Epistemic Game Theory

Lecture 8

Eric Pacuit

University of Maryland, College Park
pacuit.org
epacuit@umd.edu

April 1, 2014

Deliberation in Games

B. Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, 1990.

Deliberation in Games

"The economist's predilection for equilibria frequently arises from the belief that some underlying dynamic process (often suppressed in formal models) moves a system to a point from which it moves no further."

(pg. 1008)

B. D. Bernheim. *Rationalizable Strategic Behavior*. Econometrica, 52, 4, pgs. 1007 - 1028.

"It is not just a question of what common knowledge obtains at the moment of truth, but also how common knowledge is preserved, created, or destroyed in the deliberational process which leads up to the moment of truth." (pg. 160)

B. Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, 1990.

Information Feedback

In the simplest case, deliberation is trivial; one calculates expected utility and maximizes

Information feedback: "the very process of deliberation may generate information that is relevant to the evaluation of the expected utilities. Then, processing costs permitting, a Bayesian deliberator will feed back that information, modifying his probabilities of states of the world, and recalculate expected utilities in light of the new knowledge."

Skyrms' Model of Deliberation

A decision maker has to choose between n acts: s_1, s_2, ..., s_n

Skyrms' Model of Deliberation

A decision maker has to choose between n acts: s_1, s_2, ..., s_n

State of indecision: $\mathbf{P} = \langle p_1, \ldots, p_n \rangle$ of probabilities for each act ($\sum_i p_i = 1$). The *default mixed act* is the mixed act corresponding to the state of indecision (decision makers always make a decision).

Skyrms' Model of Deliberation

A decision maker has to choose between n acts: s_1, s_2, …, s_n

State of indecision: $\mathbf{P} = \langle p_1, \ldots, p_n \rangle$ of probabilities for each act ($\sum_i p_i = 1$). The *default mixed act* is the mixed act corresponding to the state of indecision (decision makers always make a decision).

Status quo: $EU(\mathbf{P}) = \sum_i p_i \cdot u_i(s_i)$

The Dynamical Rule "Seeks the good"

As a rational player deliberates, she updates here state of indecision according to some dynamical rule that "seeks the good":

1. the dynamical rule raises the probability of an act only if that act has utility greater than the status quo
2. the dynamical rule raises the sum of the probability of all acts with utility greater than the status quo (if any)

The Dynamical Rule "Seeks the good"

As a rational player deliberates, she updates here state of indecision according to some dynamical rule that "seeks the good":

1. the dynamical rule raises the probability of an act only if that act has utility greater than the status quo
2. the dynamical rule raises the sum of the probability of all acts with utility greater than the status quo (if any)

Fact. All dynamical rules that seek the good have the same fixed points: those states in which the expected utility of the status quo is maximal.

Games played by Bayesian deliberators

For each player, the decisions of the other players constitute the relevant state of the world, which together with her decision, determines the consequences in accordance with the payoff matrix.

Games played by Bayesian deliberators

For each player, the decisions of the other players constitute the relevant state of the world, which together with her decision, determines the consequences in accordance with the payoff matrix.

1. Start from the initial position, player i calculates expected utility and moves by her adaptive rule to a new state of indecision.

Games played by Bayesian deliberators

For each player, the decisions of the other players constitute the relevant state of the world, which together with her decision, determines the consequences in accordance with the payoff matrix.

1. Start from the initial position, player i calculates expected utility and moves by her adaptive rule to a new state of indecision.
2. She knows that the other players are Bayesian deliberators who have just carried out a similar process.

Games played by Bayesian deliberators

For each player, the decisions of the other players constitute the relevant state of the world, which together with her decision, determines the consequences in accordance with the payoff matrix.

1. Start from the initial position, player *i* calculates expected utility and moves by her adaptive rule to a new state of indecision.
2. She knows that the other players are Bayesian deliberators who have just carried out a similar process.
3. So, she can simply go through their calculations to see their new states of indecision and update her probabilities for their acts accordingly (*update by emulation*).

Let G be a strategic game for two players with n strategies and $\langle r_{ij}, c_{ij} \rangle$ be the payoff matrix for G.

$\mathbf{P}_{col}(t)$, $\mathbf{P}_{row}(t)$ are row and columns states of indecision at stage t of the deliberational process.

