Epistemic Game Theory

Lecture 7

Eric Pacuit

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

March 24, 2014

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

Deliberational Decision Theory

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

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.

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.

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.

From Substantive to Procedural Rationality

"Discussions of substantive rationality take place in an essentially *static* framework. Thus, equilibrium is discussed without explicit reference to any dynamic process by means of which the equilibrium is achieved. Similarly, prior beliefs are taken as given, without reference to how a rational agent acquires these beliefs. Indeed, all questions of the procedure by means of which rational behavior is achieved are swept aside by a methodology that treats this procedure as completed and reifies the supposed limiting entities by categorizing them axiomatically." (pg. 180)

K. Binmore. *Modeling Rational Players: Part I*. Economics and Philosophy, 3, pgs. 179 - 214, 1987.

From Substantive to Procedural Rationality

What does it mean to choose "rationally"?

"A glance at any dictionary will confirm that economists, firmly entrenched in the static viewpoint described above, have hijacked this word and used it to mean something for which the word *consistent* would be more appropriate. " (pg. 181)

K. Binmore. *Modeling Rational Players: Part I*. Economics and Philosophy, 3, pgs. 179 - 214, 1987.

Decision Theory

Choose an act at time t that maximizes (causal, evidential) expected utility with respect to a probability Pr_t that characterizes your beliefs at time 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.

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

$$\underset{\text{initial model}}{\mathcal{M}_0} \Longrightarrow \mathcal{M}_1 \Longrightarrow \mathcal{M}_2 \Longrightarrow \cdots$$

Modeling Rational Deliberation

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

initial
model

Each $\mathcal{M}_i$ describes what the decision maker believes, *including beliefs about what they are going to do* (at the end of deliberation).

Modeling Rational Deliberation

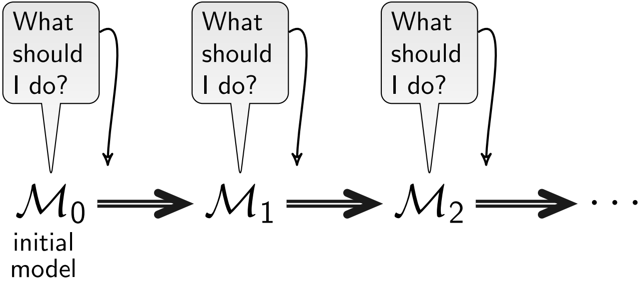

Dynamical rules transform the decision maker's beliefs, given her evaluation of the available acts.

Modeling Rational Deliberation

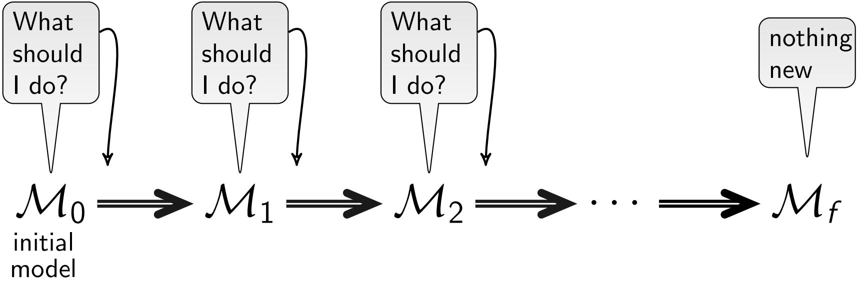

Deliberations stops when a “fixed-point” is reached.

Updated Probability on Acts

If $Pr_t(A) < 1$ and $x > y$, then
$Pr_t(A \mid EU_t(A) = x \;\; \& \;\; EU_t({\sim}A) = y) > Pr_t(A)$

Skyrms' Model of Deliberation

For each t, $\mathcal{M}_t$ consists of two components: a probability on the set of acts and a probability over the state of the world.

Skyrms' Model of Deliberation

For each t, $\mathcal{M}_t$ consists of two components: a probability on the set of acts and a probability over the state of the world. At each time t,

Step 1: The agent assesses the utilities of acts in light of her beliefs at time t about the state of the world.

Step 2: The agents alters her probabilities for acts and states in light of utilities using an *update rule* that "seeks the good"

Skyrms' Model of Deliberation

For each t, $\mathcal{M}_t$ consists of two components: a probability on the set of acts and a probability over the state of the world. At each time t,

Step 1: The agent assesses the utilities of acts in light of her beliefs at time t about the state of the world.

Step 2: The agents alters her probabilities for acts and states in light of utilities using an *update rule* that "seeks the good" by increasing probabilities of acts with utilities above the status quo, decreasing probabilities of acts with utilities below the status quo and leaving probabilities of acts equal to the status quo unchanged.

Information Feedback

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."

Deliberational Equilibrium

The decision maker cannot decide to do an act that is not an equilibrium of the deliberational process.

(*provided we neglect processing costs...the implementations use a "satisficing level"*)

Deliberational Equilibrium

The decision maker cannot decide to do an act that is not an equilibrium of the deliberational process.

