Reasoning in Games

Lecture 4

Eric Pacuit

University of Maryland, College Park

pacuit.org

epacuit@umd.edu

August 13, 2015

Plan

✓ Day 1: Decision Theory

✓ Day 2: From Decisions to Games

✓ Day 3: Game Models

- Day 4: Modeling Deliberation (in Games)

- Day 5: Backward and Forward Induction, Concluding Remarks (Language-Based Games/ Variable Frame Theory, Behavioral Game Theory, ...)

Summary

- Game models describe the *informational context* of a game.

- Interpreting mixed strategies: Epistemic interpretation, *purification theorem*

- Game models can be used to *characterize* different solution concepts (e.g., iterated strict dominance, iterated weak dominance, Nash equilibrium, correlated equilibrium,...)

Next Steps

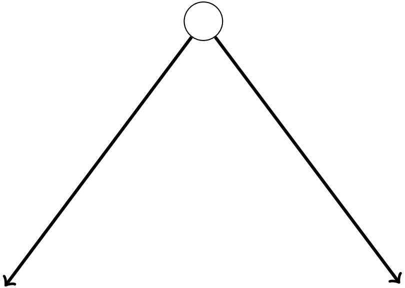

Moving beyond the basic model: Belief revision

Modeling deliberation (in games)

Next Steps

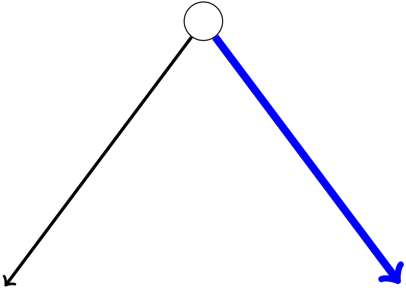

Moving beyond the basic model: Belief revision

Modeling deliberation (in games)

Strategic Reasoning

"The word *eductive* will be used to describe a dynamic process by means of which equilibrium is achieved through careful reasoning on the part of the players. Such reasoning will usually require an attempt to simulate the reasoning processes of the other players. Some measure of pre-play communication is therefore implied, although this need not be explicit. To reason along the lines "if I think that he thinks that I think..." requires that information be available on how an opponent thinks."

(pg. 184)

K. Binmore. *Modeling Rational Players*. Economics and Philosophy, 3,179 - 21, 1987.

Deliberational Decision Theory

F. Arntzenius. *No Regrest, or: Edith Piaf Revamps Decision Theory*. Erkenntnis, 68, pgs. 277 - 297, 2008.

J. Joyce. *Regret and Instability in Causal Decision Theory*. Synthese, 187: 1, pgs. 123 - 145, 2012.

I. Douven. *Decision theory and the rationality of further deliberation*. Economics and Philosophy, 18, pgs. 303 - 328, 2002.

Deliberational Decision Theory

Current Evaluation: If Pr_t characterizes your beliefs at time t, then at t you should *evaluate* each act by its (causal, evidential) expected utility computed using Pr_t.

Deliberational Decision Theory

Current Evaluation: If Pr_t characterizes your beliefs at time t, then at t you should *evaluate* each act by its (causal, evidential) expected utility computed using Pr_t.

Full Information: You should act on your time-t utility assessments only if those assessments are based on beliefs that incorporate *all* the evidence that is both freely available to you at t and relevant to the question about what your acts are likely to cause.

Deliberational Decision Theory

Current Evaluation: If Pr_t characterizes your beliefs at time t, then at t you should *evaluate* each act by its (causal, evidential) expected utility computed using Pr_t.

Full Information: You should act on your time-t utility assessments only if those assessments are based on beliefs that incorporate *all* the evidence that is both freely available to you at t and relevant to the question about what your acts are likely to cause.

Sometimes initial opinions fix actions, *but not always* (e.g., Murder Lesion, Psychopath Button)

Modeling Rational Deliberation

$$\mathcal{M}_0 \Longrightarrow \mathcal{M}_1 \Longrightarrow \mathcal{M}_2 \Longrightarrow \cdots$$

initial model

Modeling Rational Deliberation

$$\mathcal{M}_0 \Longrightarrow \mathcal{M}_1 \Longrightarrow \mathcal{M}_2 \Longrightarrow \cdots$$

initial model

Each $\mathcal{M}_i$ describes the decision maker's current thoughts about what might happen during a play of the game (her beliefs and "inclinations").