For example, a state of indecision for the row player is

$$\mathbf{P}_{row}(t) = \langle p^1_{row}(t), \ldots, p^n_{row}(t) \rangle$$

where $p^j_{row}(t)$ is the probability that row assigns to her strategy j at time t.

$$EU_{row}(i, t) = \sum_{k=1}^{n} p_{col}^{k}(t) \cdot r_{ik}$$

$$EU_{row}(i,t) = \sum_{k=1}^{n} p_{col}^{k}(t) \cdot r_{ik}$$

$$SQ_{row}(t) = \sum_{i=1}^{n} p_{row}^{i}(t) \cdot EU_{row}(i,t)$$

$$EU_{row}(i,t) = \sum_{k=1}^{n} p_{col}^{k}(t) \cdot r_{ik}$$

$$SQ_{row}(t) = \sum_{i=1}^{n} p_{row}^{i}(t) \cdot EU_{row}(i,t)$$

$$Cov_{row}(i,t) = \max\{EU_{row}(i,t) - SQ_{row}(t), 0\}$$

$$EU_{col}(i,t) = \sum_{k=1}^{n} p^k_{row}(t) \cdot c_{ki}$$

$$SQ_{col}(t) = \sum_{i=1}^{n} p^i_{col}(t) \cdot EU_{col}(i,t)$$

$$Cov_{col}(i,t) = \max\{EU_{col}(i,t) - SQ_{col}(t), 0\}$$

$$\mathbf{P}_{row}(t+1) = D(\mathbf{P}_{row}(t), \mathbf{P}_{col}(t))$$

$$\mathbf{P}_{row}(t+1) = D(\mathbf{P}_{row}(t), \mathbf{P}_{col}(t))$$

Dynamical rule

$$\mathbf{P}_{row}(t+1) = D(\mathbf{P}_{row}(t), \mathbf{P}_{col}(t))$$

Dynamical rule

State of indecision

$$\mathbf{P}_{row}(t+1) = D(\mathbf{P}_{row}(t), \mathbf{P}_{col}(t))$$

Dynamical rule State of indecision Beliefs about the state of nature

$$\mathbf{P}_{col}(t+1) = D(\mathbf{P}_{col}(t), \mathbf{P}_{row}(t))$$

Dynamical rule State of indecision Beliefs about the state of nature

Dynamical Rules

Nash: $p^i_{row}(t+1) = \frac{k \cdot p^i_{row}(t) + Cov_{row}(i,t)}{k + \sum_i Cov_{row}(i,t)}$

Bayes: $p^i_{row}(t+1) = p^i_{row}(t) + \frac{1}{k} \cdot p^i_{row}(t) \cdot \frac{EU_{row}(i,t) - SQ_{row}(t)}{SQ_{row}(t)}$

Bayes2: $p^i_{row}(t+1) = p^i_{row}(t) \cdot \frac{EU_{row}(i,t)}{SQ_{row}(t)}$

$k > 0$ is an **index of caution** (slowing down the rate of convergence)

		col L	col R
row	U	-10,-10	5,-5
row	D	-5,5	0,0

$$\mathbf{P}_{row}(0) = \langle 0.2, 0.8 \rangle \text{ and } \mathbf{P}_{col}(0) = \langle 0.4, 0.6 \rangle$$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(0) = \langle 0.2, 0.8\rangle \text{ and } \mathbf{P}_{col}(0) = \langle 0.4, 0.6\rangle$$

$EU_{row}(U, 0) = 0.4 \cdot -10 + 0.6 \cdot 5 = -1$
$EU_{row}(D, 0) = 0.4 \cdot -5 + 0.6 \cdot 0 = -2$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(0) = \langle 0.2, 0.8\rangle \text{ and } \mathbf{P}_{col}(0) = \langle 0.4, 0.6\rangle$$

$EU_{row}(U,0) = 0.4 \cdot -10 + 0.6 \cdot 5 = -1$
$EU_{row}(D,0) = 0.4 \cdot -5 + 0.6 \cdot 0 = -2$
$SQ_{row}(0) = 0.2 \cdot EU_{row}(U,0) + 0.8 \cdot EU_{row}(D,0) = 0.2 \cdot -1 + 0.8 \cdot -2 = -1.8$

		col L	col R
row	U	-10,-10	5,-5
row	D	-5,5	0,0

$$\mathbf{P}_{row}(0) = \langle 0.2, 0.8\rangle \text{ and } \mathbf{P}_{col}(0) = \langle 0.4, 0.6\rangle$$

$EU_{row}(U, 0) = 0.4 \cdot -10 + 0.6 \cdot 5 = -1$
$EU_{row}(D, 0) = 0.4 \cdot -5 + 0.6 \cdot 0 = -2$
$SQ_{row}(0) = 0.2 \cdot EU_{row}(U, 0) + 0.8 \cdot EU_{row}(D, 0) = 0.2 \cdot -1 + 0.8 \cdot -2 = -1.8$
$Cov_{row}(U, 0) = \max\{EU_{row}(U, 0) - SQ_{row}(0), 0\} = 0.8$
$Cov_{row}(D, 0) = \max\{EU_{row}(D, 0) - SQ_{row}(0), 0\} = 0$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(1) = \langle \frac{k \cdot 0.2 + 0.8}{k + 0.8}, \frac{k \cdot 0.8 + 0}{k + 0.8} \rangle \text{ and } \mathbf{P}_{col}(0) = \langle 0.4, 0.6 \rangle$$