(*provided we neglect processing costs...the implementations use a "satisficing level"*)

This sort of equilibirium requirement can be seen *as a consequence of the expected utility principle* (dynamic coherence).

It is usually neglected because the process of informational feedback is usually neglected.

Details of the Model

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

Details of the Model

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).

Details of the Model

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

A person's state of indecision evolves during deliberation.

After computing expected utility, she will believe more strongly that she will ultimately do the acts (or one of those acts) that are ranked more highly than her current state of indecision.

The decision maker follows a "simple dynamical rule" for "making up one's mind"

The Dynamical Rule “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"

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.

decision maker's personal state: $\langle x, y\rangle$ where x is the state of indecision and the probabilities she assigns to the "states of nature"

decision maker's personal state: $\langle x, y\rangle$ where x is the state of indecision and the probabilities she assigns to the "states of nature"

Dynamics: $\varphi(\langle x, y\rangle) = \langle x', y'\rangle$ consisting of

1. An "adaptive dynamic map" D sending $\langle x, y\rangle$ to x'
2. the informational feedback process I sending $\langle x, y\rangle$ to y'

decision maker's personal state: $\langle x, y\rangle$ where x is the state of indecision and the probabilities she assigns to the "states of nature"

Dynamics: $\varphi(\langle x, y\rangle) = \langle x', y'\rangle$ consisting of

1. An "adaptive dynamic map" D sending $\langle x, y\rangle$ to x'
2. the informational feedback process I sending $\langle x, y\rangle$ to y'

A personal state $\langle x, y\rangle$ is a **deliberational equilibrium** iff $\varphi(\langle x, y\rangle) = \langle x, y\rangle$

Nash Dynamics

covetability of act A: given a state of indecision $\mathbf{P}$

$cov(A) = \max(EU(A) - EU(\mathbf{P}), 0)$

Nash Dynamics

covetability of act A: given a state of indecision $\mathbf{P}$
$cov(A) = \max(EU(A) - EU(\mathbf{P}), 0)$

Nash map: $\mathbf{P} \mapsto \mathbf{P}'$ where each component p_i' is calculated as follows:

$$p_i' = \frac{p_i + cov(A_i)}{1 + \sum_i cov(A_i)}$$

Nash Dynamics

covetability of act A: given a state of indecision $\mathbf{P}$
$cov(A) = \max(EU(A) - EU(\mathbf{P}), 0)$

Nash map: $\mathbf{P} \mapsto \mathbf{P}'$ where each component p_i' is calculated as follows:

$$p_i' = \frac{p_i + cov(A_i)}{1 + \sum_i cov(A_i)}$$

More generally, for $k > 0$,

$$p_i' = \frac{k \cdot p_i + cov(A_i)}{k + \sum_i cov(A_i)}$$

where k is the "index of caution". The higher the k the more slowly the decision maker moves in the direction of acts that look more attractive than the status quo.

Fact. If D seeks the good and I is continuous, then there is a delbierational equilibrium, $\langle x, y\rangle$, for $\langle D, I\rangle$. If D' also seeks the good, then $\langle x, y\rangle$ is also a deliberational equilibrium for $\langle D', I\rangle$. The default mixed act corresponding to x maximizes expected utility at $\langle x, y\rangle$.

Tension with the Logical Omniscience Assumption?

"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.

Tension with the Logical Omniscience Assumption?

"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.

Tension with the Logical Omniscience Assumption?

“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.

Deliberation Crowds out Prediction?

No probabilities of acts!

F. Schick. *Self-Knowledge, Uncertainty and Choice*. The British Journal for the Philosophy of Science, 30:3, pgs. 235 - 252, 1979.

I. Levi. *Rationality, prediction and autonomous choice*. in *The Covenant of Reason*.

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

J. Joyce. *Levi on Predicting One's Own Actions*. Philosophical Studies, 110, pgs. 69 - 102, 2002.

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.*

Tracing Procedure

J.C. Harsanyi. *The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-Person Noncooperative Games*. International Journal of Game Theory, 4, pgs. 61 - 94, 1975.

J. C. Harsanyi and R. Selten. *A general theory of equilibrium selection in games*. The MIT Press, 1988.

Tracing Procedure

J.C. Harsanyi. *The Tracing Procedure: A Bayesian Approach to Defining a Solution for n-Person Noncooperative Games*. International Journal of Game Theory, 4, pgs. 61 - 94, 1975.

J. C. Harsanyi and R. Selten. *A general theory of equilibrium selection in games*. The MIT Press, 1988.

P. Jean-Jacques Herings. *Two simple proofs of the feasibility of the linear tracing procedure*. Economic Theory, 15, pgs. 485 - 490, 2000.

S. H. Schanuel, L.K. Simon, and W. R. Zame. *The algebraic geometry of games and the tracing procedure*. in *Game equilibrium models II: methods, morals and markets*, pgs. 9 - 43, Springer, 1991.