Modeling Rational Deliberation

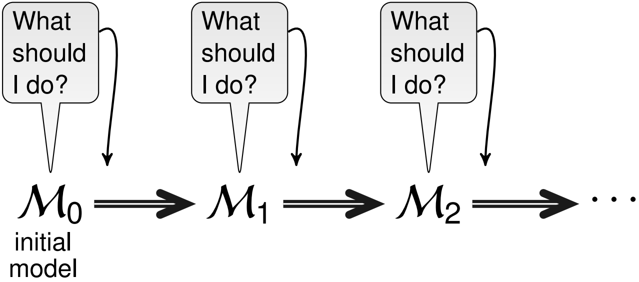

Dynamical rules transform the decision maker's beliefs, given her evaluation of the available acts.

Modeling Rational Deliberation

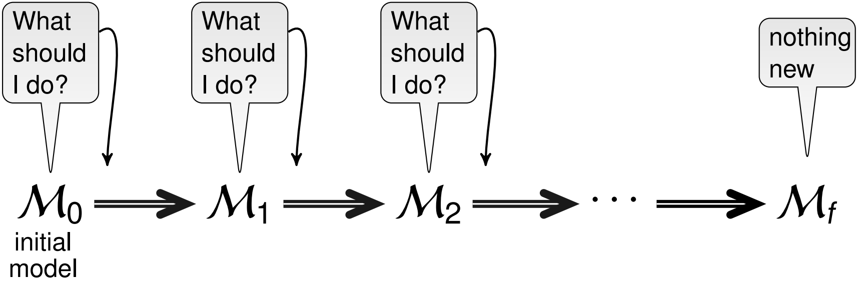

Deliberations stops when a “fixed-point” is reached.

Deliberation in games

- The Harsanyi-Selten tracing procedure
- Brian Skyrms' model of "dynamic deliberation"
- 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.*

- What operations transform the models?
- Where does the “new information” come from? What are player *i*’s opponents thinking about doing? (“update by emulation”)
- Why keep deliberating?

Information Feedback

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

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.”

Deliberation in Games

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

		B	
		L	R
A	U	2,1	0,0
	D	0,0	1,2

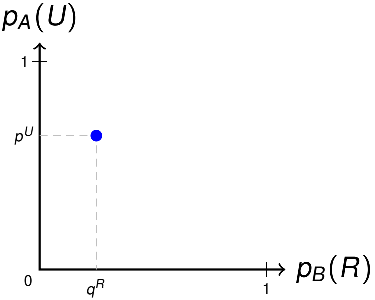

A's current state of indecision

		B	
		L	*R*
A	*U*	2,1	0,0
	D	0,0	1,2

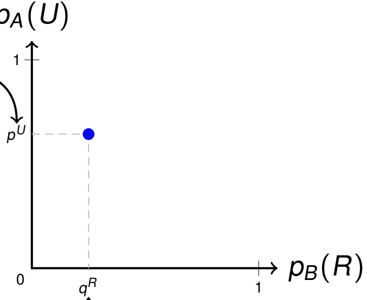

A's current belief about what *B* is going to do

B's current belief about what *A* is going to do

		B	
		L	*R*
A	*U*	2,1	0,0
	D	0,0	1,2

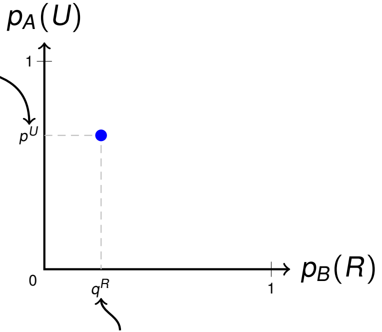

B's current state of indecision

A \ B	L	R
U	2,1	0,0
D	0,0	1,2

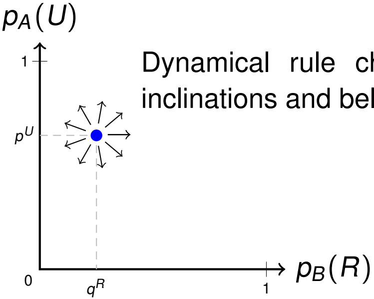

Dynamical rule changing inclinations and beliefs

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 their payoffs.

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 their payoffs.

1. Starting from an initial position, player i calculates her expected utility and moves by her dynamical 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 their payoffs.

1. Starting from an initial position, player i calculates her expected utility and moves by her dynamical 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 their payoffs.

1. Starting from an initial position, player *i* calculates her expected utility and moves by her dynamical 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 strategy j at time t.

