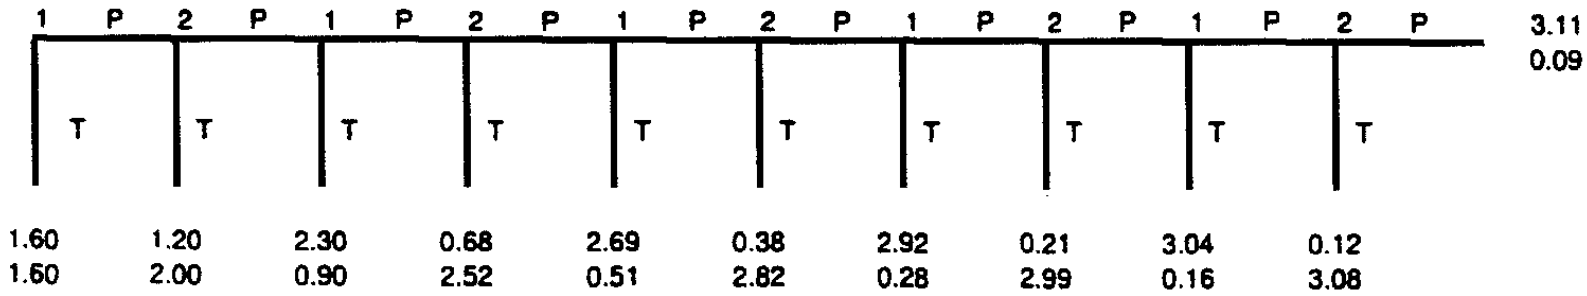


Fig. 1. A ten-move constant-sum centipede game

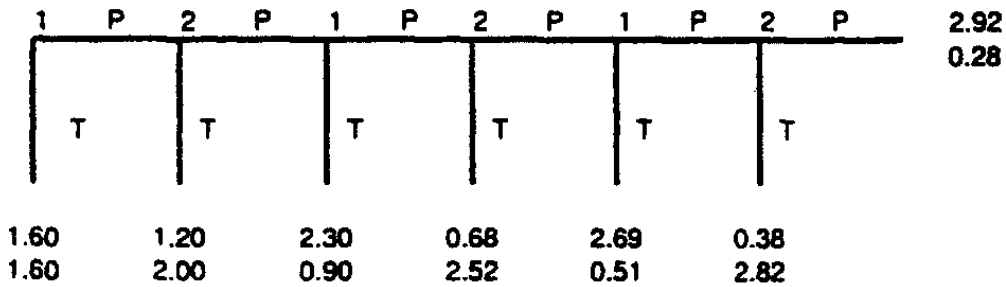


Fig. 2. A six-move constant-sum centipede game

Table 2. Proportions of matches ending at each outcome

Exp. #	Number of Passes* 0	1	2	3	4	5	6	7	8	9	10
1	.57	.31	.11	.00	.01	0	0	0	0	0	0
CIT-10	(57)	(31)	(11)	(0)	(1)	0	0	0	0	0	0
2	.51	.36	.12	.00	.01	0	0	0	0	0	0
CIT-10	(51)	(36)	(12)	(0)	(1)	0	0	0	0	0	0
3	.60	.28	.12	0	0	0	0	0	0	0	0
UI-10	(60)	(28)	(12)	0	0	0	0	0	0	0	0
4	.38	.36	.18	.06	.01	.00	.01	0	0	0	0
UI-10	(38)	(36)	(18)	(6)	(1)	(0)	(1)	0	0	0	0
5	.42	.40	.09	.05	.02	.02	0	0	0	0	0
PCC-10	(34)	(32)	(7)	(4)	(2)	(2)	0	0	0	0	0
6	.22	.23	.26	.13	.08	.06	.01	.00	.00	.01	0
PCC-10	(22)	(23)	(26)	(13)	(8)	(6)	(1)	(0)	(0)	(1)	0
Pooled	.45	.32	.14	.04	.02	.01	.003	.00	.00	.001	0
10 move	(262)	(186)	(86)	(23)	(13)	(8)	(2)	(0)	(0)	(1)	0

7	.62	.31	.07	0	0	0	0
CIT-6	(62)	(31)	(7)	0	0	0	0
8	.77	.23	0	0	0	0	0
UI-6	(77)	(23)	0	0	0	0	0
9	.33	.48	.15	.02	.01	0	0
PCC-6	(27)	(39)	(12)	(2)	(1)	0	0
Pooled	.59	.33	.07	.007	.003	0	0
6 move	(166)	(93)	(19)	(2)	(1)	0	0

* The number in parentheses is the number of observations at that node in the game tree.

