Strategic reasoning in games: From beliefs about mistakes to mistaken beliefs

Eric Pacuit

Department of Philosophy
University of Maryland, College Park
pacuit.org

February 20, 2017

Plan

- Just enough game theory
- Deliberation in decision and game theory
- A model of deliberation in games
- Reasoning about extensive form games
- Discussion

Just Enough Game Theory

A **game** is a mathematical model of a social interaction that includes

- the players (N);
- the actions (strategies) the players *can* take (for $i \in N$, $S_i \neq \emptyset$);
- the players' interests (for $i \in N$, $u_i : \Pi_{i \in N} S_i \rightarrow \mathbb{R}$); and
- the "structure" of the decision problem.

Just Enough Game Theory

A **game** is a mathematical model of a social interaction that includes

- the players (N);
- the actions (strategies) the players *can* take (for $i \in N$, $S_i \neq \emptyset$);
- the players' interests (for $i \in N$, $u_i : \Pi_{i \in N} S_i \rightarrow \mathbb{R}$); and
- the "structure" of the decision problem.

It does not specify the actions that the players do take.

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

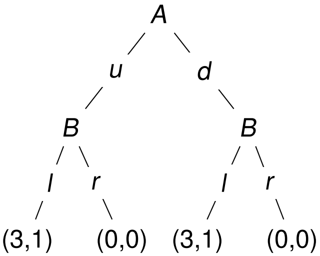

		B	
		l	r
A	u	3, 1	0, 0
	d	0, 0	1, 3

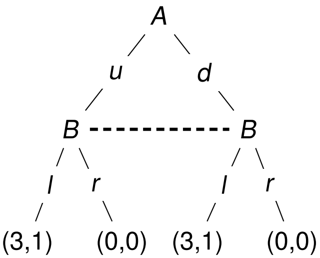

A \ B	l	r
u	3, 1	0, 0
d	0, 0	1, 3

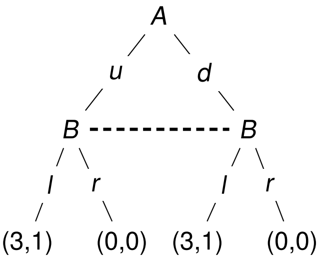

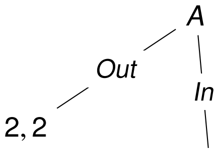

A \ B	l	r
u	3, 1	0, 0
d	0, 0	1, 3

Just Enough Game Theory: Solution Concepts

A **solution concept** is a systematic description of the outcomes that may emerge in a family of games.

This is the starting point for most of game theory and includes many variants: Nash equilibrium (and refinements), backwards induction, or iterated dominance of various kinds.

They are usually thought of as the embodiment of “rational behavior” in some way and used to analyze game situations.

Game Models

- A game is a *partial* description of a set (or sequence) of interdependent **(Bayesian) decision problems**.

Game Models

- A game is a *partial* description of a set (or sequence) of interdependent **(Bayesian) decision problems**.

  A game will not normally contain enough information to determine what the players *believe* about each other.

Game Models

- A game is a *partial* description of a set (or sequence) of interdependent **(Bayesian) decision problems**.

  A game will not normally contain enough information to determine what the players *believe* about each other.

- A **model of a game** is a completion of the partial specification of the Bayesian decision problems *and* a representation of a particular play of the game.

Game Models

- A game is a *partial* description of a set (or sequence) of interdependent **(Bayesian) decision problems**.

  A game will not normally contain enough information to determine what the players *believe* about each other.

- A **model of a game** is a completion of the partial specification of the Bayesian decision problems *and* a representation of a particular play of the game.

- There are no special rules of rationality telling one what to do in the absence of degrees of belief except: decide what you believe, and then **maximize (subjective) expected utility**.

Bayesian Decision Problem

	it rains	it does not rain
take umbrella	encumbered, dry	encumbered, dry
leave umbrella	wet	free, dry

States: it rains; it does not rain

Outcomes: encumbered, dry; wet; free, dry

Actions: take umbrella; leave umbrella

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

Utility of outcome o

$$EU(A) = \sum_{o \in O} P_A(o) \times U(o)$$

Expected utility of action A

Utility of outcome o

Probability of outcome o conditional on A

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \,\&\, A)}{P(A)}$

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \;\&\; A)}{P(A)}$

Classical: $P_A(o) = \sum_{s \in S} P(s) f_{A,s}(o)$, where

$$f_{A,s}(o) = \begin{cases} 1 & A(s) = o \\ 0 & A(s) \neq o \end{cases}$$

$P_A(o)$: probability of o conditional on A — how likely it is that outcome o will occur, on the supposition that the agent chooses act A.

Evidential: $P_A(o) = P(o \mid A) = \frac{P(o \,\&\, A)}{P(A)}$

Classical: $P_A(o) = \sum_{s \in S} P(s) f_{A,s}(o)$, where

$$f_{A,s}(o) = \begin{cases} 1 & A(s) = o \\ 0 & A(s) \neq o \end{cases}$$

Causal: $P_A(o) = P(A \mathrel{\Box\!\!\rightarrow} o)$

P("if A were performed, outcome o would ensue")

(Lewis, 1981)

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

A		$Play_B(L)$	$Play_B(R)$
	U	2	0
	D	0	1

B		$Play_A(U)$	$Play_A(D)$
	L	1	0
	R	0	2

Models of Games

Suppose that G is a game.

- Outcomes of the game: $S = \Pi_{i \in N} S_i$
- Player i's partial beliefs (or conjecture): $P_i \in \Delta(S_{-i})$

  $\Delta(X)$ *is the set of probabilities measures over* X

"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, 1984.

Reasoning in Games

"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.

Reasoning in Games

"*The* fundamental insight of game theory [is] that a rational player must take into account that the players reason about each other in deciding how to play" (pg. 81)

R. Aumann and J. Dreze. *Rational expectations in games*. American Economic Review, Vol. 98, pgs. 72 – 86 (2008).

Reasoning in Games

"*The* fundamental insight of game theory [is] that a rational player must take into account that the players reason about each other in deciding how to play" (pg. 81)

R. Aumann and J. Dreze. *Rational expectations in games*. American Economic Review, Vol. 98, pgs. 72 – 86 (2008).

Exactly *how* the players incorporate the fact that they are interacting with other (actively reasoning) rational agents is the subject of much debate.