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

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

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

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

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

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

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

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

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

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

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

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

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)

		Bob	
		L	R
Ann	U	2,1	0,0
	D	0,0	1,2

$\mathbf{P}_A = \langle 0.2, 0.8\rangle$ and $\mathbf{P}_B = \langle 0.4, 0.6\rangle$

$$EU(U) = 0.4 \cdot 2 + 0.6 \cdot 0 = 0.8$$
$$EU(D) = 0.4 \cdot 0 + 0.6 \cdot 1 = 0.6$$
$$EU(L) = 0.2 \cdot 1 + 0.8 \cdot 0 = 0.2$$
$$EU(R) = 0.2 \cdot 0 + 0.8 \cdot 2 = 1.6$$
$$SQ_A = 0.2 \cdot EU(U) + 0.8 \cdot EU(D) = 0.2 \cdot 0.8 + 0.8 \cdot 0.6 = 0.64$$
$$SQ_B = 0.4 \cdot EU(L) + 0.6 \cdot EU(R) = 0.4 \cdot 0.2 + 0.6 \cdot 1.6 = 1.04$$

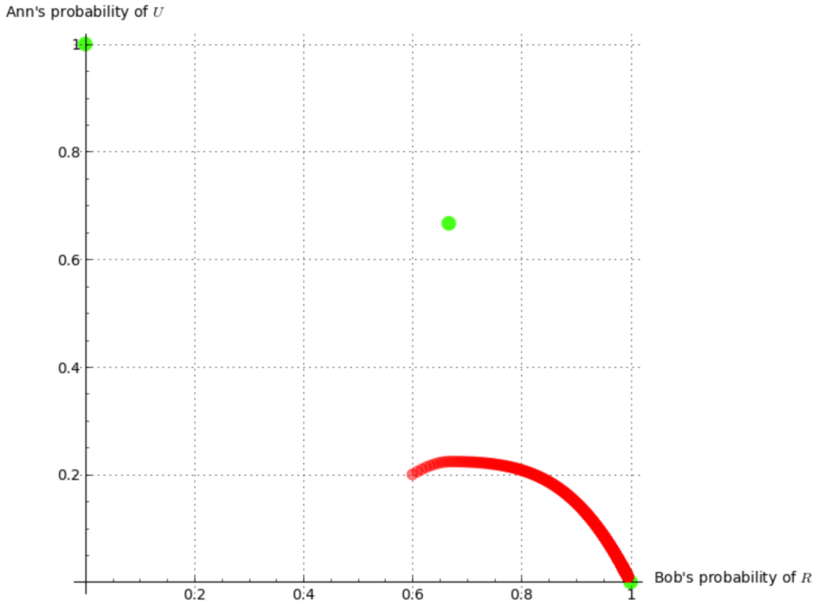

Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Ann \ Bob	L	R
U	2,1	0,0
D	0,0	1,2

$\mathbf{P}_A = \langle 0.2, 0.8 \rangle$ and $\mathbf{P}_B = \langle 0.1, 0.9 \rangle$

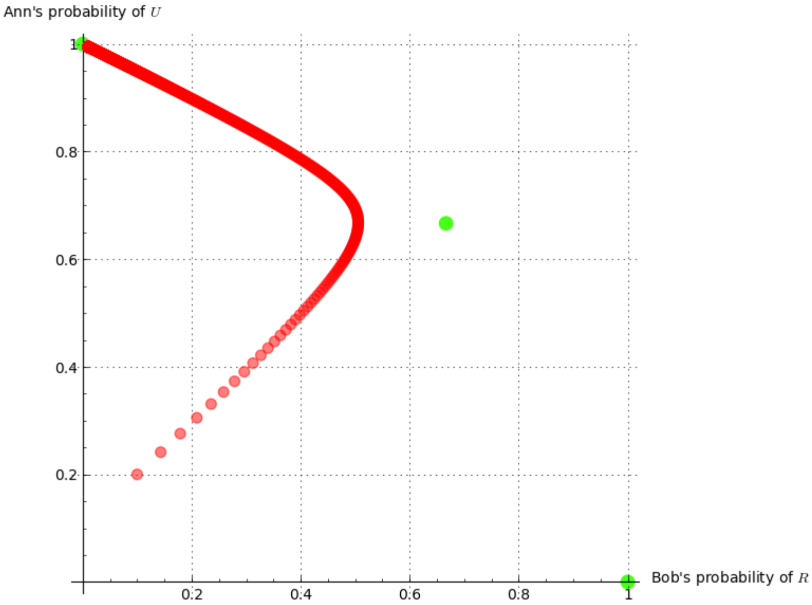
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Nash Dynamics