Table 3. Implied take probabilities across matches

Exp. #	Match Group	Conditional Probabilities* p_0	p_1	p_2	p_3	p_4	p_5	p_6	p_7	p_8	p_9
1	<6	.46	.70	.88	.00	1.00					
		(50)	(27)	(8)	(1)	(1)					
CIT-10	>5	.68	.75	1.00							
		(50)	(16)	(4)							
2	<6	.34	.64	.92	.00	1.00					
		(50)	(33)	(12)	(1)	(1)					
CIT-10	>5	.68	.94	1.00							
		(50)	(16)	(1)							
3	<6	.48	.65	1.00							
		(50)	(26)	(9)							
UI-10	>5	.72	.79	1.00							
		(50)	(14)	(3)							
4	<6	.32	.38	.67	.71	.50	.00	1.00			
		(50)	(34)	(21)	(7)	(2)	(1)	(1)			
UI-10	>5	.44	.82	.80	1.00						
		(50)	(28)	(5)	(1)						
5	<6	.40	.67	.56	.50	.50	1.00				
		(45)	(27)	(9)	(4)	(2)	(1)				
PCC-10	>5	.44	.70	.33	.50	.50	1.00				
		(36)	(20)	(6)	(4)	(2)	(1)				
6	<6	.22	.15	.36	.42	.50	.67	.50	.00	.00	1.00
		(50)	(39)	(33)	(21)	(12)	(6)	(2)	(1)	(1)	(1)
PCC-10	>5	.22	.44	.64	.50	.50	1.00				
		(50)	(39)	(22)	(8)	(4)	(2)				

7	<6	.48	.73	1.00		
		(50)	(26)	(7)		
CIT-6	>5	.76	1.00			
		(50)	(12)			
8	<6	.70	1.00			
		(50)	(15)			
UI-6	>5	.84	1.00			
		(50)	(8)			
9	<6	.20	.67	.75	.67	1.00
		(45)	(36)	(12)	(3)	(1)
PCC-6	>5	.50	.83	1.00		
		(36)	(18)	(3)		

* The number in parentheses is the number of observations at that node in the game tree.

The always Take Model

Backward induction implies "always Take"

This is a point prediction, in that the model rules out any other outcome as impossible.

Econometrically, this makes it impossible to specify the likelihood of our data, which contain outcomes contradicting "always Take".

Fey et al. consider two models of mistakes: an AQRE model and a trembling-hand model

In the trembling-hand version of "always Take":

- Every player, at every node reached in every match, chooses Take with a fixed probability p and chooses Pass with probability $1 - p$.
- Estimates of p close to 1 would indicate that the "Always Take" model does fairly well in explaining the data.

The always Take Model

- The DM chooses "Take" at any given node with probability

$$p_t = 1 - (1 - p_0)e^{-\alpha t}$$

  with p_0 is the probability of Take used by subjects in the first match of the experiment and α represents the learning rate of the subjects.
- M be the total number of matches in the experiment, indexed by $i = 1....., M$.
- n_i is the terminal node in match i, t_i is the round in which match i is played
  The log-likelihood is given by

$$\log\ L = \sum_{i=1}^{M} \log\left[1 - (1 - p_0)e^{-\alpha t_i}\right] + \sum_{i=1}^{M} n_i \log(1 - p_0)e^{-\alpha t_i}$$

The AQRE with learning

- A player at node j in match number t faces a choice between Take, with payoff $\hat{u}_T = u_T + \epsilon_T$, and Pass, with expected payoff $\hat{u}_P + \epsilon_P$ where $\epsilon = \epsilon_T - \epsilon_P$ is drawn from a logistic distribution. So, the probability that a player at node j in the match t chooses *Take* is

$$\begin{aligned} p_t^j &= \mathbb{P}(\hat{u}_T > \hat{u}_P) \\ &= \mathbb{P}(\epsilon > u_P - u_T) = \frac{1}{1 + e^{\lambda(u_P - u_T)}} \end{aligned}$$