Deliberational Decision Theory

F. Arntzenius. *No Regrets, 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 P_t characterizes your beliefs at time t, then at t you should *evaluate* each act by its (causal, evidential) expected utility computed using P_t.

Deliberational Decision Theory

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

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.” (Skyrms)

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

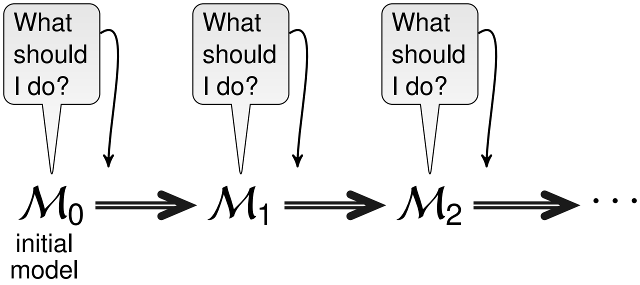

Dynamical rules transform the decision maker's beliefs, given her evaluation of the available acts.

Modeling Rational Deliberation

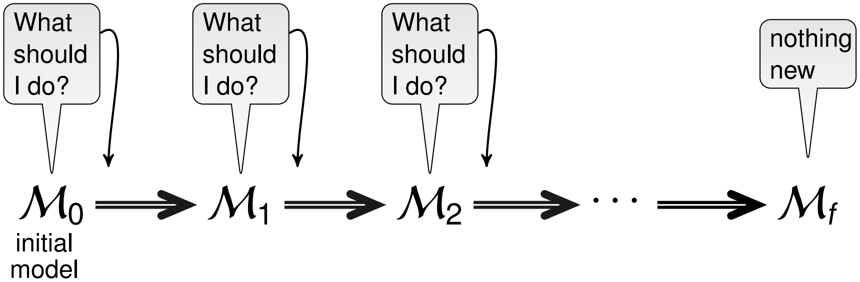

Deliberations stops when a “fixed-point” is reached.

States of Indecision

A **state of indecision** is a probability distribution over the set of acts.

States of Indecision

A **state of indecision** is a probability distribution over the set of acts.

“My tentative view is that to make a certain mixed decision is just to have certain credences in ones acts at the end of a rational deliberation.

States of Indecision

A **state of indecision** is a probability distribution over the set of acts.

"My tentative view is that to make a certain mixed decision is just to have certain credences in ones acts at the end of a rational deliberation. On this view, mixed decisions are not decisions to perform certain acts with certain probabilities....Decision theory, on this view, does not evaluate the rationality of actions. Rather it evaluates the rationality of credences in actions. On this understanding of mixed acts decision theory is theory of what credences one ought to have in ones actions, it is not a theory that tells one which actions are rational and which are not, nor does it even evaluate how rational each possible act is." (Arntzenius)