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

- Suppose there is a common prior that Ann will choose U with probability 0.5 and Bob will choose L with probability 0.5

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

	L	R
U	2,0.5	2,2
D	0.5,0.5	0.5,2

- Suppose there is a common prior that Ann will choose U with probability 0.5 and Bob will choose L with probability 0.5
- Consider the modified game (G^0) where the utilities are the *expected utilities* of the first game (G^1)

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

	L	R
U	2,0.5	2,2
D	0.5,0.5	0.5,2

- Suppose there is a common prior that Ann will choose U with probability 0.5 and Bob will choose L with probability 0.5
- Consider the modified game (G^0) where the utilities are the *expected utilities* of the first game (G^1)
- This game as a unique Nash equilibrium

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

	L	R
U	2,0.5	2,2
D	0.5,0.5	0.5,2

- For each $t \in [0,1]$, the game G^t is defined so that the payoffs of $u_i^t(x,y) = t \cdot u_i^1(x,y) + (1-t) \cdot u_i^0(x,y)$

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

	L	R
U	2,0.5	2,2
D	0.5,0.5	0.5,2

- For each $t \in [0,1]$, the game G^t is defined so that the payoffs of $u_i^t(x,y) = t \cdot u_i^1(x,y) + (1-t) \cdot u_i^0(x,y)$
- A graph of the equilibrium points as t varies from 0 to 1 will show a connected path from equilibria in G^0 to equilibria in G^1

The Tracing Procedure

	L	R
U	4,1	0,0
D	0,0	1,4

	L	R
U	2,0.5	2,2
D	0.5,0.5	0.5,2

- For each $t \in [0,1]$, the game G^t is defined so that the payoffs of $u_i^t(x,y) = t \cdot u_i^1(x,y) + (1-t) \cdot u_i^0(x,y)$
- A graph of the equilibrium points as t varies from 0 to 1 will show a connected path from equilibria in G^0 to equilibria in G^1
- This process almost always leads to a unique equilibrium in G^1 (modifying the payoffs with a a logarithmic term guarantees uniqueness)

Harsanyi and Selten also propose a method to *calculate* a prior given the game.

Harsanyi and Selten also propose a method to *calculate* a prior given the game.

The picture is of deliberators who are “computationally adept but, initially at least, strategically naive”

Harsanyi and Selten also propose a method to *calculate* a prior given the game.

The picture is of deliberators who are “computationally adept but, initially at least, strategically naive”

They can identify game-theoretic equilibria instantaneously.

Harsanyi and Selten also propose a method to *calculate* a prior given the game.

The picture is of deliberators who are "computationally adept but, initially at least, strategically naive"

They can identify game-theoretic equilibria instantaneously.

At $t = 0$ each contemplates jumping to the conclusion that the act the with maximum expected utility according to the common prior is the correct one.

Harsanyi and Selten also propose a method to *calculate* a prior given the game.

The picture is of deliberators who are "computationally adept but, initially at least, strategically naive"

They can identify game-theoretic equilibria instantaneously.

At $t = 0$ each contemplates jumping to the conclusion that the act the with maximum expected utility according to the common prior is the correct one.

At later times, the hypothesis that the other players will make their best response gets stronger and stronger, until at $t = 1$ only an equilibrium point of the original game remains.

Skyrms' model of deliberation in games.

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*).

Games played by Bayesian deliberators

Under suitable conditions of common knowledge, a joint deliberational equilibrium on the part of all players corresponds to a Nash equilibrium point of the game.

Games played by Bayesian deliberators

Under suitable conditions of common knowledge, a joint deliberational equilibrium on the part of all players corresponds to a Nash equilibrium point of the game.

Strengthening the assumptions slightly leads in a natural way to refinements of the Nash equilibrium.

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

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

$$\begin{aligned} EU(U) &= 0.8 \\ EU(D) &= 0.6 \\ EU(L) &= 0.2 \\ EU(R) &= 1.6 \\ SQ_A &= 0.64 \\ SQ_B &= 1.04 \end{aligned}$$

$$\begin{aligned} COV(U) &= \max(0.8 - 0.64, 0) = 0.16 \\ COV(D) &= \max(0.6 - 0.64, 0) = 0 \\ COV(L) &= \max(0.28 - 1.04, 0) = 0 \\ COV(R) &= \max(1.6 - 1.04, 0) = 0.56 \end{aligned}$$

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

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

$$\begin{array}{ll}
EU(U) = 0.8 & COV(U) = \max(0.8 - 0.64, 0) = 0.16 \\
EU(D) = 0.6 & COV(D) = \max(0.6 - 0.64, 0) = 0 \\
EU(L) = 0.2 & COV(L) = \max(0.28 - 1.04, 0) = 0 \\
EU(R) = 1.6 & COV(R) = \max(1.6 - 1.04, 0) = 0.56 \\
p_U = \frac{k\cdot 0.2 + 0.16}{k + 0.16} & \\
p_L = \frac{k\cdot 0.4 + 0}{k + 0.56} &
\end{array}$$