- p_t^N is the probability that Player 2 chooses *Take* at N. Given this decision probability, we can compute the expected payoff a Pass by a Player 1 at the node $N-1$ and determine: p_t^{N-1}.
- By backward induction, we thus obtain

$$\mathbf{p}_t = (p_t^1, p_t^2, ..., p_t^N)$$

The AQRE with learning

- To capture learning, we assume that

$$\lambda_t = \lambda_0 + \beta t$$

  which implies that decision probabilities depend on round t of a match

- Let $\hat{f}_t^1 = p_t^1$, $\hat{f}_t^2 = (1 - p_t^1)p_t^2$, $\hat{f}_t^3 = (1 - p_t^1)(1 - p_t^2)p_t^3$, etc., capture the predicted probabilities that the match in round t terminates at the nodes 1, 2, ...

- The log likelihood is

$$\log\ L = \sum_{i=1}^{M} \log(\hat{f}_{t_i}^{n_i}(\lambda_{t_i}, \beta))$$

Results of AQRE vs "Always Take"

- In eight of the nine individual experiments, the Agent Quantal Response Equilibrium model does better
- AQRE captures behavior significantly better in the 6-move games
- AQRE fits about as good as "always Take" in the 10-move games
- suggests that AQRE works in short games, but 10 moves seem too long to capture behavior reliably with a backward induction concept

ML estimates for the "always Take" model

Table 5. MLE of learning model parameters

Exp. #	Subject Pool	MLE p_0	α	$-\log L$
1	CIT-10	0.526	0.0683	101.19
2	CIT-10	0.416	0.117**	104.13
3	UI-10	0.496	0.1067**	94.15
4	UI-10	0.380	0.0568*	136.04
5	PCC-10	0.554	−0.0198	108.54
6	PCC-10	0.254	0.0335*	184.72
Pooled	10 move	0.410	0.0531**	764.78
7	CIT-6	0.511	0.1292**	85.61
8	UI-6	0.702	0.1225	57.62
9	PCC-6	0.367	0.0917**	102.84
Pooled	6 move	0.493	0.1212**	257.45

* Significant at the .05 level.
** Significant at the .01 level.

ML estimates for the AQRE model

Table 6. MLE for quantal response equilibrium model (logistic specification)

Exp. #	Subject Pool	MLE λ_0	β	$-\log L$	Vuong Stat.
1	CIT-10	0.745	0.128	98.62	−0.909
2	CIT-10	0.298	0.314**	97.18	−2.679**
3	UI-10	0.684	0.232**	90.56	−1.551
4	UI-10	0.032	0.142**	133.30	−1.169
5	PCC-10	0.769	−0.095	107.09	−1.021
6	PCC-10	−0.479	0.094	198.49	3.173**
Pooled	10 move	−0.028	0.103**	764.08	−0.157
7	CIT-6	0.827	0.324**	79.48	−2.949**
8	UI-6	2.184	0.343*	52.42	−3.415**
9	PCC-6	0.376	0.201*	96.96	−2.596**
Pooled	6 move	0.683	0.292**	243.57	−3.314**

* Significant at the .05 level.

** Significant at the .01 level.

Information cascades

Agent logit equilibrium

Informational cascades

"An informational cascade occurs when it is optimal for an individual, having observed the actions of those ahead of him, to follow the behavior of the preceding individual without regard to his own information."

Bikhchandani, Hirschleifer, and Welch, 1992

Information cascades capture the idea that a group of rational decision makers can generate collective behaviors that seem, *prima facie*, irrational (e.g. financial bubbles, spread of false information, fashion cycles and fades...).

Simple game with information Cascades

Basic setup:

- Urn A = (2 red balls and 1 blue ball)
- Urn B = (1 red ball and 2 blue balls)
- Nature P_0 and six players $(P_1, ..., P_6)$
- P_0 chooses A with probability $\frac{1}{2}$ and B with probability $\frac{1}{2}$ but does not reveal her choice
- Each of the six players receives a payoff of 1 if they choose the urn chosen by Player 0.