$EU_{row}(U, 0) = 0.4 \cdot -10 + 0.6 \cdot 5 = -1$
$EU_{row}(D, 0) = 0.4 \cdot -5 + 0.6 \cdot 0 = -2$
$SQ_{row}(0) = 0.2 \cdot EU_{row}(U, 0) + 0.8 \cdot EU_{row}(D, 0) = 0.2 \cdot -1 + 0.8 \cdot -2 = -1.8$
$Cov_{row}(U, 0) = \max\{EU_{row}(U, 0) - SQ_{row}(0), 0\} = 0.8$
$Cov_{row}(D, 0) = \max\{EU_{row}(D, 0) - SQ_{row}(0), 0\} = 0$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(1) = \langle \frac{k\cdot 0.2+0.8}{k+0.8}, \frac{k\cdot 0.8+0}{k+0.8}\rangle \text{ and } \mathbf{P}_{col}(1) = \langle \frac{k\cdot 0.4+1.8}{k+1.8}, \frac{k\cdot 0.6+0}{k+1.8}\rangle$$

$EU_{col}(L,0) = 0.2\cdot -10 + 0.8\cdot 5 = 2$
$EU_{col}(R,0) = 0.2\cdot -5 + 0.8\cdot 0 = -1$
$SQ_{col}(0) = 0.4\cdot EU_{col}(L,0) + 0.6\cdot EU_{row}(R,0) = 0.4\cdot 2 + 0.6\cdot -1 = 0.2$
$Cov_{col}(L,0) = \max\{EU_{col}(L,0) - SQ_{col}(0), 0\} = 1.8$
$Cov_{col}(R,0) = \max\{EU_{col}(R,0) - SQ_{col}(0), 0\} = 0$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(1) = \langle 0.2 + \frac{1}{k}\cdot 0.2\cdot\frac{-1-(-1.8)}{-1.8}, 0.8 + \frac{1}{k}\cdot 0.8\cdot\frac{-2-(-1.8)}{-1.8}\rangle$$

$EU_{row}(U,0) = 0.4\cdot -10 + 0.6\cdot 5 = -1$
$EU_{row}(D,0) = 0.4\cdot -5 + 0.6\cdot 0 = -2$
$SQ_{row}(0) = 0.2\cdot EU_{row}(U,0) + 0.8\cdot EU_{row}(D,0) = 0.2\cdot -1 + 0.8\cdot -2 = -1.8$

		col	
		L	R
row	U	-10,-10	5,-5
	D	-5,5	0,0

$$\mathbf{P}_{row}(1) = \langle 0.2 + \frac{1}{k}\cdot 0.2 \cdot \frac{0.8}{-1.8}, 0.8 + \frac{1}{k}\cdot 0.8\cdot\frac{-0.2}{-1.8}\rangle$$

$EU_{row}(U,0) = 0.4\cdot -10 + 0.6\cdot 5 = -1$
$EU_{row}(D,0) = 0.4\cdot -5 + 0.6\cdot 0 = -2$
$SQ_{row}(0) = 0.2{\cdot}EU_{row}(U,0) + 0.8{\cdot}EU_{row}(D,0) = 0.2{\cdot}-1+0.8{\cdot}-2 = -1.8$

col

row		L	R
	U	0,0	15,5
	D	5,15	10,10

$$\mathbf{P}_{row}(1) = \langle 0.2 + \frac{1}{k}\cdot 0.2 \cdot \frac{9-8.2}{8.2}, 0.8 + \frac{1}{k}\cdot 0.8 \cdot \frac{8-8.2}{8.2}\rangle$$

$EU_{row}(U,0) = 0.4 \cdot 0 + 0.6 \cdot 15 = 9$
$EU_{row}(D,0) = 0.4 \cdot 5 + 0.6 \cdot 10 = 8$
$SQ_{row}(0) = 0.2 \cdot EU_{row}(U,0) + 0.8 \cdot EU_{row}(D,0) = 0.2 \cdot 9 + 0.8 \cdot 8 = 8.2$

		Bob	
		L	R
Ann	U	0.5,0.5	0.5,0.5
	D	1,1	0,0

		Bob	
		L	R
Ann	U	0.5,0.5	0.5,0.5
	D	1,1	0,0

Darwin can lead to an imperfect equilibrium. Nash can only lead to D, L.