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

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

$$EU(U) = 0.8$$
$$EU(D) = 0.6$$
$$EU(L) = 0.2$$
$$EU(R) = 1.6$$
$$p_U = \frac{10 \cdot 0.2 + 0.16}{10 + 0.16} = 0.212598$$
$$p_L = \frac{k \cdot 0.4 + 0}{k + 0.56} = 0.378788$$

$$COV(U) = \max(0.8 - 0.64, 0) = 0.16$$
$$COV(D) = \max(0.6 - 0.64, 0) = 0$$
$$COV(L) = \max(0.28 - 1.04, 0) = 0$$
$$COV(R) = \max(1.6 - 1.04, 0) = 0.56$$

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

$\mathbf{P}_A = \langle 0.212598, 0.787402 \rangle$ and $\mathbf{P}_B = \langle 0.378788, 0.621212 \rangle$

$$EU(U) = 0.38 \cdot 2 + 0.62 \cdot 0 = 0.8$$
$$EU(D) = 0.38 \cdot 0 + 0.62 \cdot 1 = 0.6$$
$$EU(L) = 0.21 \cdot 1 + 0.78 \cdot 0 = 0.2$$
$$EU(R) = 0.21 \cdot 0 + 0.78 \cdot 2 = 1.6$$
$$SQ_A = 0.21 \cdot EU(U) + 0.78 \cdot EU(D)$$
$$SQ_B = 0.37 \cdot EU(L) + 0.62 \cdot EU(R)$$

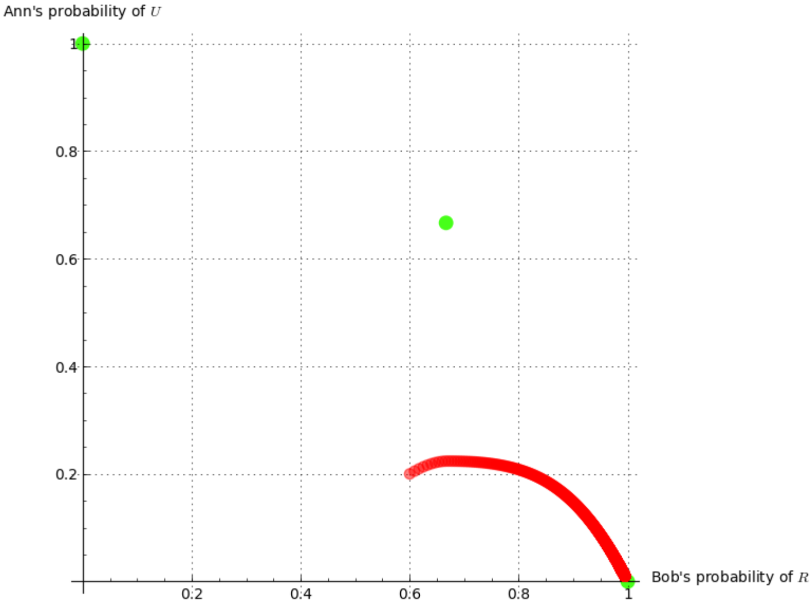
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