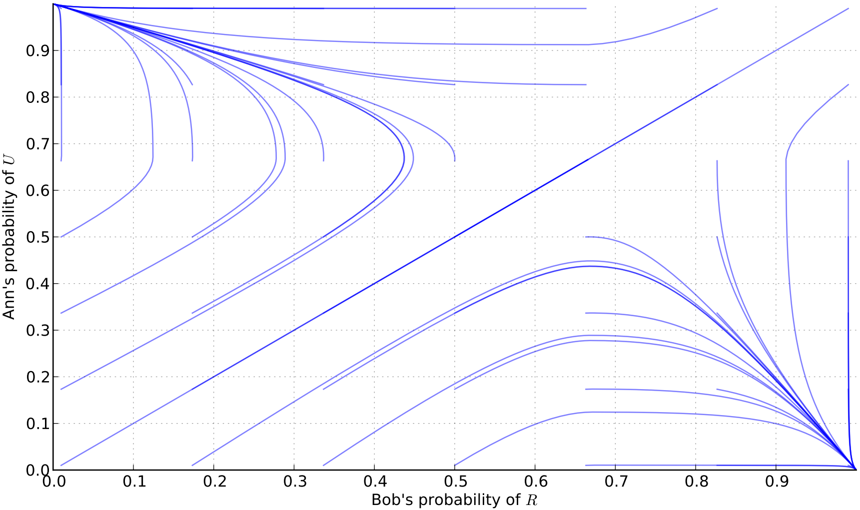

Normal form vs. Extensive form

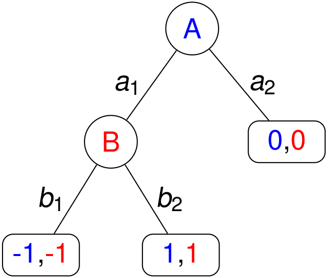

Normal form vs. Extensive form

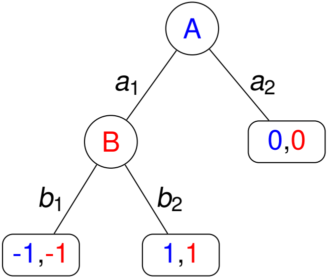

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

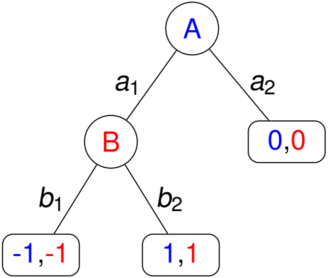

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

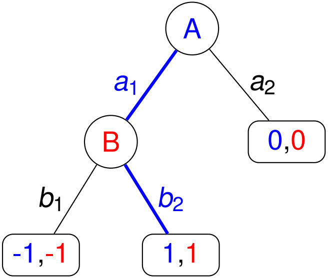

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

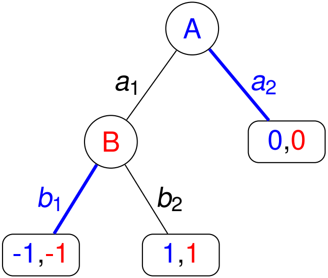

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

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

Normal form vs. Extensive form

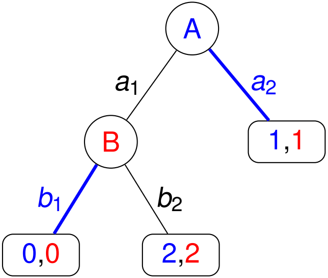

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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

Normal form vs. Extensive form

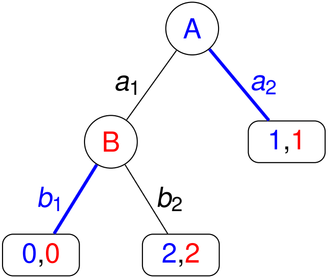

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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 $P_B(a_2 \mid a_1) = 0$).

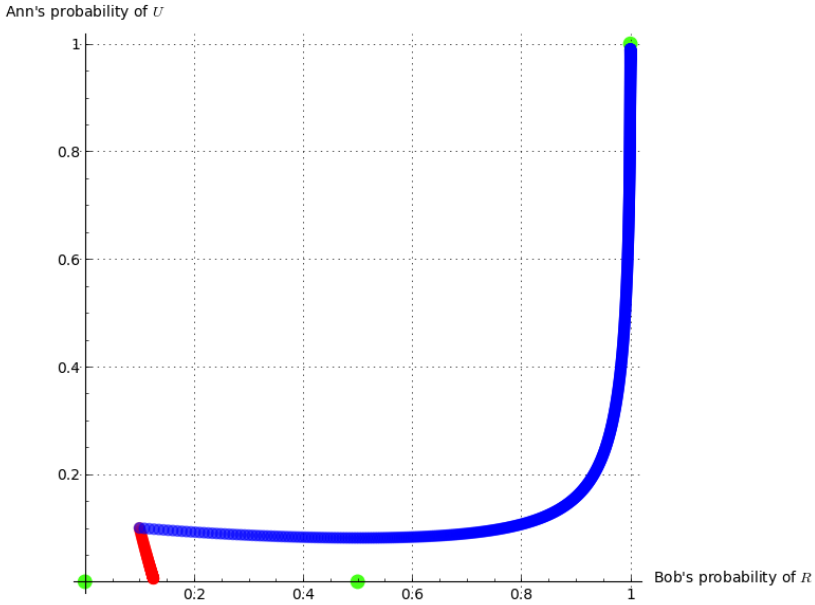
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Update by Emulation