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.*

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

Strategic Game

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

Strategic Game

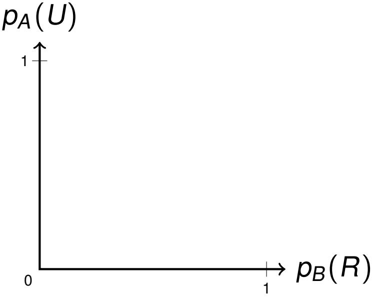

Solution Space

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

Strategic Game

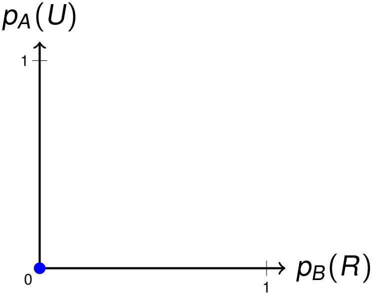

Solution Space

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

Strategic Game

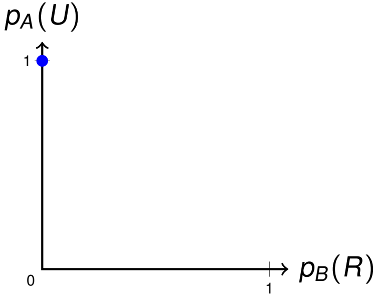

Solution Space

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

Strategic Game

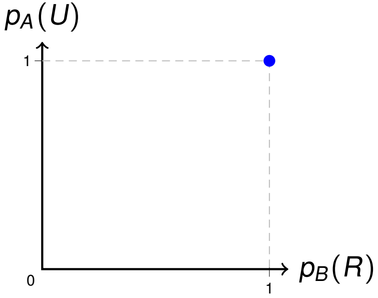

Solution Space

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

Strategic Game

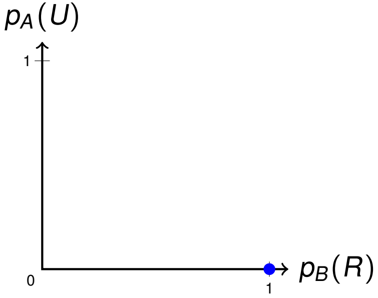

Solution Space

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

Strategic Game

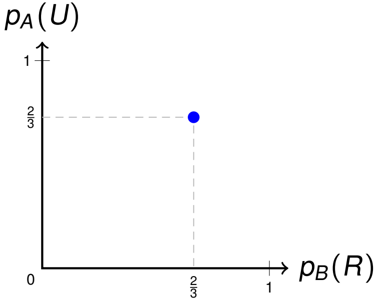

Solution Space

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

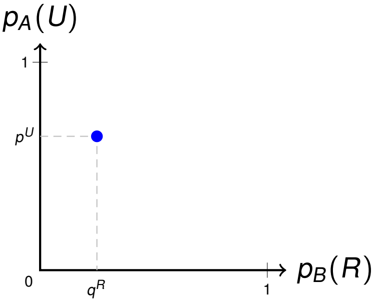

A's current state of indecision

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

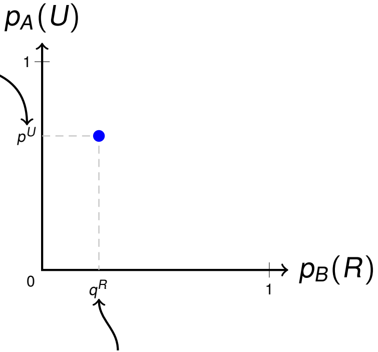

B's current state of indecision

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

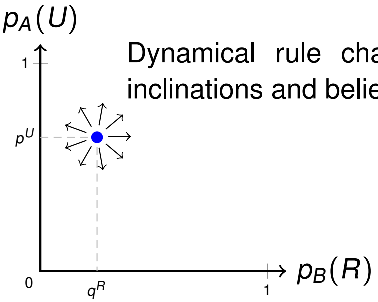

Dynamical rule changing inclinations and beliefs

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

A		Play$_B(L)$	Play$_B(R)$
A	U	2	0
A	D	0	1

B		Play$_A(U)$	Play$_A(D)$
B	L	1	0
B	R	0	2

State of indecision at time t

$\mathbf{P}_A(t) = \langle p_A^U(t), p_A^D(t) \rangle$ $\quad$ $\mathbf{P}_B(t) = \langle p_B^L(t), p_B^R(t) \rangle$

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

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

Dynamical rule

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

Dynamical rule

State of indecision

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

Dynamical rule

State of indecision

Beliefs about the state of nature

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

Dynamical rule

State of indecision

Beliefs about the state of nature

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

		$\text{Play}_B(L)$	$\text{Play}_B(R)$
A	U	2	0
A	D	0	1

		$\text{Play}_A(U)$	$\text{Play}_A(D)$
B	L	1	0
B	R	0	2

Beliefs at time t

$\mathbf{B}_A(t) = \langle p_A^L(t), p_A^R(t)\rangle$ $\mathbf{B}_B(t) = \langle p_B^U(t), p_B^D(t)\rangle$

Update by Emulation

For each player, the decisions of the other players constitute the relevant state of the world, which together with her decision, determines their payoffs.

Update by Emulation

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.

Update by Emulation

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.

Update by Emulation

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.

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

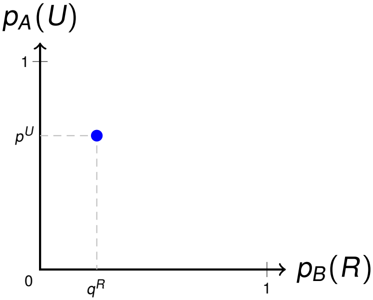

A's current state of indecision

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

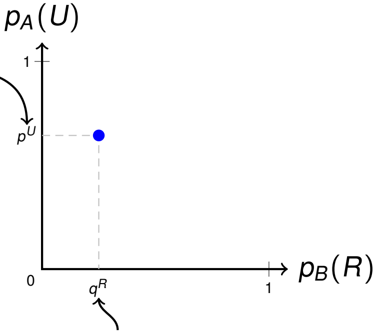

A's current belief about what *B* is going to do

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

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

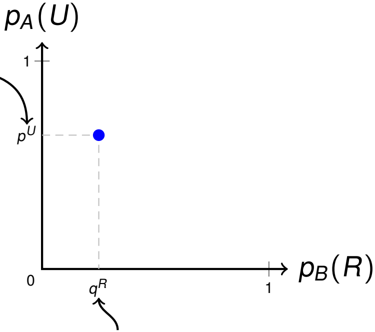

B's current state of indecision

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

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

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.

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

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

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

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

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

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

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

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

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

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)

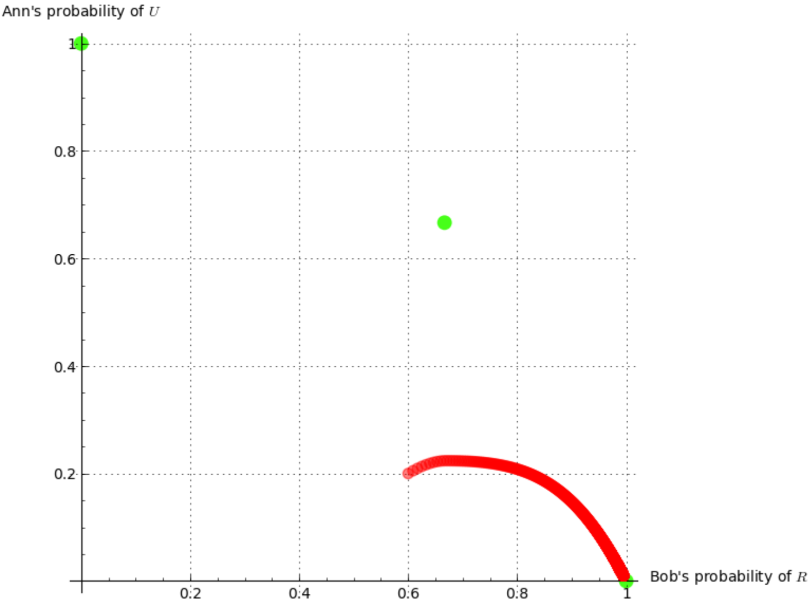
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