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

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

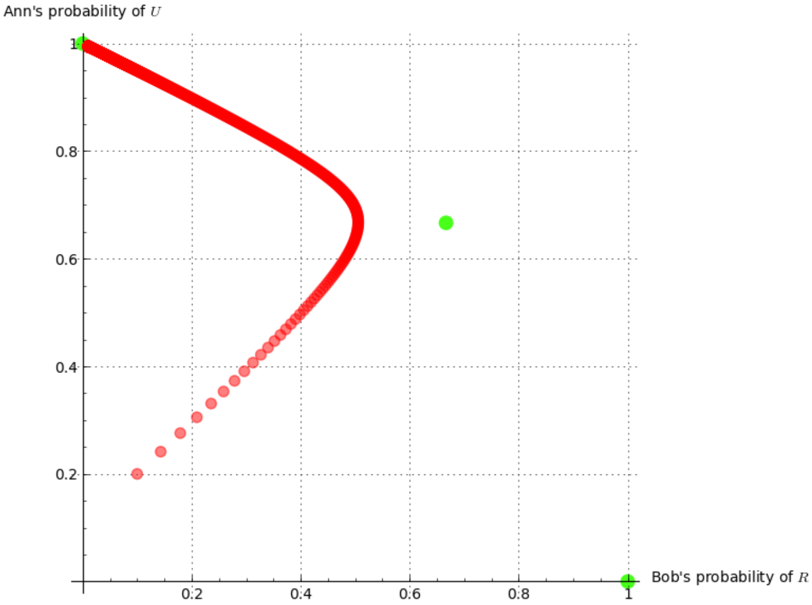

Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

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

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

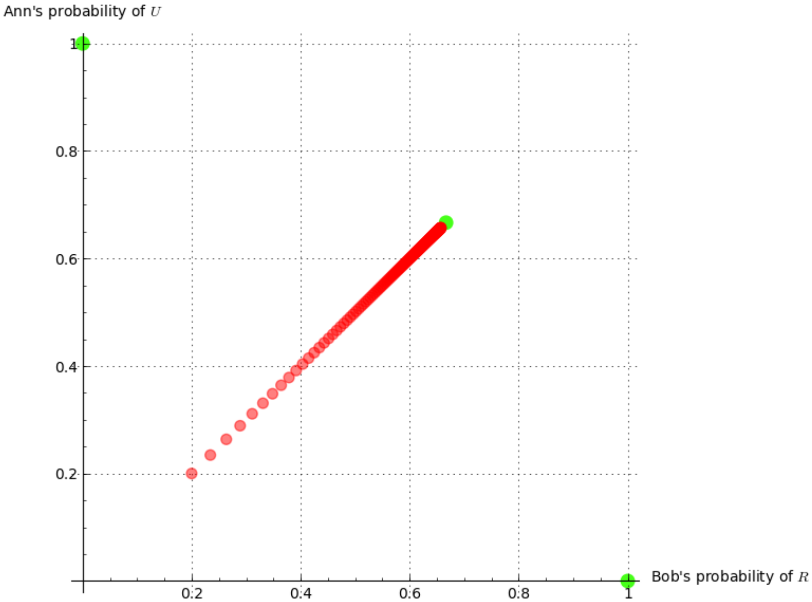
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Bayes Dynamics

If the new information that a player gets by emulating other players' calculations, updating his probabilities on their actions, and recalculating his expected utilities is e, then his new probabilities that he will do act A should be:

$$p_2(A) = p_1(A \mid e) = p_1(A) \cdot \frac{p_1(e \mid A)}{\sum_i p_1(A_i) \cdot p_1(e \mid A_i)}$$

where $\{A_i\}$ is a partition on the alternative acts.

Bayes Dynamics

If the new information that a player gets by emulating other players' calculations, updating his probabilities on their actions, and recalculating his expected utilities is e, then his new probabilities that he will do act A should be:

$$p_2(A) = p_1(A \mid e) = p_1(A) \cdot \frac{p_1(e \mid A)}{\sum_i p_1(A_i) \cdot p_1(e \mid A_i)}$$

where $\{A_i\}$ is a partition on the alternative acts.

But our deliberators do not have the appropriate proposition e in a large probability space that defines the likelihoods $p(e \mid A)$.

Is Nash a Bayes dynamics?

- If a deliberator starts with probability 1 that she will do some act that has utility less than the status quo, Nash will pull that probability down and raise the zero probabilities of competing acts.

Is Nash a Bayes dynamics?

- If a deliberator starts with probability 1 that she will do some act that has utility less than the status quo, Nash will pull that probability down and raise the zero probabilities of competing acts.

  "Indeed, one can argue that if a deliberator is absolutely sure which act he is going to do he needn't deliberate, and if he is absolutely sure he won't do an act, then his deliberation should ignore that act. "

Is Nash a Bayes dynamics?

- If a deliberator starts with probability 1 that she will do some act that has utility less than the status quo, Nash will pull that probability down and raise the zero probabilities of competing acts.

  "Indeed, one can argue that if a deliberator is absolutely sure which act he is going to do he needn't deliberate, and if he is absolutely sure he won't do an act, then his deliberation should ignore that act. "

- If two acts have expected utility less that the status quo, then they both get covetability 0, even if their expected utilities are quite different.

Pumping up the probabilities of acts that look best *and* the ratios of the probabilities of second to their place, etc. By Bayes Theorem,

$$\frac{p_2(A)}{p_2(B)} = \frac{p_1(A)}{p_1(B)} \cdot \frac{p_1(e \mid A)}{p_1(e \mid B)}$$

Tendency toward better response

Assume that the decision makers likelihoods are an increasing function of the newly calculated expected utilities.

Tendency toward better response

Assume that the decision makers likelihoods are an increasing function of the newly calculated expected utilities.

Darwin flow:

$$p_2(A) = p_1(A) + k \cdot p_1(A) \cdot \frac{EU(A) - EU(SQ)}{EU(SQ)}$$

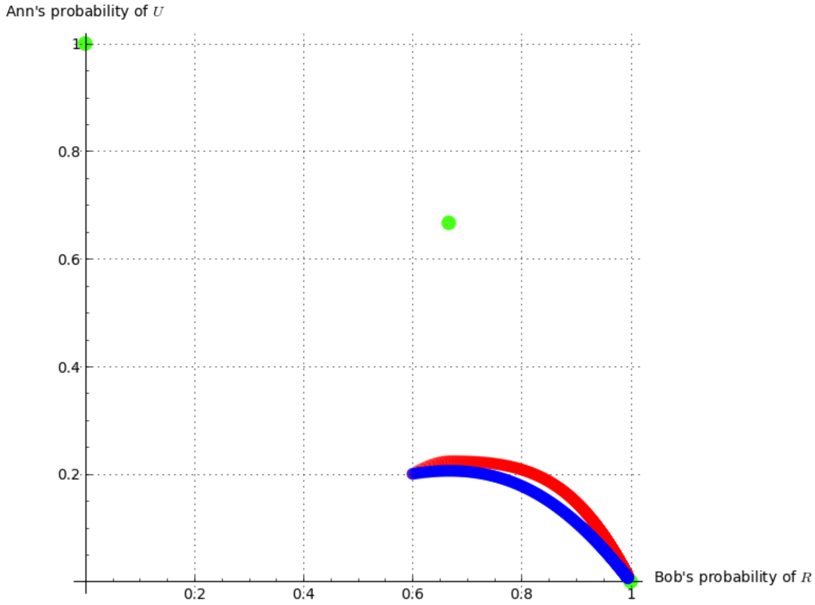

Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

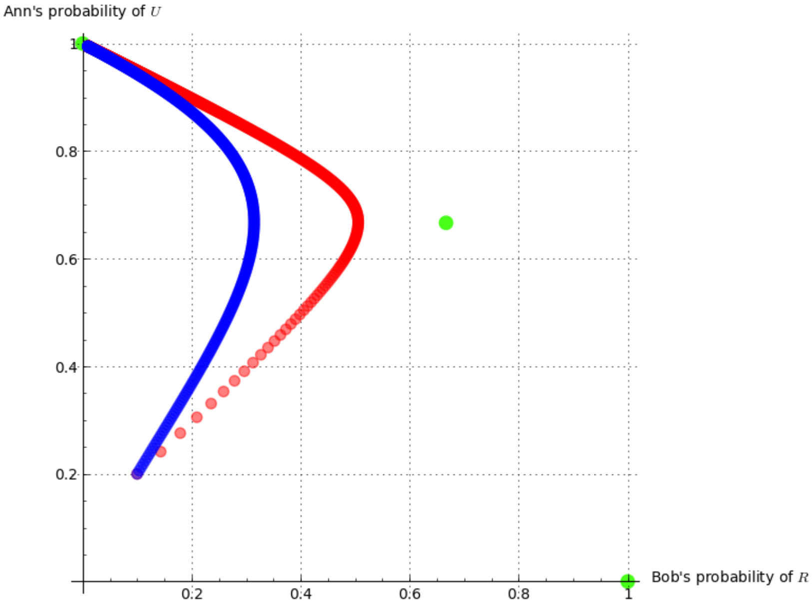
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

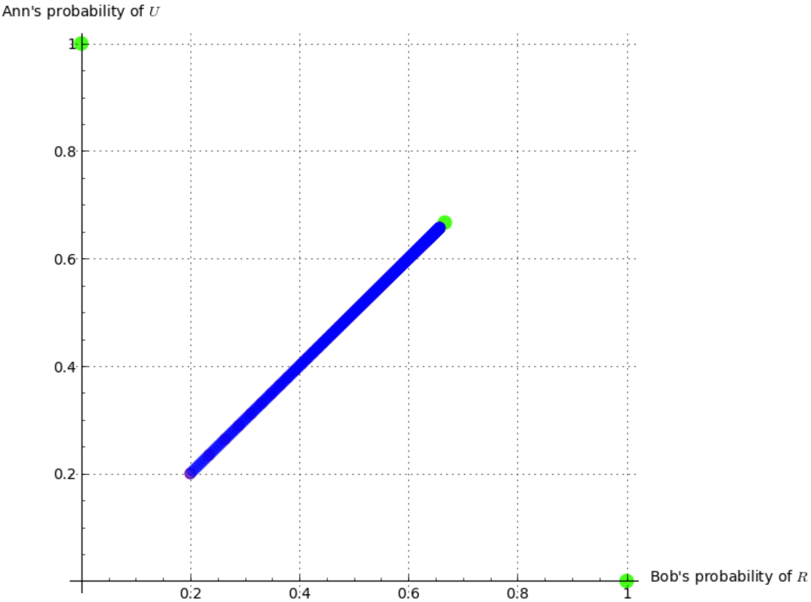
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Bob

		L	R
Ann	U	3,1	0,0
	D	0,0	1,3

$\mathbf{P}_A = \langle 0.05, 0.95 \rangle$ and $\mathbf{P}_B = \langle 0, 1 \rangle$

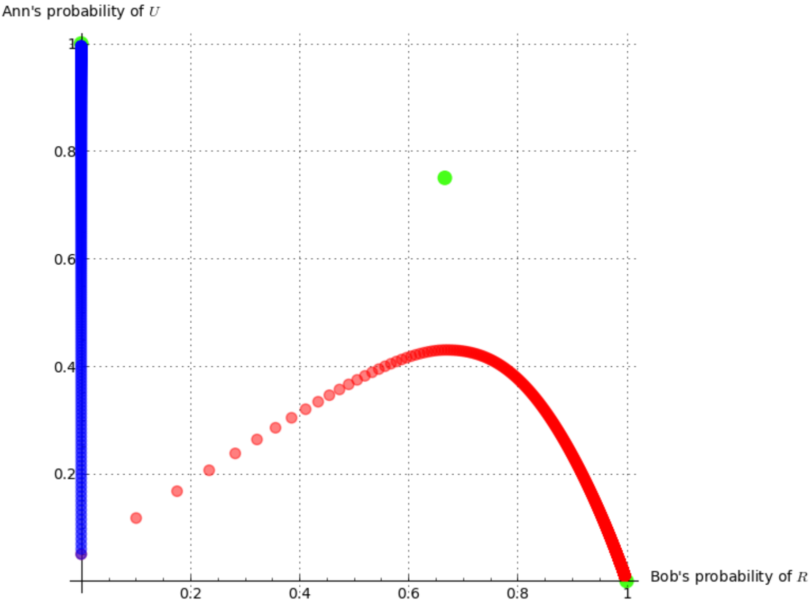
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