Order of moves:

1. P_1 draws a ball from the urn chosen by P0 and chooses an urn;
2. P_2 observes P_1's choice, draws a ball from the urn chosen by P_0, and chooses an urn.
3. P_3 observes P_1 and P_2's choices, draws a ball from the urn chosen by P_0, and chooses an urn

...

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \quad B = (\bullet, \bullet, \bullet)$$

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \quad B = (\bullet, \bullet, \bullet)$$

- Three players (P_1, P_2, P_3)

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \quad B = (\bullet, \bullet, \bullet)$$

- Three players (P_1, P_2, P_3)
- P_1 draws a ball from the chosen urn and selects an urn;

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \quad B = (\bullet, \bullet, \bullet)$$

- Three players (P_1, P_2, P_3)

- P_1 draws a ball from the chosen urn and selects an urn;
- P_2 observes P_1's choice, draws a ball, and selects an urn;

Incomplete Information: Cascade games

- One the following Urns is selected with probability $\frac{1}{2}$:

$$A = (\bullet, \bullet, \bullet) \quad B = (\bullet, \bullet, \bullet)$$

- Three players (P_1, P_2, P_3)

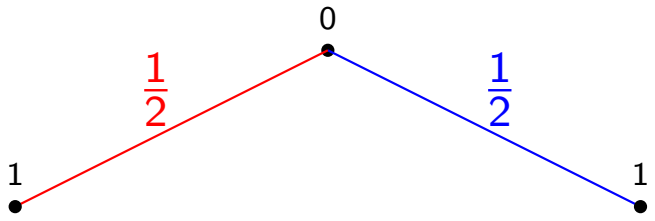
- P_1 draws a ball from the chosen urn and selects an urn;
- P_2 observes P_1's choice, draws a ball, and selects an urn;
- P_3 observes previous choices, draws a ball, and selects an urn;

Each player get something only if he correclty guesses which urn has been selected

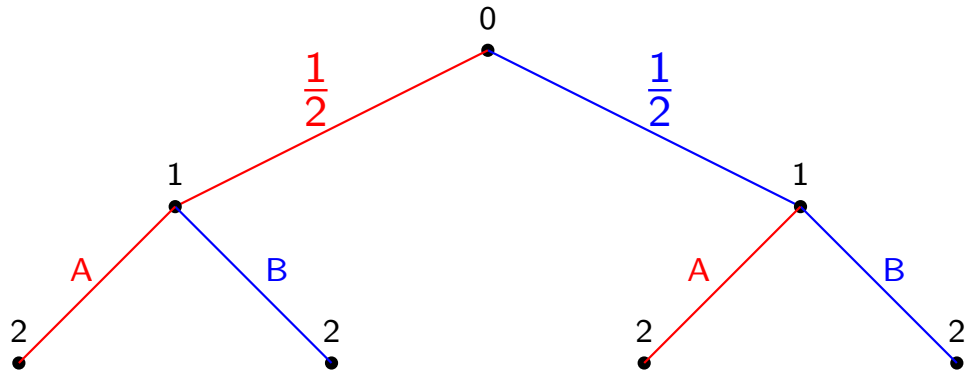
The Extensive Form Representation with three players (without showing ball-drawing)



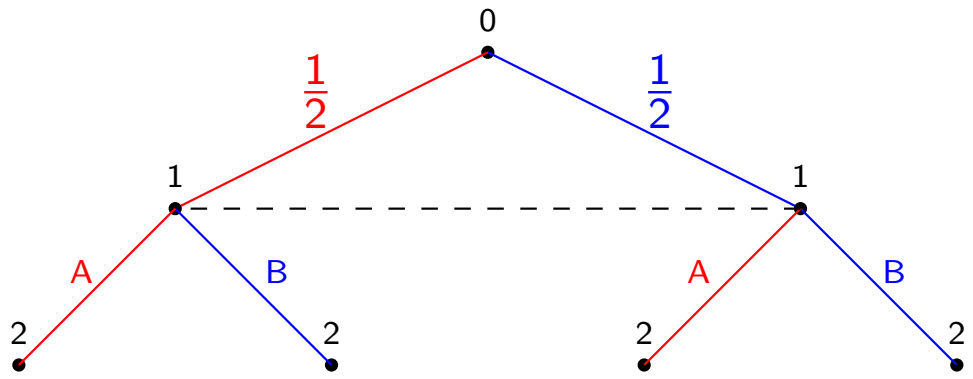
The Extensive Form Representation with three players (without showing ball-drawing)



