

Reasoning in Games: Players as Programs


Lecture 3

Eric Pacuit

Department of Philosophy

University of Maryland

pacuit.org



Plan

- ✓ **Monday** Epistemic utility theory, Decision- and game-theoretic background: Nash equilibrium
- ✓ **Tuesday** Introduction to game theory: rationalizability, epistemic game theory, introduction to backward induction
- Wednesday** backward and forward induction, Iterated games and learning, Skyrms's model of rational deliberation;
- Thursday** Introduction to webppl; Game-theoretic reasoning in webppl; Coordination games (comparing Skyrms's model of deliberation and the webppl approach)
- Friday** Models of game-theoretic reasoning

Yesterday

- ▶ Criticisms of the Nash program
- ▶ rationalizability and epistemic game theory
- ▶ Backward induction

“...no, equilibrium is not the way to look at games. Now, Nash equilibrium is king in game theory. Absolutely king. We say: No, Nash equilibrium is an interesting concept, and it's an important concept, but it's not the most basic concept. The most basic concept should be: to maximise your utility given your information. It's in a game just like in any other situation. Maximise your utility given your information!”

Robert Aumann, 5 Questions on Epistemic Logic, 2010

Theorem (Bernheim; Pearce; Brandenburger and Dekel; ...). (B_1, B_2, \dots, B_n) is a BRS for G iff there exists a model $\mathcal{M}^G = \langle W, (P_i)_{i \in N}, \mathbf{s} \rangle$ such that for all $i \in N$, $\text{Rat}_i = W$ and $[B_1 \times \dots \times B_n] = W$.

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,2	4,1
	<i>D</i>	1,4	3,3

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,1	1,0
	<i>D</i>	1,0	0,1

Game 2

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,2	4,1
	<i>D</i>	1,4	3,3

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* strictly dominates *D* and *L* strictly dominates *R*.

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,2	4,1
	<i>D</i>	1,4	3,3

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* strictly dominates *D* and *L* strictly dominates *R*.

Game 2: *U* strictly dominates *D*

		Bob	
		L	R
Ann	U	2,2	4,1
	D	1,4	3,3

Game 1

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

Game 2

Game 1: U strictly dominates D and L strictly dominates R .

Game 2: U strictly dominates D , and *after removing* D , L strictly dominates R .

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,2	4,1
	<i>D</i>	1,4	3,3

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	2,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* strictly dominates *D* and *L* strictly dominates *R*.

Game 2: *U* strictly dominates *D*, and *after removing D*, *L* strictly dominates *R*.

Theorem. In all models where the players are *rational* and there is *common belief of rationality*, the players choose strategies that survive iterative removal of strictly dominated strategies (and, conversely...).

Comparing Dominance Reasoning and MEU

$$G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$$

$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

Comparing Dominance Reasoning and MEU

$$G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$$

$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

$s, s' \in S_i$, s **strictly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, u_i(s, s_{-i}) > u_i(s', s_{-i})$$

Comparing Dominance Reasoning and MEU

$$G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$$

$X \subseteq S_{-i}$ (a set of strategy profiles for all players except i)

$s, s' \in S_i$, s **strictly dominates** s' with respect to X provided

$$\forall s_{-i} \in X, u_i(s, s_{-i}) > u_i(s', s_{-i})$$

$p \in \Delta(X)$, s is a **best response** to p with respect to X provided

$$\forall s' \in S_i, EU(s, p) \geq EU(s', p)$$

		L	Bob	R
Ann	U	$5, *$	$1, *$	
	M	$1, *$	$5, *$	
	D	$2, *$	$2, *$	

D is strictly dominated by $(0.5U, 0.5M)$.

Strict Dominance and MEU

Proposition. Suppose that $G = \langle N, (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ is a strategic game and $X \subseteq S_{-i}$. A strategy $s_i \in S_i$ is strictly dominated (possibly by a mixed strategy) with respect to X iff there is no probability measure $p \in \Delta(X)$ such that s_i is a best response to p .

Let $P \in \Delta(X)$ be a probability measure, the **support** of P is $\text{supp}(P) = \{x \in X \mid P(x) > 0\}$.

A probability measure $P \in \Delta(X)$ is said to be a **full support** probability measure on X provided $\text{supp}(P) = X$.

Strategic Reasoning and Admissibility

“The argument for deletion of a weakly dominated strategy for player i is that he contemplates the possibility that every strategy combination of his rivals occurs with positive probability. However, this hypothesis clashes with the logic of iterated deletion, which assumes, precisely, that eliminated strategies are not expected to occur.”