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

$$\mathbf{P}_A = \langle 0.65, 0.35 \rangle \text{ and } \mathbf{P}_B = \langle 0.55, 0.45 \rangle$$

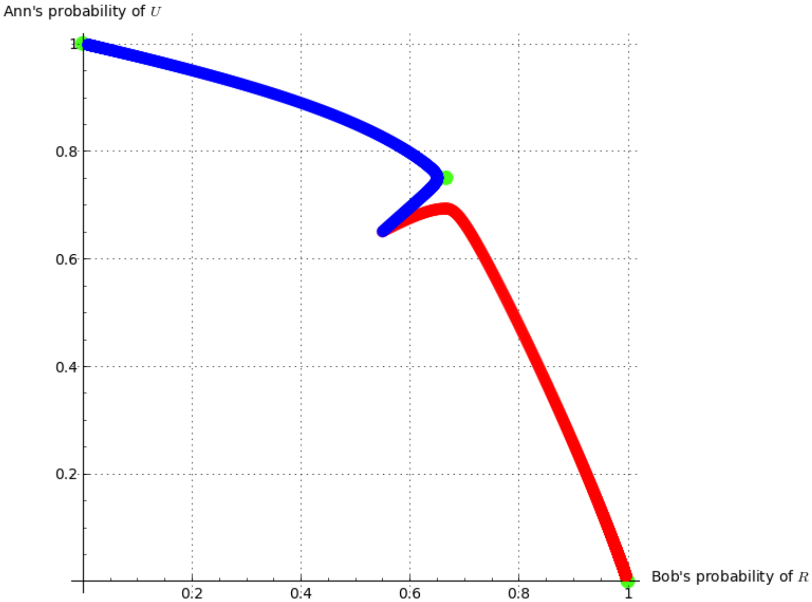
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

Normal form vs. Extensive Form

Explicitly modeling deliberation transforms a *single* choice into a situation of *sequential* choice.

Refinements for Nash equilibrium

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

Refinements for Nash equilibrium

Bob

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

If Bayesian deliberation must start in the interior of the space of indecision, then *dynamic deliberation* cannot lead to U, R.

Refinements for Nash equilibrium

Bob

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

If Bayesian deliberation must start in the interior of the space of indecision, then *dynamic deliberation* cannot lead to U, R.

Call an equilibrium *accessible* provided one can converge to it starting at a completely mixed state of indecision.

Refinements for Nash equilibrium

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

If Bayesian deliberation must start in the interior of the space of indecision, then *dynamic deliberation* cannot lead to U, R.

Call an equilibrium *accessible* provided one can converge to it starting at a completely mixed state of indecision.

Does *accessibility* correspond to perfect/proper equilibria?

Proper equilibrium

		Bob		
		L	C	R
Ann	U	-9,-9	-4,-4	-4,-4
	M	0,0	0,0	-4,-4
	D	1,1	0,0	-9,-9

Proper equilibrium

Bob

Ann	L	C	R
U	-9,-9	-4,-4	-4,-4
M	0,0	0,0	-4,-4
D	1,1	0,0	-9,-9

Proper equilibrium

		Bob L	C	R
Ann	U	-9,-9	-4,-4	-4,-4
	M	0,0	0,0	-4,-4
	D	1,1	0,0	-9,-9

ϵ-**proper equilibrium**: a completely mixed strategy profile such that if strategy s is a better response than s', then $\frac{p(s)}{p(s')} < \epsilon$

Proper equilibrium: the mixed strategy profile that is the limit as ϵ goes to 0 of ϵ-proper equilibria.

Proper equilibrium

		Bob L	C	R
Ann	U	-9,-9	-4,-4	-4,-4
	M	0,0	0,0	-4,-4
	D	1,1	0,0	-9,-9

ϵ-**proper equilibrium**: a completely mixed strategy profile such that if strategy s is a better response than s', then $\frac{p(s)}{p(s')} < \epsilon$

Proper equilibrium: the mixed strategy profile that is the limit as ϵ goes to 0 of ϵ-proper equilibria.

Proper equilibrium

Ann \ Bob	L	C	R
U	-9,-9	-4,-4	-4,-4
M	0,0	0,0	-4,-4
D	1,1	0,0	-9,-9

Starting at $\mathbf{P}_A = \langle 0.01, 0.5, 0.49 \rangle$ and $\mathbf{P}_B = \langle 0.01, 0.5, 0.49 \rangle$, Nash dynamics converges to (M, C), but Darwin converges to (D, L)

		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.

Samuelson identified adaptive rules that correspond to proper/perfect equilibrium. A key feature is:

ordinality: the velocity of probability change of a strategy depends only on the ordinal ranking among strategies according to their expected utilities.