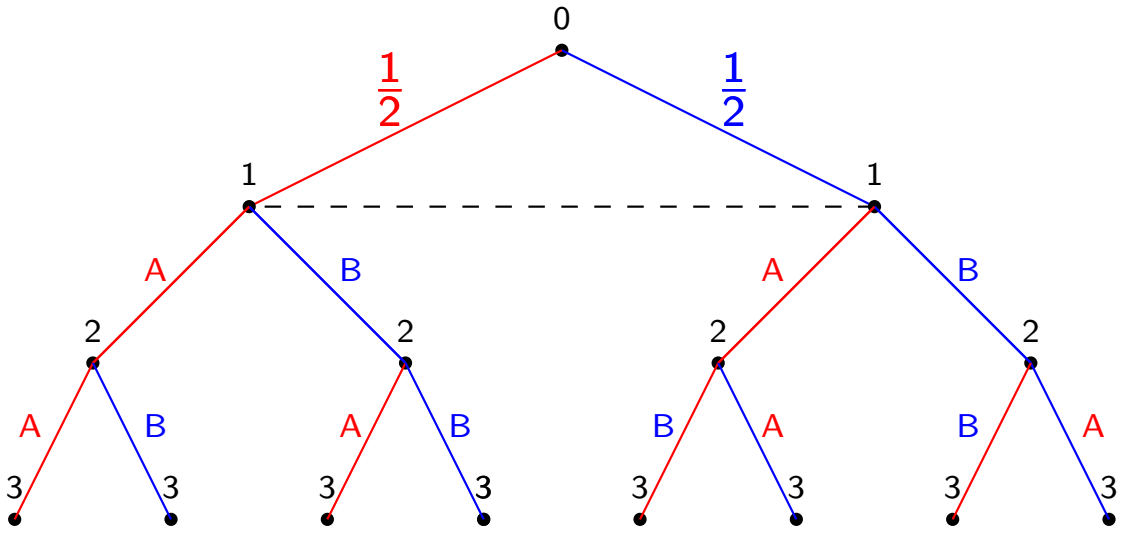
The Extensive Form Representation with three players (without showing ball-drawing)



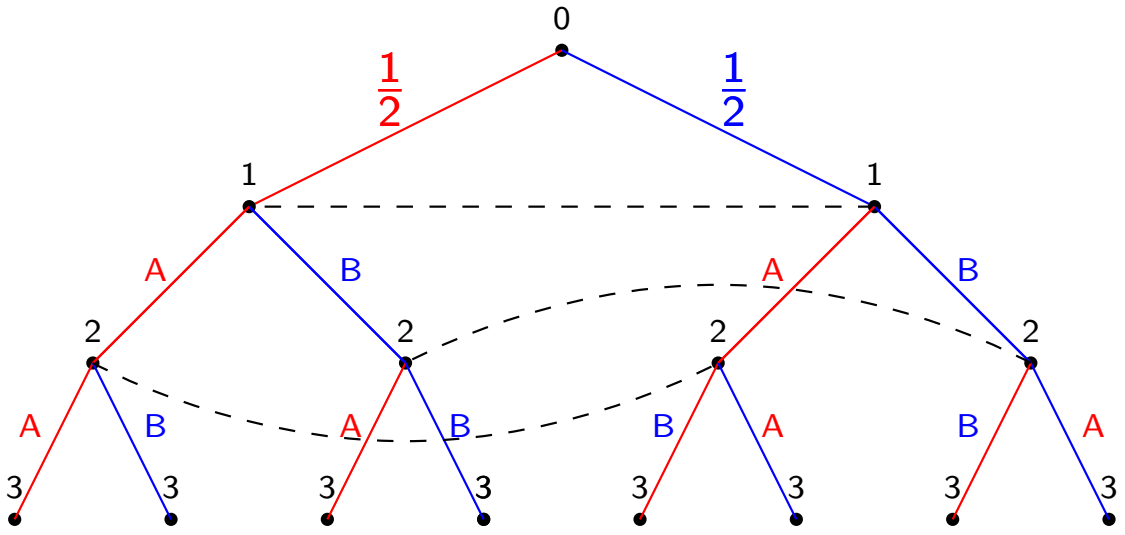
The Extensive Form Representation with three players (without showing ball-drawing)