		Bob	
		L	R
Ann	U	0.5,0.5	0.5,0.5
	D	1,1	0,0

Darwin can lead to an imperfect equilibrium. Nash can only lead to D, L.

Normal form vs. Extensive form

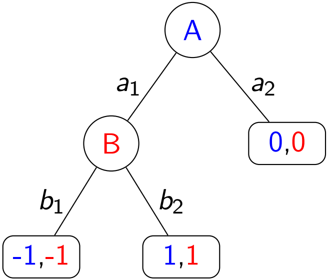

Normal form vs. Extensive form

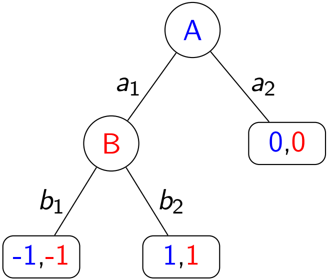

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

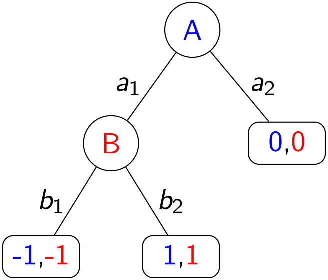

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

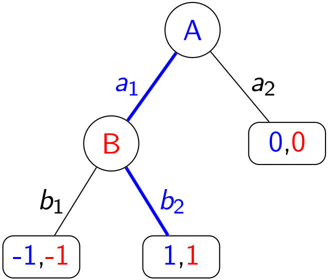

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

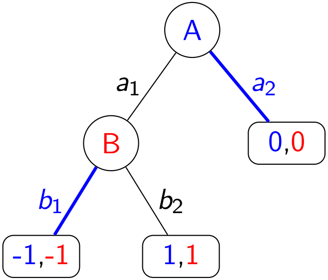

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

(Cf. the various notions of *sequential equilibrium*)

Normal form vs. Extensive form

T. Seidenfeld. *When normal and extensive form decisions differ*. in *Logic, Methodology and Philosophy of Science IX*, Elsevier, 1994.

Normal form vs. Extensive form

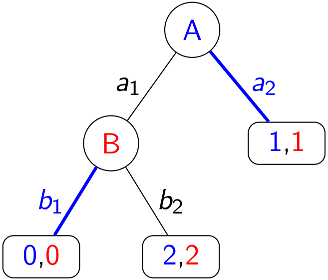

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

On the normal form, there are imperfect equilibria accessible by Darwin dynamics (e.g., $\mathbf{P}_A = \langle 0, 1 \rangle$, $\mathbf{P}_B = \langle 0.04, 0.96 \rangle$).

Normal form vs. Extensive form

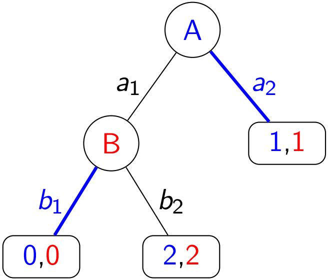

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

On the normal form, there are imperfect equilibria accessible by Darwin dynamics (e.g., $\mathbf{P}_A = \langle 0, 1\rangle$, $\mathbf{P}_B = \langle 0.97, 0.03\rangle$).

This equilibria is not accessible on the tree: Bob calculates the expected utility *at his information set* (so, $P_B(a_1 \mid a_1) = 1$ and $P_B(a_2 \mid a_1) = 0$).

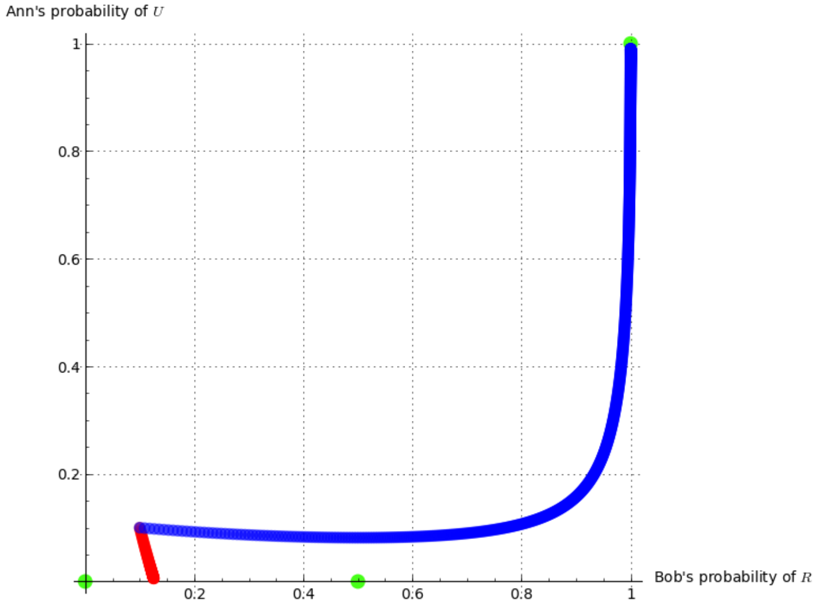
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