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

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

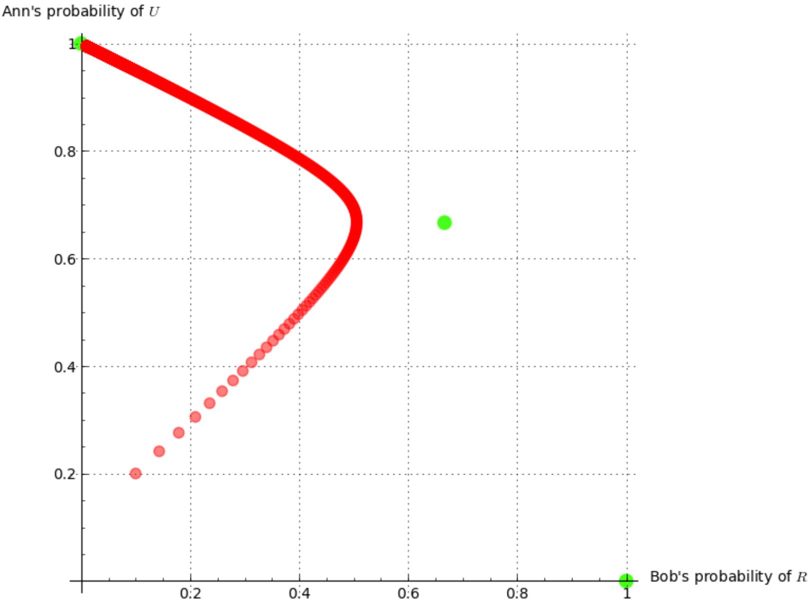

Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Nash Dynamics

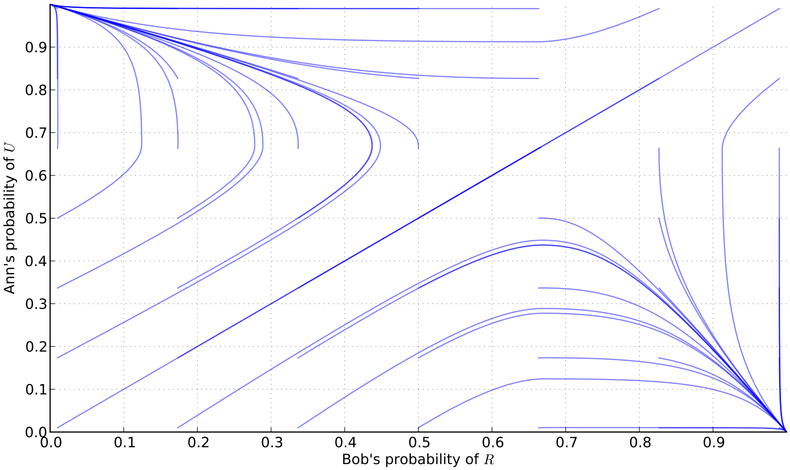

Modeling Deliberation in Games

- 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, rationality “announcements”)

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, rationality “announcements”)
- Generalize the basic model: 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, rationality "announcements")
- Generalize the basic model: 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)

Rational deliberation in extensive form games

Normal form vs. Extensive form

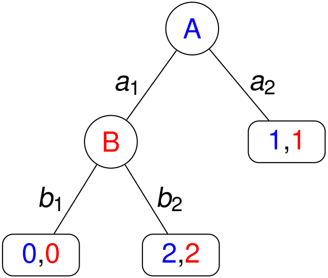

Normal form vs. Extensive form

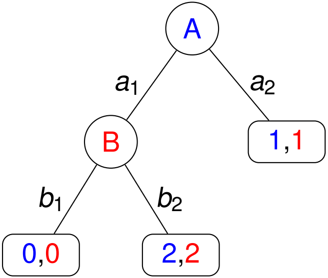

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

Normal form vs. Extensive form

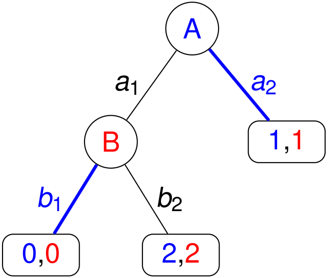

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

Normal form vs. Extensive form

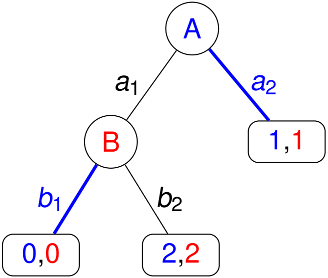

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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

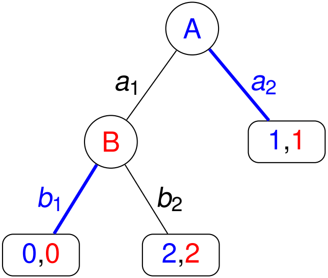

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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

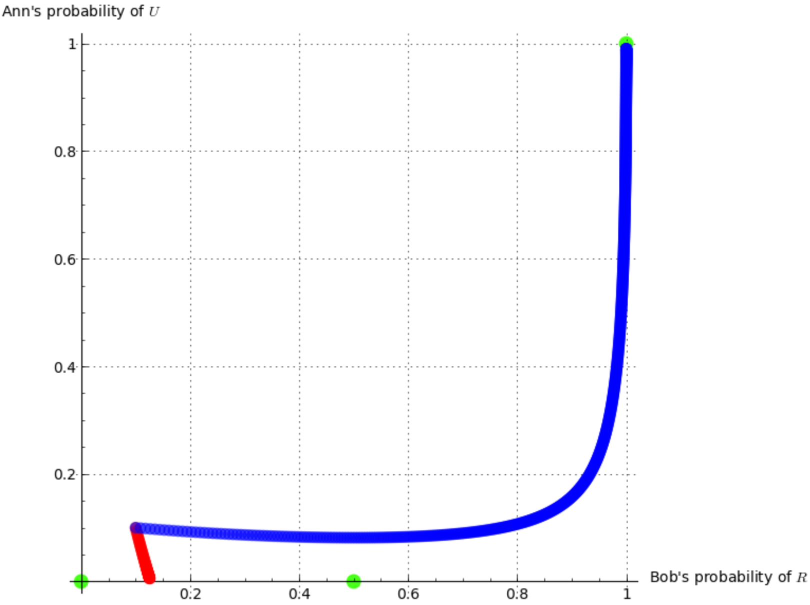
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

BI Puzzle?

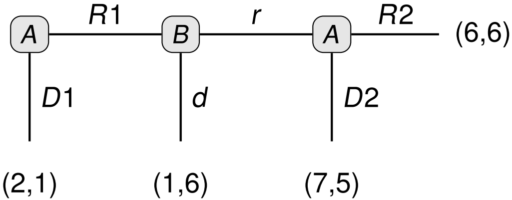

BI Puzzle?

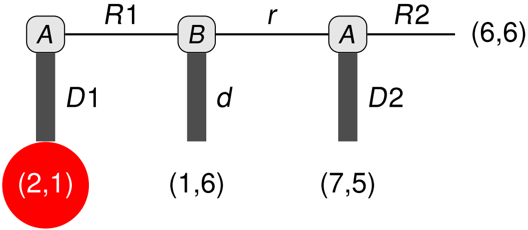

BI Puzzle?