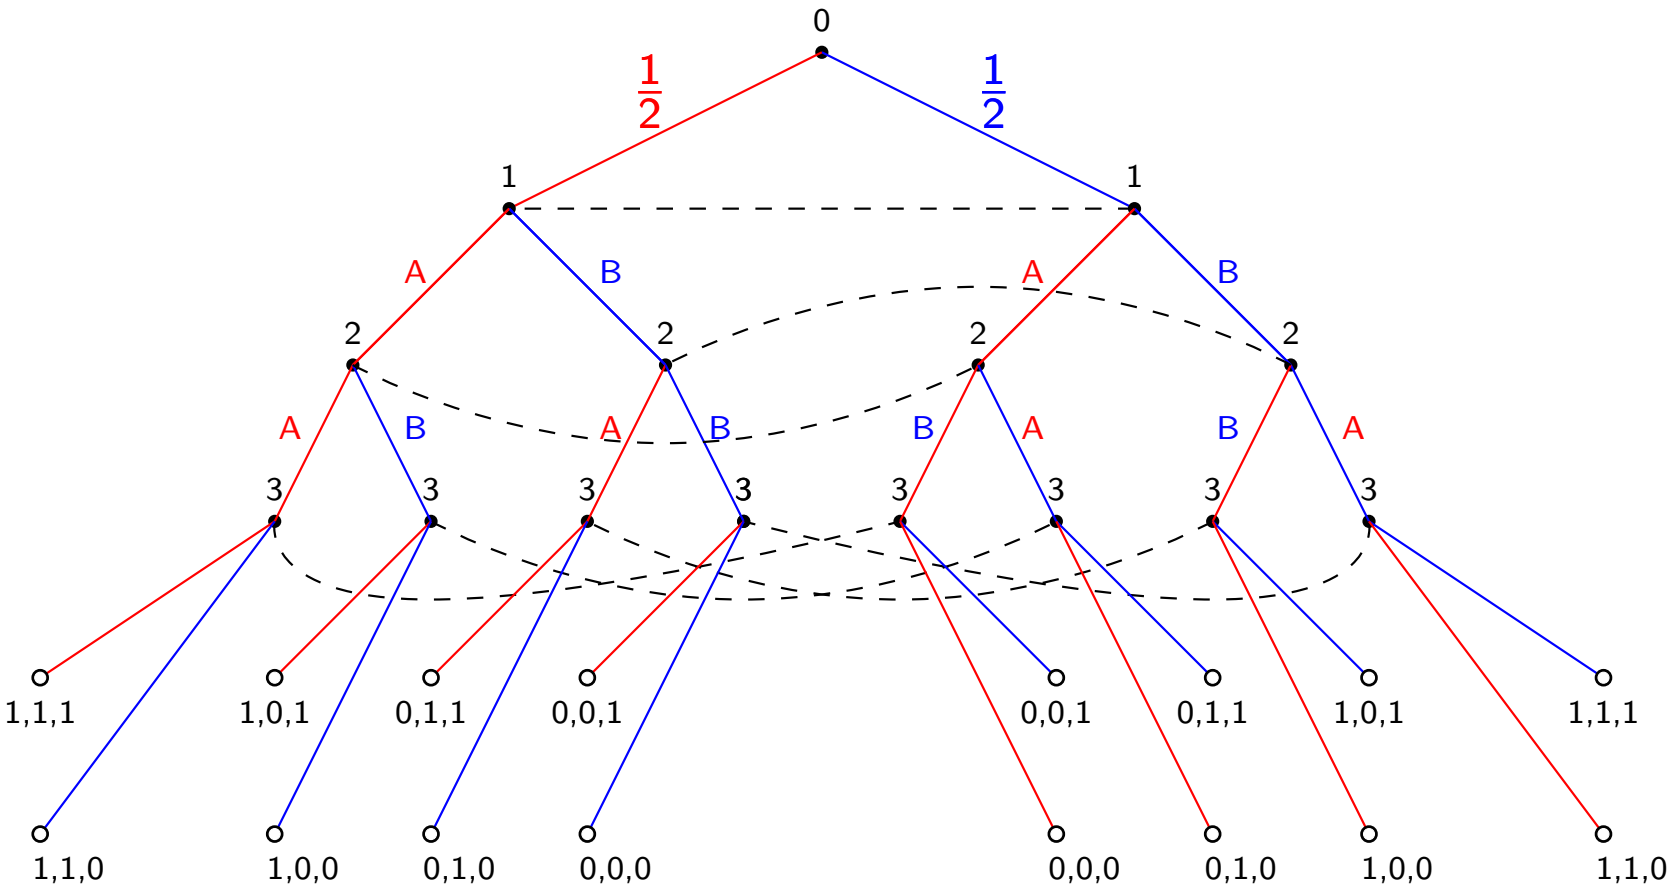
The Extensive Form Representation with three players (without showing ball-drawing)