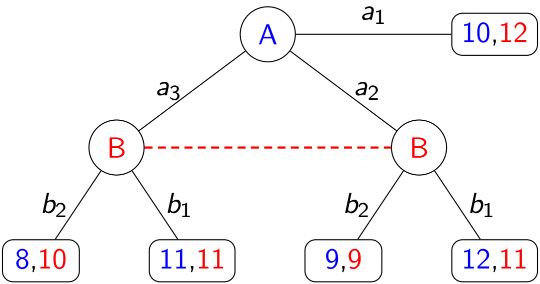

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	10,12	10,12
a_2	12,11	9,9
a_3	11,11	8,10

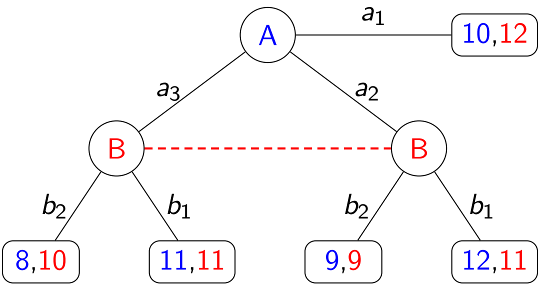

- No matter what Ann's probabilities are for playing a_2 and a_3, Bob is always better off playing b_1.
- Thus, Bob will play b_1 at his information set
- Knowing this, Ann will play a_2
- Dynamic deliberation will never lead to the "bad" equilibrium (a_1, b_2 if a_2 or a_3)

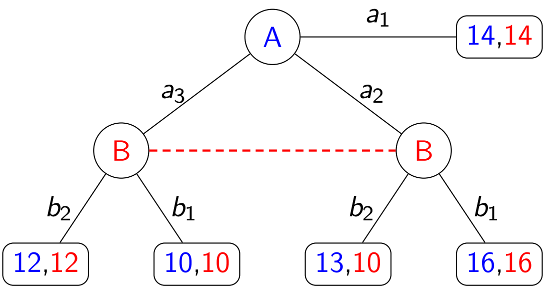

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	14,14	14,14
a_2	16,16	13,10
a_3	10,10	12,12

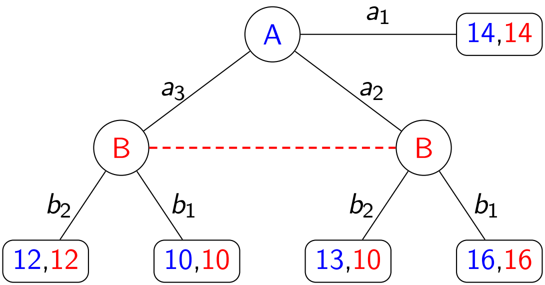

- If Ann plays a_2, Ann will get a better payoff than if Ann plays a_3 no matter what Bob does
- This will lead Bob to play b_1. Ann can figure *this* out, so she will play a_2.
- If we implement some sort of Bayes dynamics with a *tendency towards a better response*, then deliberation will lead to the "good" equilibrium.

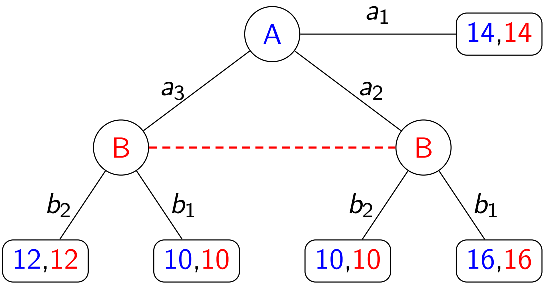

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	14,14	14,14
a_2	16,16	10,10
a_3	10,10	12,12

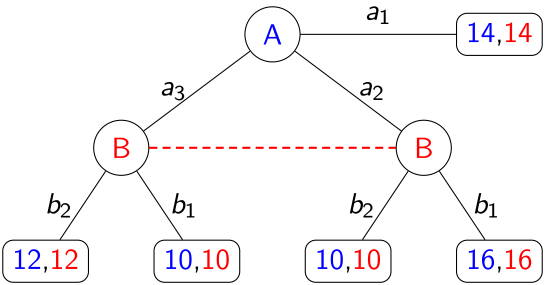

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- In the preceding example, $p_A(a_2 \mid a_2 \text{ or } a_3)$ is high because a_2 strictly dominates a_3.