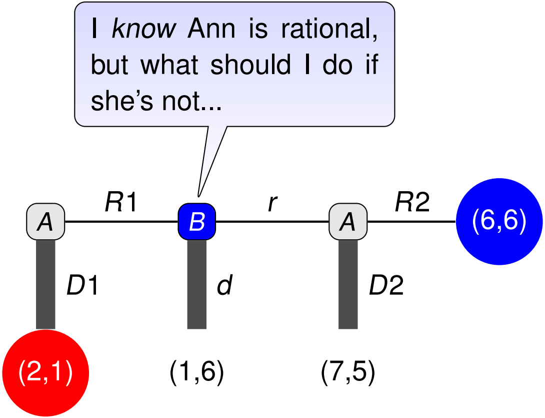
I *know* Ann is rational, but what should I do if she's not...
A
R1
B
r
A
R2
(6,6)
D1
d
D2
(2,1)
(1,6)
(7,5)

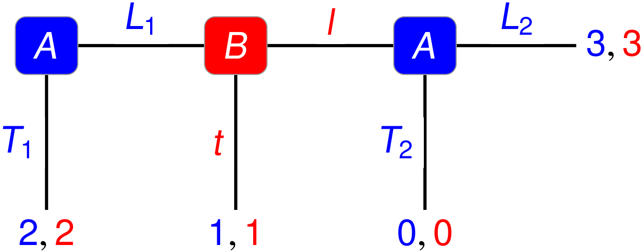
A
L1
B
l
A
L2
3, 3
T1
t
T2
2, 2
1, 1
0, 0

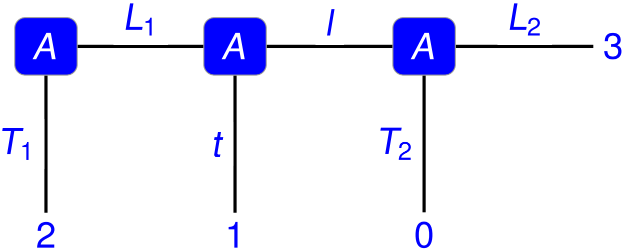
A
L_1
A
l
A
L_2
3
T_1
t
T_2
2
1
0

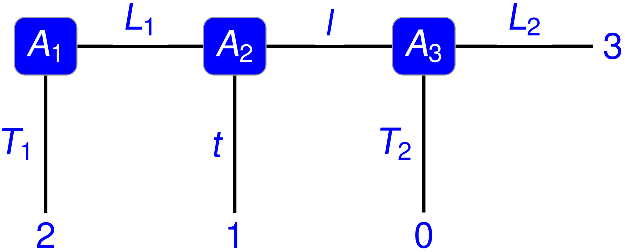

A1
L1
A2
l
A3
L2
3
T1
t
T2
2
1
0

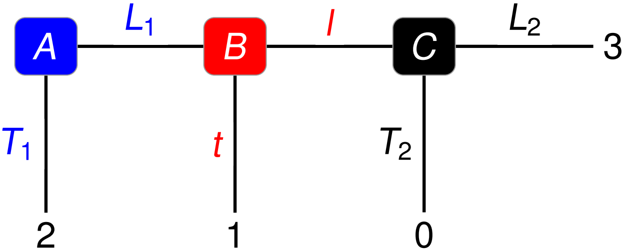

A
L_1
B
l
C
L_2
3
T_1
t
T_2
2
1
0

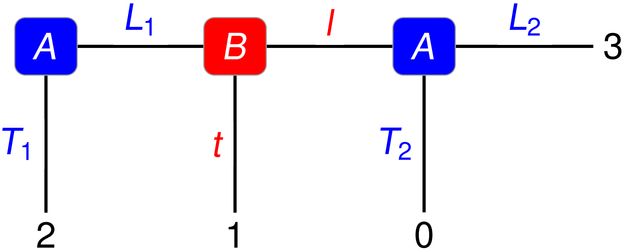

A
L_1
B
l
A
L_2
3
T_1
t
T_2
2
1
0

R. Aumann. *Backwards induction and common knowledge of rationality*. Games and Economic Behavior, 8, pgs. 6 - 19, 1995.

R. Stalnaker. *Knowledge, belief and counterfactual reasoning in games*. Economics and Philosophy, 12, pgs. 133 - 163, 1996.

J. Halpern. *Substantive Rationality and Backward Induction.* Games and Economic Behavior, 37, pp. 425-435, 1998.

Materially Rational: A player i is materially rational at a state w if every choice actually made is rational.

Substantively Rational: A player i is substantively rational at a state w if the player is materially rational and, in addition, for each *possible* choice, the player *would* have chosen rationally if she had had the opportunity to choose.

Materially Rational: A player i is materially rational at a state w if every choice actually made is rational.

Substantively Rational: A player i is substantively rational at a state w if the player is materially rational and, in addition, for each *possible* choice, the player *would* have chosen rationally if she had had the opportunity to choose.

E.g., Taking keys away from someone who is drunk.

		Bob	
		t	*l*
Ann	*T*	2,2	2,2
	LT	1,1	0,0
	LL	1,1	3,3

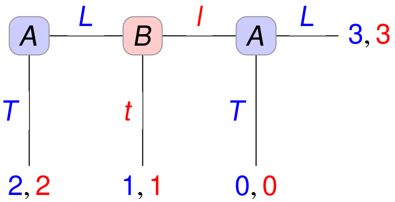

- The backward induction solution is (LL, l)
- Consider a model with a single possible world assigned the profile (TL, t).

		Bob	
		t	*l*
Ann	*T*	2,2	2,2
	LT	1,1	0,0
	LL	1,1	3,3

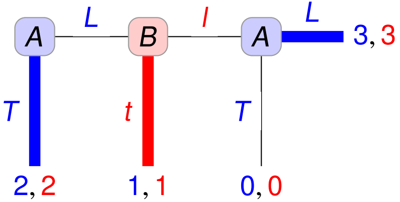

- *T* is a best response to *t*, so Ann is materially rational. She is also substantively rational. (Why?)
- Bob doesn't move, so Bob is materially rational. Is he substantively rational?