L. Samuelson. *Evolutionary foundations for solution concepts for finite, two-player, normal-form games*. Proceedings of TARK, 1988.

Normal form vs. Extensive form

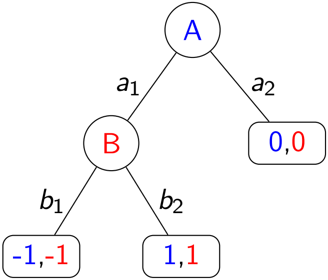

Normal form vs. Extensive form

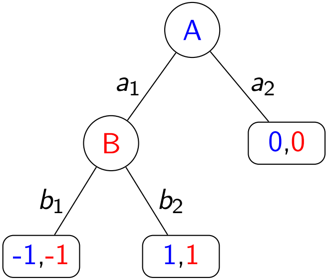

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

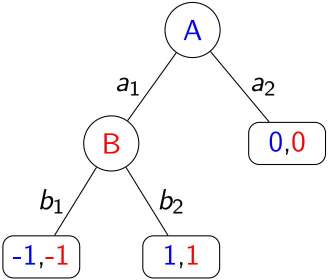

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

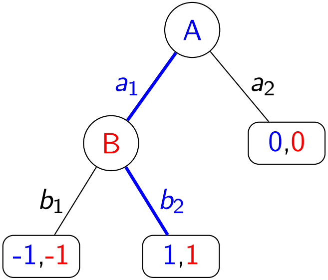

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

Normal form vs. Extensive form

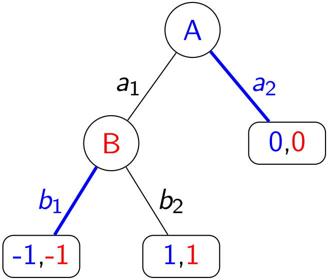

	b_1 if a_1	b_2 if a_1
a_1	-1,-1	1,1
a_2	0,0	0,0

(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

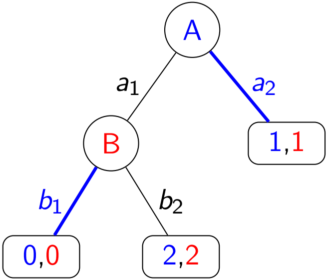

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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

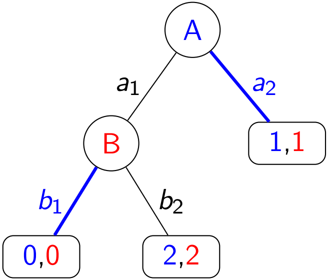

	b_1 if a_1	b_2 if a_1
a_1	0,0	2,2
a_2	1,1	1,1

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

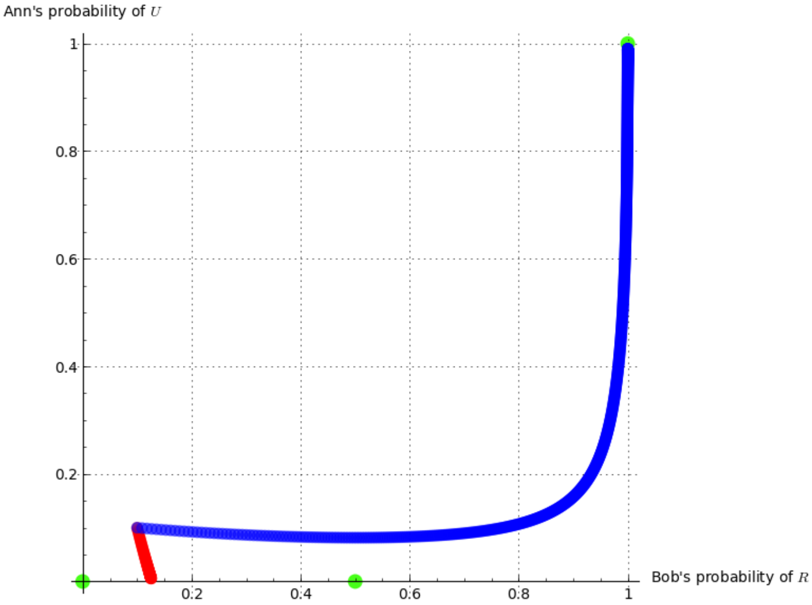
Ann's probability of U
1
0.8
0.6
0.4
0.2
0.2
0.4
0.6
0.8
1
Bob's probability of R

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.

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 $\nu_i = \{i_1, \dots i_j\}$ 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 $t+1$).

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.

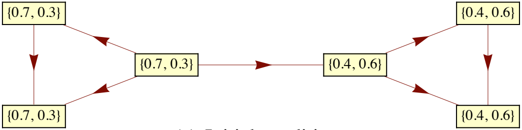

(a) Initial conditions

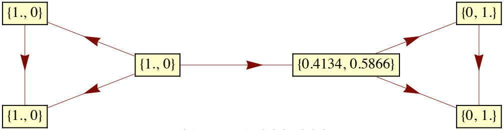

(b) $t = 1,000,000$

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}).