- Nash and Bayes
- Nash vs. Bayes
- Fixed beliefs
- Uncertainty about the other player's beliefs
- Social networks

Nash vs. Bayes, I

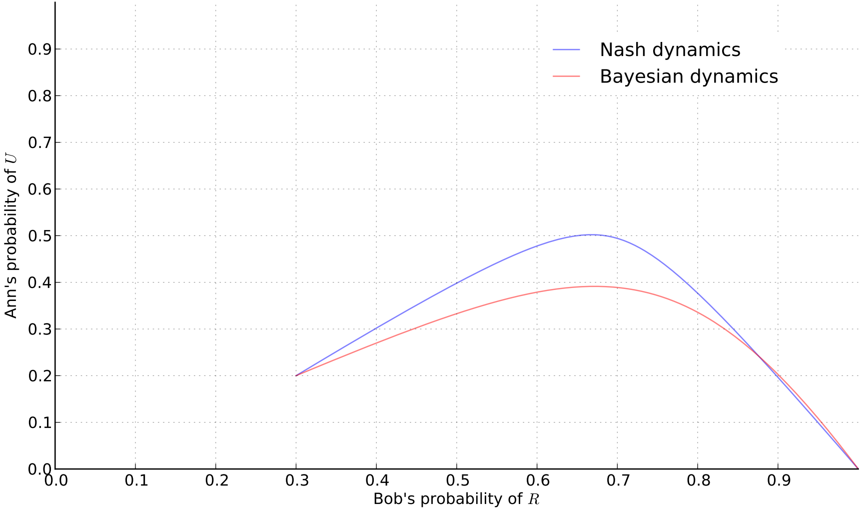

Nash vs. Bayes, I

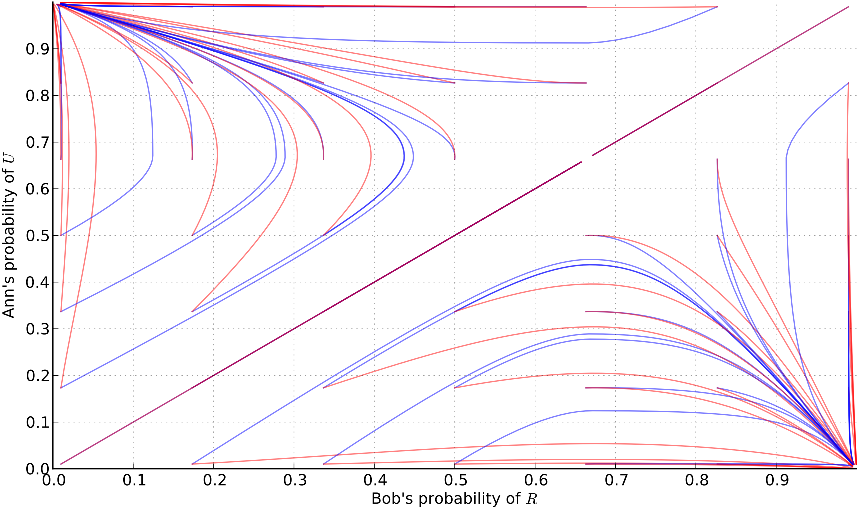

Nash vs. Bayes, II

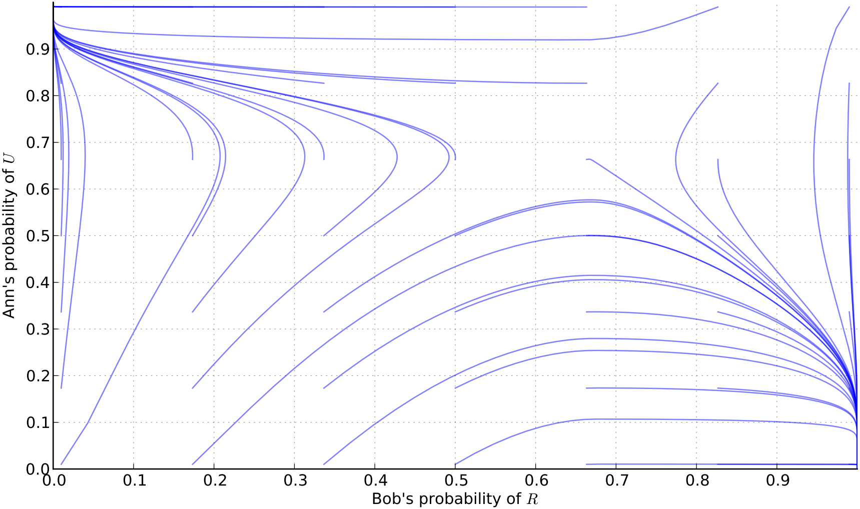

Fixed Beliefs

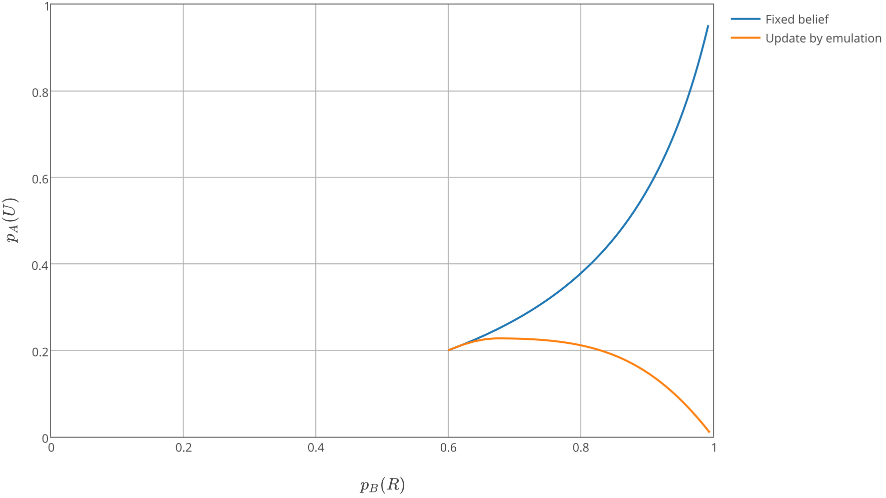