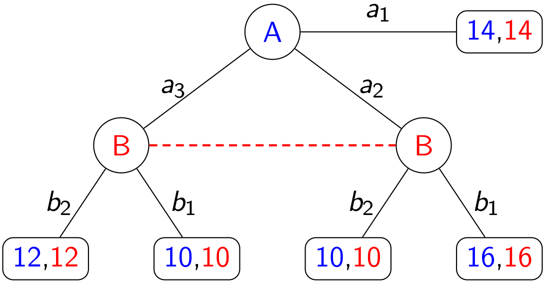

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- In this example, $p_A(a_2 \mid a_2 \text{ or } a_3)$ is high because a_1 strictly dominates a_3.

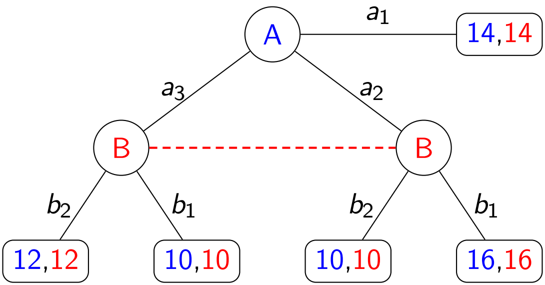

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- Unless there is some "pre-deliberational" pruning, Darwin dynamics can lead to either equilibrium.

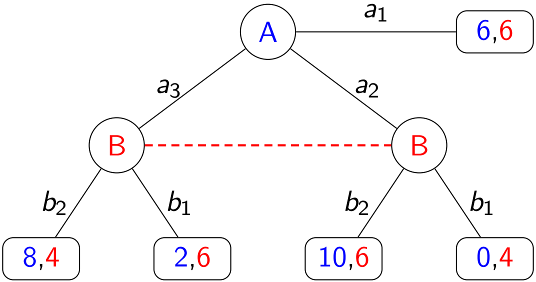

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	6,6	6,6
a_2	0,4	10,6
a_3	2,6	8,4

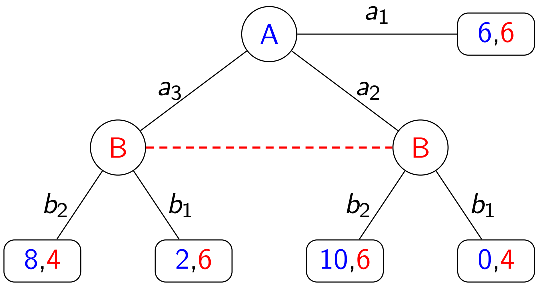

- W. Harper, Dynamic Deliberation, PSA 1992.
- (a_2, b_2 if $a2$ or a_3) and $(a_1, (0.5b_1, 0.5b_2))$ are equilibrium
- a_1 is not *ratifiable*: if Bob is given a chance to move that means Ann must be expecting Bob to choose b_2. Ann's best response to this is a_2. Knowing this Bob will choose b_2
- Starting at $(1/3a_1, 1/3a_2, 1/3a_3)$ and $(1/2b_1, 1/2b_2)$, both Darwin and Nash dynamics lead to the "bad" equilibrium.

General comments

- Extensive games (imperfect information), imprecise probabilities, other notions of stability, weaken common knowledge assumptions,...
- Generalizing the basic model.
- Relation with correlated equilibrium (correlation through rational deliberation)
- Why assume deliberators are in a "information feedback situation"?
- Deliberation in decision theory.

Stability of Equilibria

An equilibrium point e is **stable** under dynamics if points nearby remain close for all time under the action of dynamics. It is **strongly stable** if there is a neighborhood of e such that the trajectories of all points in that neighborhood converge to e.

- In the game of chicken: the two pure equilibria are strongly stable while the mixed equilibria is not stable.
- Pure equilibria can be dynamically unstable (Myerson's game)
- Mixed equilibria can be strongly stable (Matching Pennies)
- A pure strategy may be highly unstable (Moulin's game)

Moulin's Game

		Bob L	C	R
Ann	U	1,3	2,0	3,1
	M	0,2	2,2	0,2
	D	3,1	2,0	1,3

Mixed Equilibria

In static discussion of game theory, it is often remarked that mixed equilibria are unstable because if your opponent plays the equilibrium strategy, then you can always do just as well by playing any pure strategy with positive weight in your mixed equilibrium strategy than by playing the mixed equilibrium itself.