Ann \ Bob	*t*	*l*
T	2,2	2,2
LT	1,1	0,0
LL	1,1	3,3

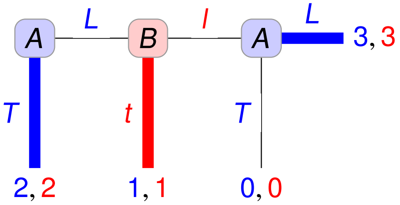

- Is Bob substantively rational? Would *t* be rational, *if* he had a chance to act?
- Suppose that Bob is disposed to revise his beliefs in such a way that if Ann acted irrationally once, she will act irrationally later in the game.

Ann \ Bob	*t*	*l*
T	2,2	2,2
LT	1,1	0,0
LL	1,1	3,3

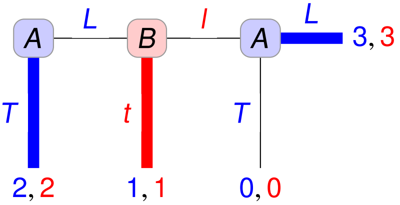

- Bob's belief in a causal counterfactual: Ann would choose *L* on her second move *if* she had a chance to move.
- But we need to ask what would Bob believe about Ann *if* he learned that he was wrong about her first choice. This is a question about Bob's belief revision policy.

Informal characterizations of BI

- Future choices are *epistemically independent* of any observed behavior
- Any “off-equilibrium” choice is interpreted simply as a mistake (which will not be repeated)
- At each choice point in a game, the players only reason about future paths

Rationalizing Observed Actions

After observing an (unexpected) move by some player, you could:

1. Change your belief about the player's rationality, but maintain your beliefs about the player's *passive beliefs*.
2. Change your belief about the player's passive beliefs, but maintain your belief in the player's rationality.
3. Conclude that the player perceives the game differently.

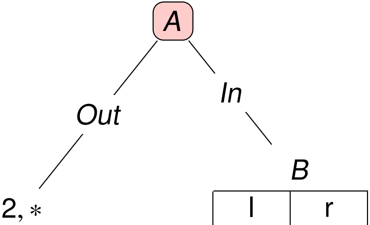

	l	r
u	4, ∗	0, ∗

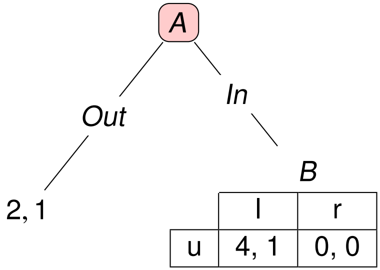
A
Out
In
2, 1
B
l
r
u
4, 1
0, 0

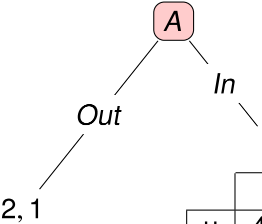

		B	
		l	r
A	u	4, 1	0, 0
	d	0, 0	1, 4

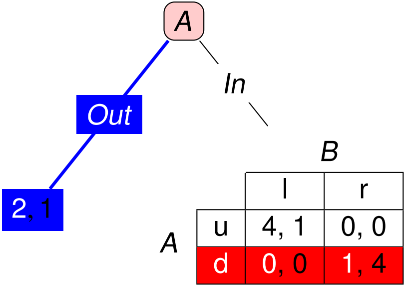
A
Out
In
2, 1
B
l
r
A
u
4, 1
0, 0
d
0, 0
1, 4

		B	
		l	r
A	*Out*	2, 1	2, 1
	u	4, 1	0, 0
	d	0, 0	1, 4

		B	
		l	r
	Out	2, 1	2, 1
A	u	4, 1	0, 0
	d	0, 0	1, 4

What is forward induction reasoning?

Forward Induction Principle: a player should use all information she acquired about her opponents' past behavior in order to improve her prediction of their future simultaneous and past (unobserved) behavior, relying on the assumption that they are rational.

P. Battigalli. *On Rationalizability in Extensive Games*. Journal of Economic Theory, 74, pgs. 40 - 61, 1997.

Issues

- The players' conditional beliefs must be *rich enough* to employ the forward induction principle.
- Do the players robustly believe the forward induction principle?
- Can players become more/less confident in the forward induction principle?

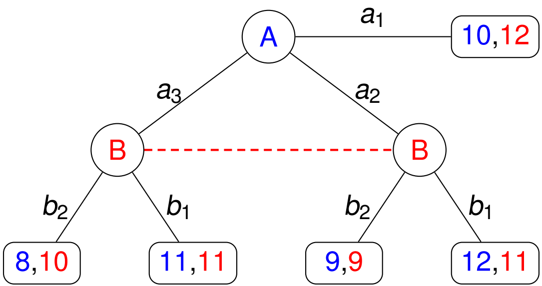

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	10,12	10,12
a_2	12,11	9,9
a_3	11,11	8,10

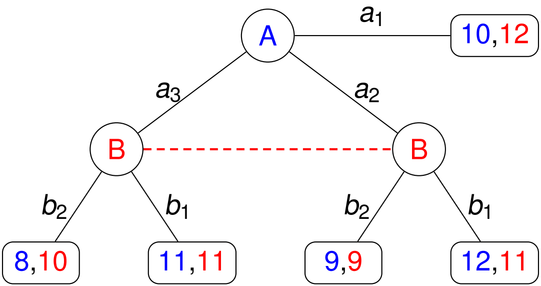

- No matter what Ann's probabilities are for playing a_2 and a_3, Bob is always better off playing b_1.
- Thus, Bob will play b_1 at his information set
- Knowing this, Ann will play a_2
- Dynamic deliberation will never lead to the "bad" equilibrium (a_1, b_2 if a_2 or a_3)

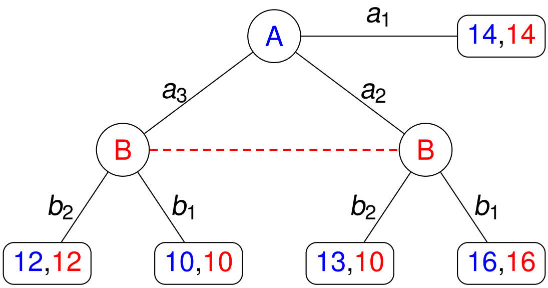

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	14,14	14,14
a_2	16,16	13,10
a_3	10,10	12,12