Errors

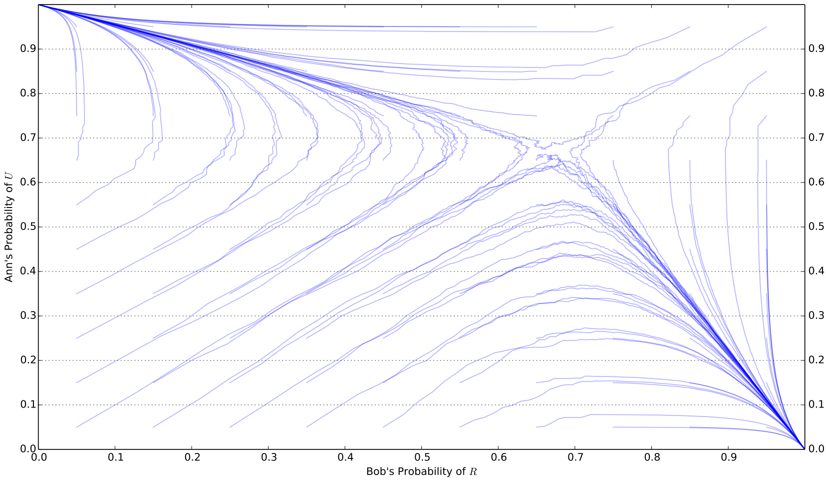

Ann's Probability of U
Bob's Probability of R
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9

Errors

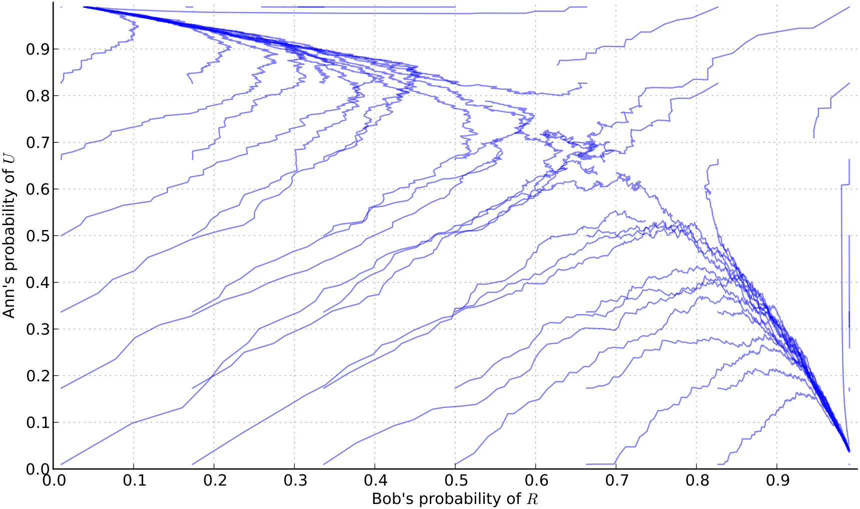

0.9
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
Ann's probability of U
0.0
0.1
0.2
0.3
0.4
0.5
0.6
0.7
0.8
0.9
Bob's probability of R