Mas-Colell, Whinston and Green. *Introduction to Microeconomics*. 1995.

A Puzzle

R. Cubitt and R. Sugden. *Rationally Justifiable Play and the Theory of Non-cooperative games*. Economic Journal, 104, pgs. 798 - 803, 1994.

R. Cubitt and R. Sugden. *Common reasoning in games: A Lewisian analysis of common knowledge of rationality*. Manuscript, 2011.

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	0,0
	<i>D</i>	0,0	0,0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	1,0
	<i>D</i>	1,0	0,1

Game 2

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	0,0
	<i>D</i>	0,0	0,0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* weakly dominates *D* and *L* weakly dominates *R*.

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	0,0
	<i>D</i>	0,0	0,0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* weakly dominates *D* and *L* weakly dominates *R*.

Game 2: *U* weakly dominates *D*

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	0,0
	<i>D</i>	0,0	0,0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* weakly dominates *D* and *L* weakly dominates *R*.

Game 2: *U* weakly dominates *D*, and *after removing D*, *L* strictly dominates *R*.

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	0,0
	<i>D</i>	0,0	0,0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1,1	1,0
	<i>D</i>	1,0	0,1

Game 2

Game 1: *U* weakly dominates *D* and *L* weakly dominates *R*.

Game 2: But, now what is the reason for not playing *D*?

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1, 1	0, 0
	<i>D</i>	0, 0	0, 0

Game 1

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	1, 1	1, 0
	<i>D</i>	1, 0	0, 1

Game 2

Game 1: *U* weakly dominates *D* and *L* weakly dominates *R*.

Game 2: But, now what is the reason for not playing *D*?

Theorem (Samuelson). There is no model of Game 2 satisfying common knowledge of rationality (where rationality incorporates weak dominance).

Both Including and Excluding a Strategy

Returning to the problem of weakly dominated strategies and rationalizability, one solution is to assume that players consider some strategies *infinitely more likely than other strategies*.

Both Including and Excluding a Strategy

Returning to the problem of weakly dominated strategies and rationalizability, one solution is to assume that players consider some strategies *infinitely more likely than other strategies*.

		Bob		
		1	[1]	
Ann		<i>L</i>	<i>R</i>	
		<i>U</i>	3,3	1,1
		<i>D</i>	2,2	2,2

L. Blume, A. Brandenburger, E. Dekel. *Lexicographic probabilities and choice under uncertainty*. *Econometrica*, 59(1), pgs. 61 - 79, 1991.

Backward and forward induction

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.

Theorem (Aumann) In any model, if there is common knowledge that the players are substantively rational at state w , the the backward induction solution is played at w .

Two propositions φ and ψ are epistemically independent for player i in world w iff $P_{i,w}(\varphi | \psi) = P_{i,w}(\varphi | \neg\psi)$ and $P_{i,w}(\psi | \varphi) = P_{i,w}(\psi | \neg\varphi)$

A possible belief revision policy: Information about different players should be epistemically independent.

Theorem (Stalnaker's interpretation of Aumann's theorem) Let G be a game of perfect information in agent form (i.e., players only move once) in which for each player different outcomes have different payoffs. Let \mathcal{M} be a model for G in which it is common belief that all agents are perfectly rational, and that all agents adopt belief revision policies that treat information about different agents as epistemically independent. Then in \mathcal{M} , the subgame perfect equilibrium strategy profile is realized.

Theorem (Stalnaker's interpretation of Aumann's theorem) Let G be a game of perfect information in **agent form (i.e., players only move once)** in which for each player different outcomes have different payoffs. Let \mathcal{M} be a model for G in which it is common belief that all agents are perfectly rational, and that all agents adopt belief revision policies that treat information about different agents as epistemically independent. Then in \mathcal{M} , the subgame perfect equilibrium strategy profile is realized.

Theorem (Stalnaker's interpretation of Aumann's theorem) Let G be a game of perfect information in agent form (i.e., players only move once) in which for each player different outcomes have different payoffs. Let \mathcal{M} be a model for G in which it is common belief that all agents are perfectly rational, and that **all agents adopt belief revision policies that treat information about different agents as epistemically independent**. Then in \mathcal{M} , the subgame perfect equilibrium strategy profile is realized.