The Extensive Form Representation with three players (without showing ball-drawing)



The Extensive Form Representation with three players (without showing ball-drawing)



Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

2. If P_2 also draws a red ball (a), then she chooses A for sure; if she draws a blue ball (b) then she is indifferent and chooses A with probability $p \in [0, 1]$.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

2. If P_2 also draws a red ball (a), then she chooses A for sure; if she draws a blue ball (b) then she is indifferent and chooses A with probability $p \in [0, 1]$.

   P_2 chose $A \Rightarrow$ probability that she drew a red ball (a) is $1/(1+p)$.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

2. If P_2 also draws a red ball (a), then she chooses A for sure; if she draws a blue ball (b) then she is indifferent and chooses A with probability $p \in [0, 1]$.

   P_2 chose $A \Rightarrow$ probability that she drew a red ball (a) is $1/(1+p)$.

3. P_1 and P_2 chose $A \Rightarrow$ both players drew a red ball with probability $1/(1+p)$
   P_1 drew a red ball (a) and P_2 a blue ball (b) with probability $p/(1+p)$.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

2. If P_2 also draws a red ball (a), then she chooses A for sure; if she draws a blue ball (b) then she is indifferent and chooses A with probability $p \in [0, 1]$.

   P_2 chose $A \Rightarrow$ probability that she drew a red ball (a) is $1/(1+p)$.

3. P_1 and P_2 chose $A \Rightarrow$ both players drew a red ball with probability $1/(1+p)$
   P_1 drew a red ball (a) and P_2 a blue ball (b) with probability $p/(1+p)$.

4. P_3 gets signal b and in total has $1 + 1/(1+p)$ signals indicating the urn is A and $1 + p/(1+p)$ signals indicating that the right urn is B.

Intuition for the information cascade

If P_1 and P_2 choose A, then P_3 chooses A, regardless her signal.

1. P_1 chooses A only if she draws a red ball $\Rightarrow$ P_1 received a red ball for sure.

2. If P_2 also draws a red ball (a), then she chooses A for sure; if she draws a blue ball (b) then she is indifferent and chooses A with probability $p \in [0, 1]$.

   P_2 chose $A \Rightarrow$ probability that she drew a red ball (a) is $1/(1+p)$.

3. P_1 and P_2 chose $A \Rightarrow$ both players drew a red ball with probability $1/(1+p)$
   P_1 drew a red ball (a) and P_2 a blue ball (b) with probability $p/(1+p)$.

4. P_3 gets signal b and in total has $1 + 1/(1+p)$ signals indicating the urn is A and $1 + p/(1+p)$ signals indicating that the right urn is B.

5. Thus, if $p < 1$, she should choose A.

Implication: All signals but the first two are ignored – highly inefficient

Does Bayesian Nash equilibrium predict subjects' behavior in Lab?

Anderson and Holt (1997) tested the predicted information cascade experimentally with 6 players.