Errors

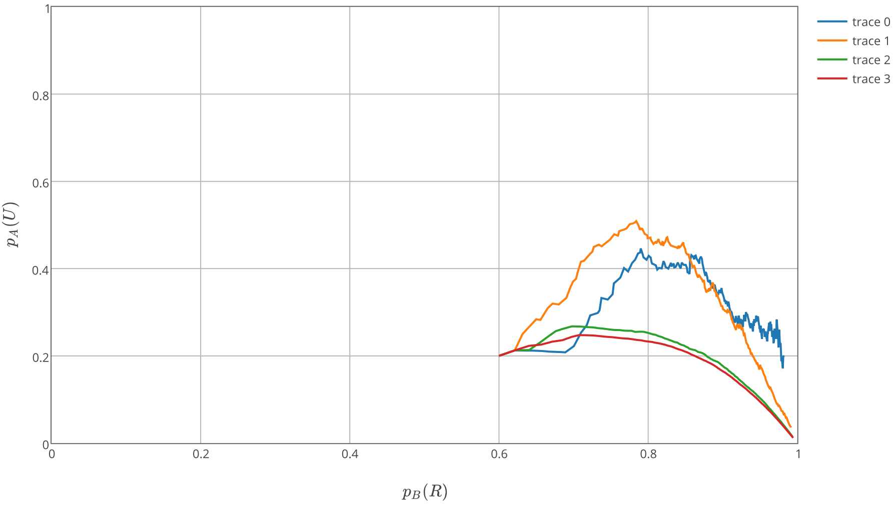

trace 0
trace 1
trace 2
trace 3
1
0.8
0.6
0.4
0.2
0
0
0.2
0.4
0.6
0.8
1
$p_A(U)$
$p_B(R)$

Learning to Play

Theorem. If players start with subjectively rational strategies, and if their individual subjective beliefs regarding opponents' strategies are "compatible with truly chosen strategies", then they must converge in a finite amount of time to play according to an ϵ-Nash in the repeated game.

E. Kalai and E. Lehrer. *Rational Learning Leads to Nash Equilibrium*. Econometrica, 61:5, pgs. 1019 - 1045, 1993.

Y. Shoham, R. Powers and T. Granager. *If multi-agent learning is the answer, what is the question?*. Artificial Intelligence, 171(7), pgs. 365 - 377, 2007.

Modeling Deliberation in Games

- Characterize outcomes in terms of *accessibility* and/or *stability*
- Relation with *correlated equilibrium* (correlation through rational deliberation)

Modeling Deliberation in Games

- Characterize outcomes in terms of *accessibility* and/or *stability*
- Relation with *correlated equilibrium* (correlation through rational deliberation)
- Comparison with other models of deliberation in games (categorize pure strategies)

Modeling Deliberation in Games

- Characterize outcomes in terms of *accessibility* and/or *stability*
- Relation with *correlated equilibrium* (correlation through rational deliberation)
- Comparison with other models of deliberation in games (categorize pure strategies)
- Generalize the basic model: extensive games (with imperfect information), imprecise probabilities, more than two players
- Weaken the common knowledge assumptions (payoffs, beliefs, dynamical rule, updating by emulation)

Modeling Deliberation in Games

- Characterize outcomes in terms of *accessibility* and/or *stability*
- Relation with *correlated equilibrium* (correlation through rational deliberation)
- Comparison with other models of deliberation in games (categorize pure strategies)
- Generalize the basic model: extensive games (with imperfect information), imprecise probabilities, more than two players
- Weaken the common knowledge assumptions (payoffs, beliefs, dynamical rule, updating by emulation)
- Deliberation in decision theory ("deliberation crowds out prediction", logical omniscience)

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

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

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

$EU_{row}(U, 0) = [0.1, 0.4]$
$EU_{row}(D, 0) = [0.6, 0.9]$
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 $\nu_i = \{i_1, \ldots i_j\}$ 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 $t+1$).

Pool this information to form your new probabilities:

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

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

Fig. 7 The game of Battle of the Sexes.

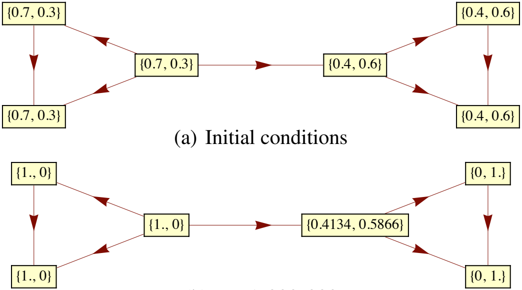

(a) Initial conditions

(b) $t = 1,000,000$

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' model of "dynamic deliberation"
- Robin Cubitt and Robert Sugden's "reasoning based expected utility procedure"
- Johan van Benthem et col.'s "virtual rationality *announcements*"

EP. *Dynamic models of rational deliberation in games.* in *Strategic Reasoning*, van Benthem, Gosh, and Verbrugge, ed., 2015.