Theorem (Stalnaker's interpretation of Aumann's theorem) Let G be a game of perfect information in agent form (i.e., players only move once) in which for each player different outcomes have different payoffs. Let \mathcal{M} be a model for G in which it is common belief that all agents are perfectly rational, and that all agents adopt belief revision policies that treat information about different agents as epistemically independent. Then in \mathcal{M} , the **subgame perfect equilibrium** strategy profile is realized.

1. Ann cheats — she has seen her opponent's cards.
2. Ann has a losing hand, since I have seen both her hand and her opponent's.
3. Ann is rational.

So, I conclude that she will not bet. But how should I revise my beliefs if I learn that Ann did bet?

1. Ann cheats — she has seen her opponent's cards.
2. Ann has a losing hand, since I have seen both her hand and her opponent's.
3. Ann is rational.

So, I conclude that she will not bet. But how should I revise my beliefs if I learn that Ann did bet?

It may be perfectly reasonable for me to be disposed to give up 2.


1. Ann cheats — she has seen her opponent's cards.
2. Ann has a losing hand, since I have seen both her hand and her opponent's.
3. Ann is rational.

So, I conclude that she will not bet. But how should I revise my beliefs if I learn that Ann did bet?

It may be perfectly reasonable for me to be disposed to give up 2.


I believe that (1) If Ann *were* to bet, she would lose (since she has a losing hand) and (2) If I were to *learn* that she *did* bet, I would conclude she will win.

		Bob	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	0,0
	<i>LL</i>	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	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	0,0
	<i>LL</i>	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?

		Bob	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	0,0
	<i>LL</i>	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.

		Bob	
		t	l
Ann	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

In a game model $\mathcal{M}^G = \langle W, \{P_i\}_{i \in N}, \mathbf{s} \rangle$, different states represent different beliefs only when the agent is doing something different.

$$P_{i,w}(E) = P_i(E \mid [\mathbf{s}_i(w)])$$

To represent different *explanations* (i.e., beliefs) for the same strategy choice, we would need a set of models $\{\mathcal{M}_1^G, \mathcal{M}_2^G, \dots\}$.

In a game model $\mathcal{M}^G = \langle W, \{P_i\}_{i \in N}, \mathbf{s} \rangle$, different states represent different beliefs only when the agent is doing something different.

$$P_{i,w}(H) = P_i(H \mid B_{i,w}), \quad B_{i,w} \subseteq [\mathbf{s}_i(w)]$$

To represent different *explanations* (i.e., beliefs) for the same strategy choice, we would need a set of models $\{\mathcal{M}_1^G, \mathcal{M}_2^G, \dots\}$.

In a game model $\mathcal{M}^G = \langle W, \{P_i\}_{i \in N}, \mathbf{s} \rangle$, different states represent different beliefs only when the agent is doing something different.

$$P_{i,w}(H) = P_i(H \mid B_{i,w}), \quad B_{i,w} \subseteq [\mathbf{s}_i(w)]$$

Two way to change beliefs: $P_i(\cdot \mid E \cap B_{i,w})$ and $P_i(\cdot \mid B'_{i,w})$ (conditioning on 0 events).

Game Models

Richer models of a game: lexicographic probabilities, conditional probability systems, non-standard probabilities, plausibility models, . . .
(type spaces)

Game Models

Richer models of a game: lexicographic probabilities, conditional probability systems, non-standard probabilities, plausibility models, . . .
(type spaces)


“The aim in giving the general definition of a model is not to propose an original explanatory hypothesis, or any explanatory hypothesis, for the behavior of players in games, but only to provide a descriptive framework for the representation of considerations that are relevant to such explanations, a framework that is as *general* and as *neutral* as we can make it.” (pg. 35)


R. Stalnaker. *Knowledge, Belief and Counterfactual Reasoning in Games*. Economics and Philosophy, 12(1), pgs. 133 - 163, 1996.


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.









		<i>B</i>	
		l	r
<i>A</i>	<i>Out</i>	2, 1	2, 1
	u	4, 1	0, 0
	d	0, 0	1, 4

		<i>B</i>	
		l	r
<i>A</i>	<i>Out</i>	2, 1	2, 1
	u	4, 1	0, 0
	d	0, 0	1, 4

		<i>B</i>	
		l	r
<i>A</i>	<i>Out</i>	2, 1	2, 1
	u	4, 1	0, 0
	d	0, 0	1, 4

		<i>B</i>	
		l	r
<i>A</i>	<i>Out</i>	2, 1	2, 1
	u	4, 1	0, 0
	d	0, 0	1, 4


		<i>B</i>	
		l	r
	Out	2, 1	2, 1
<i>A</i>	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.