Mixed Equilibria

In static discussion of game theory, it is often remarked that mixed equilibria are unstable because if your opponent plays the equilibrium strategy, then you can always do just as well by playing any pure strategy with positive weight in your mixed equilibrium strategy than by playing the mixed equilibrium itself.

If opponents are understood as dynamic deliberators, then a mixed strategy may or may not be stable.

Imprecise Priors

It is assumed that the players precise states of indecision are common knowledge at the onset of deliberation.

Imprecise Prior: Each players prior is a convex set of probability measures over her actions space.

Restrict attention to games with two players where each players has two strategies.

A precise state of indecision for the row player is

$$\mathbf{P}_{row}(t) = \langle p^1_{row}(t), \ldots, p^n_{row}(t) \rangle$$

where $p^j_{row}(t)$ is the probability that row assigns to her strategy j at time t.

An imprecise state of indecision has $p^1_{row} = [lp, up]$ and $p^2_{row} = [1 - up, 1 - lp]$. For example, if $p^1_{row} = [0.6, 0.7]$, then $p^2_{row} = [0.3, 0.4]$.

Row (Col) has an expected utility for each probability measure in Col's (Row's) interval. Row (Col) need only compute expected utilities with respect to the endpoints of columns interval.

		col	
		L	R
row	U	0,0	1,1
	D	1,1	0,0

$$p^U_{row}(0) = [0.6, 0.8] \text{ and } p^L_{col}(0) = [0.6, 0.9]$$

		col	
		L	R
row	U	0,0	1,1
	D	1,1	0,0

$$p^U_{row}(0) = [0.6, 0.8] \text{ and } p^L_{col}(0) = [0.6, 0.9]$$

$EU_{row}(U, 0) = [0.1, 0.4]$
$EU_{row}(D, 0) = [0.6, 0.9]$

	col	
row	L	R
U	0,0	1,1
D	1,1	0,0

$$p^U_{row}(0) = [0.6, 0.8] \text{ and } p^L_{col}(0) = [0.6, 0.9]$$

$EU_{row}(U, 0) = [0.1, 0.4]$
$EU_{row}(D, 0) = [0.6, 0.9]$
How should you calculate $\mathbf{P}_{row}(1)$ and $\mathbf{P}_{col}(1)$?

1. $p^U_{row} = 0.6$, $p^L_{col} = 0.6$: $SQ_{row} = 0.30$, $Cov_{row}(U) = 0$, $Cov_{row}(D) = 0.30$. $p^U_{row}(1) = \frac{0.6+0}{1+0.3} = 0.4615$

2. $p^U_{row} = 0.6$, $p^L_{col} = 0.9$: $SQ_{row} = 0.40$, $Cov_{row}(U) = 0$, $Cov_{row}(D) = 0.20$. $p^U_{row}(1) = \frac{0.6+0}{1+0.4} = 0.4286$

3. $p^U_{row} = 0.8$, $p^L_{col} = 0.6$: $SQ_{row} = 0.32$, $Cov_{row}(U) = 0$, $Cov_{row}(D) = 0.28$. $p^U_{row}(1) = \frac{0.8+0}{1+0.32} = 0.6061$

4. $p^U_{row} = 0.8$, $p^L_{col} = 0.9$: $SQ_{row} = 0.20$, $Cov_{row}(U) = 0$, $Cov_{row}(D) = 0.7$. $p^U_{row}(1) = \frac{0.8+0}{1+0.7} = 0.4706$

$p^U_{row} = [0.4286, 0.6061]$

- The area of a rectangle of indecision need not be preserved by deliberational dynamics

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49], [0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49], [0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.
- In matching pennies, the mixed strategy is strongly stable. However, starting from $[0.51, 0.49], [0.51, 0.49]$, the imprecision explodes to cover the whole space (see Figure 3.8, pg. 72)

- The area of a rectangle of indecision need not be preserved by deliberational dynamics
- For example, players may start out with imprecise prior probabilities and deliberation results in point probabilities (E.g., Figure 3.4, 3.5 on pgs. 68, 69)
- The pure mixed strategy in the game of Chicken is not stable for precise probabilities. Starting from $[0.51, 0.49], [0.51, 0.49]$, the orbit explodes to a state of mutual total bewilderment.
- In matching pennies, the mixed strategy is strongly stable. However, starting from $[0.51, 0.49], [0.51, 0.49]$, the imprecision explodes to cover the whole space (see Figure 3.8, pg. 72)
- When analyzed in terms of precise priors, the pure coordination game and Chicken were both seen to be situations in which coordination could arise spontaneously. This is not true when starting with imprecise probabilities.

J. McKenzie Alexander. *Local interactions and the dynamics of rational deliberation*. Philosophical Studies 147 (1), 2010.