Reasoning Based Expected Utility Procedure

R. Cubitt and R. Sugden. *The reasoning-based expected utility procedure*. Games and Economic Behavior, 2010.

Reasoning-Based Solution Concepts

Reasoning-Based Solution Concepts

A **categorization** is a *ternary* partition of the players choices (rather than a binary partition of what is in and what is out):

Reasoning-Based Solution Concepts

A **categorization** is a *ternary* partition of the players choices (rather than a binary partition of what is in and what is out): strategies are accumulated, deleted or neither.

Reasoning-Based Solution Concepts

A **categorization** is a *ternary* partition of the players choices (rather than a binary partition of what is in and what is out): strategies are accumulated, deleted or neither.

Example: RBEU (reasoning based expected utility):

Reasoning-Based Solution Concepts

A **categorization** is a *ternary* partition of the players choices (rather than a binary partition of what is in and what is out): strategies are accumulated, deleted or neither.

Example: RBEU (reasoning based expected utility):

- accumulate strategies that maximize expected utility for every possibly probability distribution
- delete strategies that do not maximize probability against any probability distribution

Reasoning-Based Solution Concepts

A **categorization** is a *ternary* partition of the players choices (rather than a binary partition of what is in and what is out): strategies are accumulated, deleted or neither.

Example: RBEU (reasoning based expected utility):

- accumulate strategies that maximize expected utility for every possibly probability distribution
- delete strategies that do not maximize probability against any probability distribution
- accumulated strategies must receive positive probability, deleted strategies must receive zero probability

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$$S^+ = \{L\}$$
$$S^- = \{B\}$$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$$S^+ = \{L\}$$
$$S^- = \{B\}$$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L\}$
$S^- = \{B\}$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L\}$
$S^- = \{B\}$

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L, R\}$
$S^- = \{B, M_1\}$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$$S^+ = \{L\}$$
$$S^- = \{B\}$$

$$S^+ = \{L, R\}$$
$$S^- = \{B, M_1\}$$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$$S^+ = \{L\}$$
$$S^- = \{B\}$$

$$S^+ = \{L, R\}$$
$$S^- = \{B, M_1\}$$

RBEU: Example

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L\}$
$S^- = \{B\}$

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L, R\}$
$S^- = \{B, M_1\}$

	L	R
U	1,1	1,1
M_1	0,0	1,0
M_2	2,0	0,0
B	0,2	0,0

$S^+ = \{L, R\}$
$S^- = \{B, M_1\}$

RBEU Example 2

	l	*r*
u	1,1	0,0
d	0,0	0,0

RBEU Example 2

	l	r
u	1,1	0,0
d	0,0	0,0

$$S^+ = \{u, l\}$$
$$S^- = \emptyset$$

RBEU Example 2

	l	r
u	1,1	0,0
d	0,0	0,0

$$S^+ = \{u, l\}$$
$$S^- = \emptyset$$

	l	r
u	1,1	0,0
d	0,0	0,0

$$S^+ = \{u, l\}$$
$$S^- = \{d, r\}$$

RBEU Example 3

	L	R
u	1,1	1,0
d	1,0	0,1

RBEU Example 3

	L	R
u	1,1	1,0
d	1,0	0,1

$S^+ = \{u\}$
$S^- = \emptyset$

RBEU Example 3

	L	R
u	1,1	1,0
d	1,0	0,1

	L	R
u	1,1	1,0
d	1,0	0,1

$$S^+ = \{u\}$$
$$S^- = \emptyset$$

RBEU Example 3

	L	R
u	1,1	1,0
d	1,0	0,1

$$S^+ = \{u\}$$
$$S^- = \emptyset$$

	L	R
u	1,1	1,0
d	1,0	0,1

$$S^+ = \{u\}$$
$$S^- = \emptyset$$

R. Cubitt and R. Sugden. *Common reasoning in games: A Lewisian analysis of common knowledge of rationality*. Economics and Philosophy, 2015.

Deliberation in games

- The Harsanyi-Selten tracing procedure
- Brian Skyrms' model of "dynamic deliberation"
- Robin Cubitt and Robert Sugden's "reasoning based expected utility procedure"
- Johan van Benthem et col.'s "virtual rationality *announcements*"

EP. *Dynamic models of rational deliberation in games.* in *Strategic Reasoning*, van Benthem, Gosh, and Verbrugge, ed., 2015.

Tomorrow: Backward and Forward Induction, Concluding Remarks