They conclude that

- "Individuals generally used information efficiently and followed the decisions of others when it was rational."
- "However, some errors, which tended to make subjects rely more on their own private information."

Questions

- Can we explain the deviations from Nash equilibrium with Quantal Response Equilibrium (AQRE)?
- Or should we explain the data with violation of Rational Expectations? I.e. that the subjects do not weigh the sources of information symmetrically
- Do subjects acknowledge their opponents' rationality; or do they systematically overweigh their own information relative to the information they could infer from others' behavior?

Conclusion

Agent logit equilibrium

Conclusion

- QRE can be extended straightforwardly to allow for sequential choice, maintaining the standard assumptions of logit choice and rational expectations
- Can capture behavior in some cases reasonably well, but not in general
- One recurrent theme is as in normal-form games: Subjects seem to underestimate the rationality of others
  - Explains deviations from rational-expectation play in cascade formation
  - May explain difficulties of AQRE of capturing behavior in long centipede games
- Subjects seem to even stronger underestimate how others estimate rationality of others
  - Explains why subjects seem to look ahead rather than induce backwards

References I

Repetition

Agent logit equilibrium

Repetition questions

Repetition questions

Repetition questions

Repetition questions

Repetition questions

- What is the centipede game?

Repetition questions

- What is the centipede game?
  - Alternating moves, two actions per stage (take, pass), players have incentive to pass iff opponent passes, too Jump

Repetition questions

- What is the centipede game?
  - Alternating moves, two actions per stage (take, pass), players have incentive to pass iff opponent passes, too Jump
- What are the central experimental observations and potential explanations?

Repetition questions

- What is the centipede game?
  - Alternating moves, two actions per stage (take, pass), players have incentive to pass iff opponent passes, too Jump
- What are the central experimental observations and potential explanations?
  - potential explanations: noise, strategic uncertainty, altruism
  - very few subjects stop in round 1 $\rightsquigarrow$ backward induction or max-pay not satisfied Jump
  - probability of stopping is increasing as game is played $\rightsquigarrow$ not simply mistakes
  - subjects stop later in first five games than in last five: learning, less altruism? Jump

Repetition questions II

- What is the agent quantal response equilibrium?

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?
  - Same as a standard centipede game, but sum of payoffs is constant
  - Interesting, as constant sum rules out altruism as driving force Jump

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?
  - Same as a standard centipede game, but sum of payoffs is constant
  - Interesting, as constant sum rules out altruism as driving force Jump
- What are the central observations?

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?
  - Same as a standard centipede game, but sum of payoffs is constant
  - Interesting, as constant sum rules out altruism as driving force Jump
- What are the central observations?
  - Subjects tend to take a little earlier than in standard centipede games, but still rather late
  - Probability of pass again not constant, so not simply mistakes, and some learning Jump

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?
  - Same as a standard centipede game, but sum of payoffs is constant
  - Interesting, as constant sum rules out altruism as driving force Jump
- What are the central observations?
  - Subjects tend to take a little earlier than in standard centipede games, but still rather late
  - Probability of pass again not constant, so not simply mistakes, and some learning Jump
- Fey et al. evaluate AQRE compared to the always-take model. What is the latter and what do they find?

Repetition questions II

- What is the agent quantal response equilibrium?
  - in general: players choose logit responses in all information sets (correctly anticipating strategies of others) Jump
  - in games with complete information: players play logit equilibria in all subgames Jump
- What is the constant-sum centipede game and why is it interesting?
  - Same as a standard centipede game, but sum of payoffs is constant
  - Interesting, as constant sum rules out altruism as driving force Jump
- What are the central observations?
  - Subjects tend to take a little earlier than in standard centipede games, but still rather late
  - Probability of pass again not constant, so not simply mistakes, and some learning Jump
- Fey et al. evaluate AQRE compared to the always-take model. What is the latter and what do they find?
  - always-take: DM chooses "Take" at any given node with probability $p_t = 1 - (1 - p_0)e^{-\alpha t}$; interesting benchmark, allows for learning Jump
  - AQRE captures behavior significantly better in the 6-move games, and about as good as "always Take" in the 10-move games Jump
  - AQRE/backward induction does not seem to capture behavior in long games