Consider a social network $\langle N, E \rangle$ (connected graph)

Consider a social network $\langle N, E\rangle$ (connected graph)

Convention: If there is a directed edge from A to B, then A always plays row and B always play column, and the interactions of Row and Column are symmetric in the available strategies.

Consider a social network $\langle N, E \rangle$ (connected graph)

Convention: If there is a directed edge from A to B, then A always plays row and B always play column, and the interactions of Row and Column are symmetric in the available strategies.

Let $\nu_i = \{i_1, \ldots i_j\}$ be i's neighbors

Consider a social network $\langle N, E\rangle$ (connected graph)

Convention: If there is a directed edge from A to B, then A always plays row and B always play column, and the interactions of Row and Column are symmetric in the available strategies.

Let $\nu_i = \{i_1, \ldots i_j\}$ be i's neighbors

$\mathbf{p'_{a,b}(t+1)}$ is represents the incremental refinement of player a's state of indecision given his knowledge about player b's state of indecision (at time $t+1$).

Consider a social network $\langle N, E\rangle$ (connected graph)

Convention: If there is a directed edge from A to B, then A always plays row and B always play column, and the interactions of Row and Column are symmetric in the available strategies.

Let $\nu_i = \{i_1, \dots i_j\}$ be i's neighbors

$\mathbf{p'_{a,b}(t+1)}$ is represents the incremental refinement of player a's state of indecision given his knowledge about player b's state of indecision (at time $t+1$).

Pool this information to form your new probabilities:

$$\mathbf{p}_i(t+1) = \sum_{j=1}^{k} w_{i,i_j}\mathbf{p}'_{i,i_j}(t+1)$$

- Allowing for local interactions in the dynamics of rational deliberation breaks the link between convergent points of the deliberative dynamics and Nash equilibrium points of the underlying game.
- It is no longer true that all dynamical rules have fixed points that maximize expected utility of the status quo.
- The effect of local interactions reveals reasons for preferring the Bayesian dynamics over the Nash dynamics.

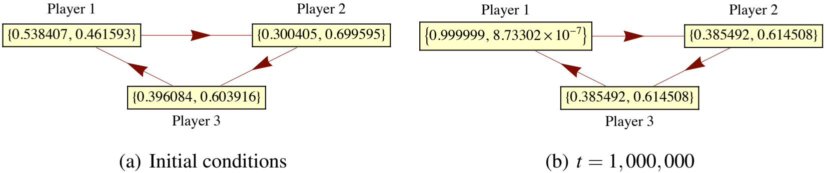

Fig. 4 The game of Chicken played on a three-person directed cycle with Nash deliberators having an index of caution of 1. Probabilities shown as (Don't Swerve, Swerve).

Out of one thousand simulations using Darwin dynamics having an index of caution 100, *all* of them converge to a state where one player assigned probability 1 to Don't Swerve and the other two assign probability 1 to Swerve.

		Billy	
		Boxing	Ballet
Maggie	Boxing	$(2,1)$	$(0,0)$
	Ballet	$(0,0)$	$(1,2)$

Fig. 7 The game of Battle of the Sexes.

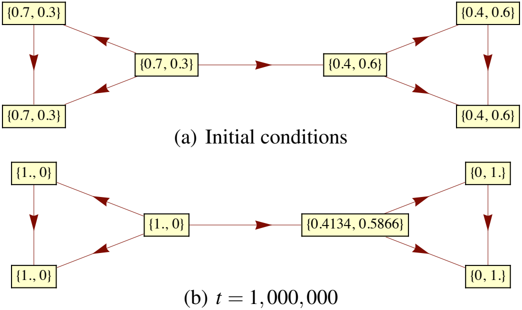

Fig. 8 Battle of the Sexes played by Nash deliberators ($k = 25$) on two cycles connected by a bridge edge (values rounded to the nearest 10^{-4}).

Deliberation in Games

- The Harsanyi-Selten tracing procedure
- Brian Skyrms' models of "dynamic deliberation"
- Ken Binmore's analysis using Turing machines to "calculate" the rational choice
- Robin Cubitt and Robert Sugden's "reasoning based expected utility procedure"
- Johan van Benthem et col.'s "virtual rationality *announcements*"

Deliberation in Games

- The Harsanyi-Selten tracing procedure
- Brian Skyrms' models of "dynamic deliberation"
- Ken Binmore's analysis using Turing machines to "calculate" the rational choice
- Robin Cubitt and Robert Sugden's "reasoning based expected utility procedure"
- Johan van Benthem et col.'s "virtual rationality *announcements*"

Different frameworks, common thought: *the "rational solutions" of a game are the result of individual deliberation about the "rational" action to choose.*