Backward *versus* Forward Induction




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

Backward *versus* Forward Induction




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

Backward *versus* Forward Induction




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

Backward versus Forward Induction




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

Backward versus Forward Induction




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

Backward *versus* Forward Induction



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

Backward versus Forward Induction



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

Backward and forward induction reasoning

We develop a Bayesian model of deliberation in extensive games.


Backward and forward induction reasoning


We develop a Bayesian model of deliberation in extensive games. Our main objective is to characterize backward and forward induction reasoning in terms of the assumptions that players make about the context of the game and the update mechanisms that drive the players' deliberations.


A. Knoks and EP. *Interpreting Mistakes in Games: From Beliefs about Mistakes to Mistaken Beliefs*. proceedings of TARK, 2015.


A. Knoks and EP. *Deliberational dynamics in context*. proceedings of LOFT, 2018.


	<i>L</i>	<i>R</i>
<i>U</i>	3, 1	0, 0
<i>D</i>	0, 0	1, 3























Taking stock

- ✓ Expected utility reasoning
- ✓ Introduction to game-theoretic reasoning: mixed strategies, Nash equilibrium, rationalizability
- ✓ A brief introduction to epistemic game theory
- ✓ Backward and forward induction
 - ▶ Prisoner's dilemma and repeated games

- ▶ Athletes using performance-enhancing drugs

- ▶ Athletes using performance-enhancing drugs
- ▶ Two competing companies deciding advertising budgets

- ▶ Athletes using performance-enhancing drugs
- ▶ Two competing companies deciding advertising budgets
- ▶ Nation-states deciding to restrict CO2 emissions

- ▶ Athletes using performance-enhancing drugs
- ▶ Two competing companies deciding advertising budgets
- ▶ Nation-states deciding to restrict CO2 emissions
- ▶ Two people meet and exchange closed bags, with the understanding that one of them contains money, and the other contains a purchase. Either player can choose to honor the deal by putting into his or her bag what he or she agreed, or he or she can defect by handing over an empty bag.

- ▶ Athletes using performance-enhancing drugs
- ▶ Two competing companies deciding advertising budgets
- ▶ Nation-states deciding to restrict CO2 emissions
- ▶ Two people meet and exchange closed bags, with the understanding that one of them contains money, and the other contains a purchase. Either player can choose to honor the deal by putting into his or her bag what he or she agreed, or he or she can defect by handing over an empty bag.
- ▶ <http://www.radiolab.org/story/golden-rule/>

Prisoner's Dilemma

Two people commit a crime.

Prisoner's Dilemma

Two people commit a crime. They are arrested by the police, who are quite sure they are guilty but cannot prove it without at least one of them confessing.

Prisoner's Dilemma

Two people commit a crime. They are arrested by the police, who are quite sure they are guilty but cannot prove it without at least one of them confessing. The police offer the following deal. Each one of them can confess and get credit for it.

Prisoner's Dilemma

Two people commit a crime. They are arrested by the police, who are quite sure they are guilty but cannot prove it without at least one of them confessing. The police offer the following deal. Each one of them can confess and get credit for it. If only one confesses, he becomes a state witness and not only is he not punished, he gets a reward.

Prisoner's Dilemma

Two people commit a crime. They are arrested by the police, who are quite sure they are guilty but cannot prove it without at least one of them confessing. The police offer the following deal. Each one of them can confess and get credit for it. If only one confesses, he becomes a state witness and not only is he not punished, he gets a reward. If both confess, they will be punished but will get reduced sentences for helping the police.

Prisoner's Dilemma

Two people commit a crime. They are arrested by the police, who are quite sure they are guilty but cannot prove it without at least one of them confessing. The police offer the following deal. Each one of them can confess and get credit for it. If only one confesses, he becomes a state witness and not only is he not punished, he gets a reward. If both confess, they will be punished but will get reduced sentences for helping the police. If neither confesses, the police honestly admit that there is no way to convict them, and they are set free.

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Possible outcomes:

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Possible outcomes: Both cooperate (C, C),

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Possible outcomes: Both cooperate (C, C), I cooperate but my partner doesn't (C, D),

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Possible outcomes: Both cooperate (C, C), I cooperate but my partner doesn't (C, D), My partner cooperates but I don't (D, C),

Prisoner's Dilemma

Two options: Cooperate with each other by not confessing (C), Defect by confessing (D)

Possible outcomes: Both cooperate (C, C), I cooperate but my partner doesn't (C, D), My partner cooperates but I don't (D, C), both defect (D, D).