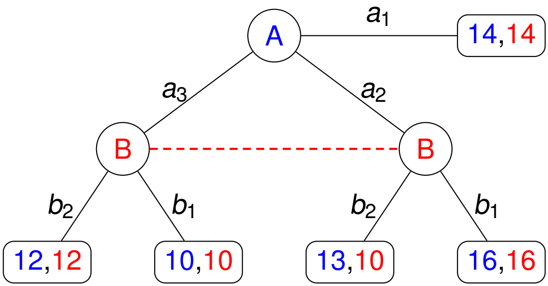

- If Ann plays a_2, Ann will get a better payoff than if Ann plays a_3 no matter what Bob does
- This will lead Bob to play b_1. Ann can figure *this* out, so she will play a_2.
- If we implement some sort of Bayes dynamics with a *tendency towards a better response*, then deliberation will lead to the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

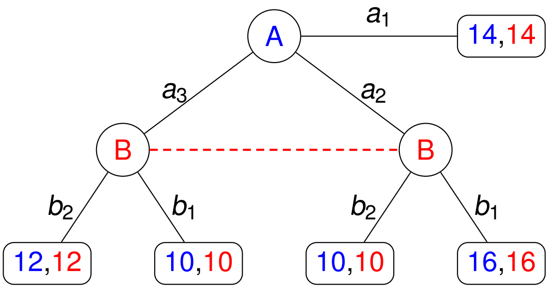

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	14,14	14,14
a_2	16,16	10,10
a_3	10,10	12,12

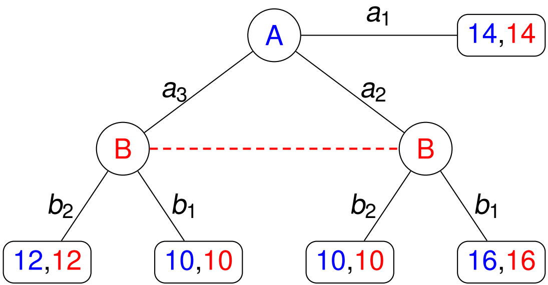

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- In the preceding example, $p_A(a_2 \mid a_2 \text{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

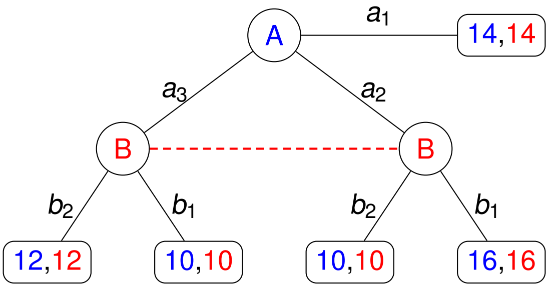

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- In this example, $p_A(a_2 \mid a_2 \text{ or } a_3)$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

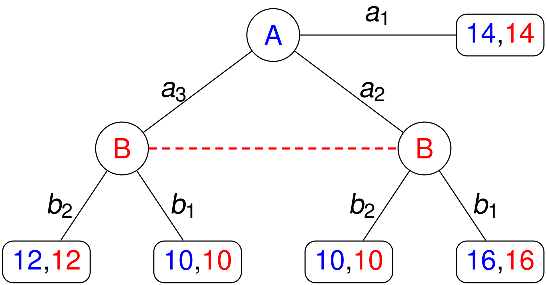

- a_1 gives Ann a higher payoff than a_3 no matter what Bob does
- Therefore, Bob should know that Ann will only play "a_2 or a_3" if she plays a_2.
- Accordingly Bob will play b_1 rather than b_2, and knowing this Ann will play a_2 rather than a_1
- Unless there is some "pre-deliberational" pruning,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

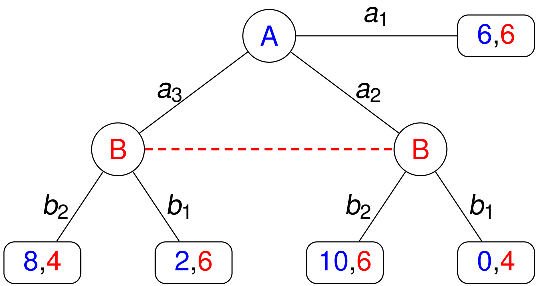

	b_1 if a_2 or a_3	b_2 if a_2 or a_3
a_1	6,6	6,6
a_2	0,4	10,6
a_3	2,6	8,4

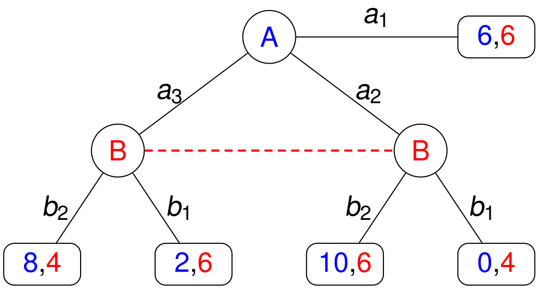

- $(a_2, b_2$ if $a2$ or $a_3)$ and $(a_1, (0.5b_1, 0.5b_2))$ are equilibrium
- a_1 is not *ratifiable*: if Bob is given a chance to move that means Ann must be expecting Bob to choose b_2. Ann's best response to this is a_2. Knowing this Bob will choose b_2
- Starting at $(1/3a_1, 1/3a_2, 1/3a_3)$ and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

Discussion

- Deliberation for logically omniscient agents
- Credences on own acts
- Backward induction vs. forward induction

Discussion: Logical Omniscience

"Any context where an agent engages in reasoning is a context that is distorted by the assumption of deductive omniscience, since reasoning (at least deductive reasoning) is an activity that deductively omniscient agents have no use for.

Discussion: Logical Omniscience

"Any context where an agent engages in reasoning is a context that is distorted by the assumption of deductive omniscience, since reasoning (at least deductive reasoning) is an activity that deductively omniscient agents have no use for. Deliberation, to the extent that it is thought of as a rational process of figuring out what one should do given one's priorities and expectations is an activity that is unnecessary for the deductively omniscient.

Discussion: Logical Omniscience

"Any context where an agent engages in reasoning is a context that is distorted by the assumption of deductive omniscience, since reasoning (at least deductive reasoning) is an activity that deductively omniscient agents have no use for. Deliberation, to the extent that it is thought of as a rational process of figuring out what one should do given one's priorities and expectations is an activity that is unnecessary for the deductively omniscient. In fact any kind of information processing or computation is unintelligible as an activity of a deductively omniscient agent. " (pp. 428,429)

R. Stalnaker. *The Problem of Logical Omniscience, I*. Sythese, 89:3, 1991, pp. 425 - 440.

Discussion: The Cost of Thinking

L. J. Savage. *Difficulties in the theory of personal probability*. Philosophy of Science, 34(4), pgs. 305 - 310, 1967.

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

Credences of own acts

"...the relevant distinction is between the *first-person* perspective of a practical deliberator and the *third-person* perspective of an observer. While the observer can predict what I will do, I can't, insofar as I deliberate upon what is to be done. Deliberating in this way is incompatible with predicting the outcome of deliberation. To put is shortly, *deliberation crowds out prediction.*" (pg. 91)

W. Rabinowicz. *Does Practical Deliberation Crowd Out Self-Prediction?*. Erkenntnis, 57, pgs. 91 = 122, 2002.

H. Gaifman. *Self-reference and the acyclicity of rational choice*. Annals of Pure and Applied Logic, 96, pgs. 117 - 140, 1999.

The Irrational Choice

Mr. Z offers Adam two boxes, each containing \$10. Adam can choose either $S1$: to take the leftmost box and get \$10, or $S2$: to take the two boxes and get \$20. Before making his decision, Adam is informed by Z. that if he acts irrationally, Z will give him a bonus of \$100. (...to eliminate noise factors, assume that Adam believes that Z. is serious, has the relevant knowledge, is a perfect reasoner and is completely trustworthy.)

“...the bonus condition in Z’s statement has truth-conditions, and once Adam has chosen it can be evaluated...It is only from the perspective of *Adam* qua *deliberating rational agent* that the bonus condition must be excluded as meaningless.”

The Rational Choice

Mr. Z offers Adam two boxes, each containing \$10. Adam can choose either $S1$: to take the leftmost box and get \$10, or $S2$: to take the two boxes and get \$20. Before making his decision, Adam is informed by Z. that if he acts **rationally**, Z will give him a bonus of \$100. (...to eliminate noise factors, assume that Adam believes that Z. is serious, has the relevant knowledge, is a perfect reasoner and is completely trustworthy.)

Backward *versus* Forward Induction

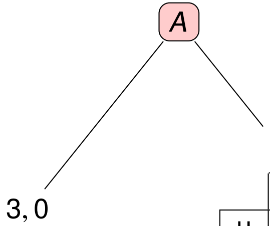

Bob

Ann	l	c	r
u	2, 2	2, 1	0, 0
d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

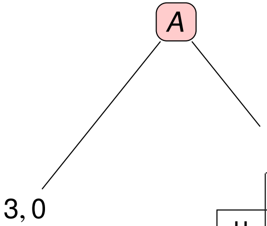

Bob

Ann	l	c	r
u	2, 2	2, 1	0, 0
d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

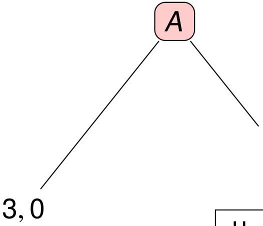

Bob

Ann		l	c	r
	u	2, 2	2, 1	0, 0
	d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

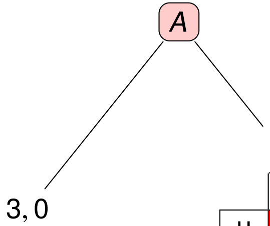

		Bob		
		l	c	r
Ann	u	2, 2	2, 1	0, 0
	d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

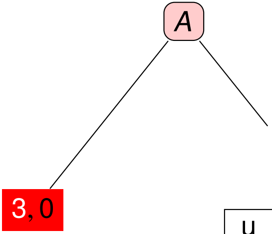

Bob

Ann	l	c	r
u	2, 2	2, 1	0, 0
d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

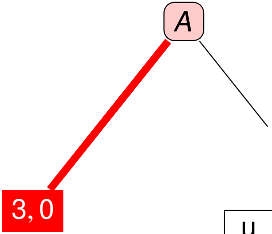

Bob

Ann		l	c	r
	u	2, 2	2, 1	0, 0
	d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Backward *versus* Forward Induction

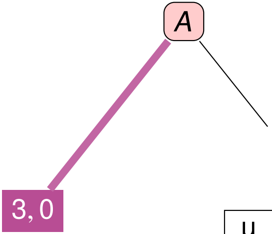

		Bob		
		l	c	r
Ann	u	2, 2	2, 1	0, 0
	d	1, 1	1, 2	4, 0

A. Perea. *Backward Induction* versus *Forward Induction Reasoning*. Games, 1, pgs. 168 - 188, 2010.

Reasoning about mistakes

A. Knoks and EP. *Reasoning about mistakes in games*. manuscript.

Reasoning about mistakes

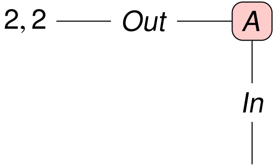

		B: l	B: r
A	u	4, 1	0, 0
A	d	0, 0	1, 4

Reasoning about mistakes

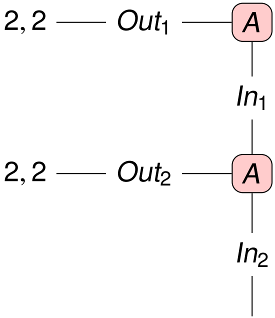

		B: l	B: r
A	u	4, 1	0, 0
A	d	0, 0	1, 4

Reasoning about mistakes

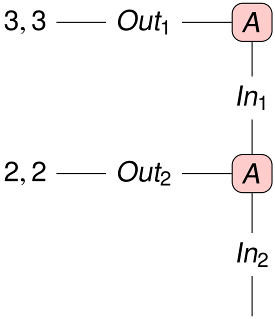

B

		l	r
A	u	4, 1	0, 0
	d	0, 0	1, 4

Thank you!