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>		
	<i>D</i>		

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3	1
	<i>D</i>	4	2

Ann's preferences

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3	4
	<i>D</i>	1	2


Bob's preferences

Prisoner's Dilemma


		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do?


Dominance Reasoning



Dominance Reasoning



Dominance Reasoning



Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

A nasty nephew wants inheritance from his rich Aunt.

Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

A nasty nephew wants inheritance from his rich Aunt. The nephew wants the inheritance, but other things being equal, does not want to apologize.

Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

A nasty nephew wants inheritance from his rich Aunt. The nephew wants the inheritance, but other things being equal, does not want to apologize. Does dominance give the nephew a reason to not apologize?

Dominance reasoning is appropriate only when probability of outcome is *independent of choice*.

A nasty nephew wants inheritance from his rich Aunt. The nephew wants the inheritance, but other things being equal, does not want to apologize. Does dominance give the nephew a reason to not apologize? *Whether or not the nephew is cut from the will may depend on whether or not he apologizes.*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do?

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do? *Dominance reasoning*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do? *Dominance reasoning*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do? *Dominance reasoning* is not **Pareto!**

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3	2.5
	<i>D</i>	2.5	2

What should Ann (Bob) do? *Think as a group!*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do? *Play against your mirror image!*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	3,3	1,4
	<i>D</i>	4,1	2,2

What should Ann (Bob) do? *Play against your mirror image!*

Prisoner's Dilemma

		Bob	
		<i>C</i>	<i>D</i>
Ann	<i>C</i>	€ , €	1 , 4
	<i>D</i>	4 , 1	2 , 2

What should Ann (Bob) do? *Change the game...*

“Game theorists think it just plain wrong to claim that the Prisoners’ Dilemma embodies the essence of the problem of human cooperation.

“Game theorists think it just plain wrong to claim that the Prisoners’ Dilemma embodies the essence of the problem of human cooperation. On the contrary, it represents a situation in which the dice are as loaded against the emergence of cooperation as they could possibly be. If the great game of life played by the human species were the Prisoner’s Dilemma, we wouldn’t have evolved as social animals!

“Game theorists think it just plain wrong to claim that the Prisoners’ Dilemma embodies the essence of the problem of human cooperation. On the contrary, it represents a situation in which the dice are as loaded against the emergence of cooperation as they could possibly be. If the great game of life played by the human species were the Prisoner’s Dilemma, we wouldn’t have evolved as social animals! No paradox of rationality exists. Rational players don’t cooperate in the Prisoners’ Dilemma, because the conditions necessary for rational cooperation are absent in this game.” (pg. 63)

K. Binmore. *Natural Justice*. Oxford University Press, 2005.

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

...

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

Iterated Prisoner's Dilemma

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

...

"A fascinating perspective and important book"
—Douglas R. Hofstadter, author of *Gödel, Escher, Bach*

REVISED EDITION

THE EVOLUTION OF
COOPERATION

Robert Axelrod



Strategies

- ▶ Periodic: All-C, All-D, CD, CCD, CDD, CCDD, ...
- ▶ Random
- ▶ Memory: Tit-for-Tat, Two-Tit-for-Tat, ...

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

	<i>C</i>	<i>D</i>
<i>C</i>	3,3	0,4
<i>D</i>	4,0	1,1

...

Additional Reading

- ▶ S. Kuhn, Prisoner's Dilemma, Stanford Encyclopedia of Philosophy, plato.stanford.edu/entries/prisoner-dilemma/
- ▶ W. Poundstone, Prisoner's Dilemma, Anchor, 1993
- ▶ Online Game Theory Course (M. Jackson, K. Leyton-Brown and Y. Shoham): game-theory-class.org
- ▶ <http://axelrod.readthedocs.io/en/stable/>

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 Regret, 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)

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.”

Rational deliberation in games

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


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. 159)

		b	
		L	R
a	U	3,1	0,0
	D	0,0	1,3



a 's current state of indecision


		b	
		L	R
a	U	3,1	0,0
	D	0,0	1,3



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	3,1	0,0
	D	0,0	1,3



b 's current state of indecision

		b	
		L	R
a	U	3,1	0,0
	D	0,0	1,3

