

Reasoning in Games: Players as Programs


Lecture 2

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, forward and backward induction; Iterated games and learning, Skyrms's model of rational deliberation I

Wednesday Skyrms's model of rational deliberation II; Introduction to webppl; Game-theoretic reasoning in webppl

Thursday Coordination games (comparing Skyrms's model of deliberation and the webppl approach)

Friday Models of game-theoretic reasoning

Yesterday


- ▶ Guess $2/3$ average, Traveler's Dilemma
- ▶ Expected utility theory
- ▶ Introduction to game theory
- ▶ Zero-sum games, Nash equilibrium

Games

		<i>B</i>	
		l	r
<i>A</i>	u	3, 1	0, 0
	d	0, 0	1, 3


Games

		<i>B</i>	
		<i>l</i>	<i>r</i>
<i>A</i>	<i>u</i>	3, 1	0, 0
	<i>d</i>	0, 0	1, 3





Games

		<i>B</i>	
		<i>l</i>	<i>r</i>
<i>A</i>	<i>u</i>	3, 1	0, 0
	<i>d</i>	0, 0	1, 3



Games

		<i>B</i>	
		<i>l</i>	<i>r</i>
<i>A</i>	<i>u</i>	3, 1	0, 0
	<i>d</i>	0, 0	1, 3



From Decisions to Games, II

“*[T]he* 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.”

R. Aumann and J. Dreze. *Rational Expectations in Games*. *American Economic Review*, 98, pp. 72-86, 2008.

Let $G = \langle (S_i)_{i \in N}, (u_i)_{i \in N} \rangle$ be a finite strategic game.

$$\Sigma_i = \{p \mid p : S_i \rightarrow [0, 1] \text{ and } \sum_{s_i \in S_i} p(s_i) = 1\}$$

The **mixed extension** of G is the game $\langle (\Sigma_i)_{i \in N}, (U_i)_{i \in N} \rangle$ where for $\sigma \in \Sigma = \Sigma_1 \times \cdots \times \Sigma_n$:

$$U_i(\sigma) = \sum_{(s_1, \dots, s_n) \in S} \sigma_1(s_1) \sigma_2(s_2) \cdots \sigma_n(s_n) u_i(s_1, \dots, s_n)$$

Theorem (Nash). Every finite game G has a Nash equilibrium in mixed strategies (i.e., there is a Nash equilibrium in the mixed extension G).

Not all equilibrium are created equal...

Perfect equilibrium

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

Perfect equilibrium

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

Perfect equilibrium

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

Isn't (U, L) more "reasonable" than (D, R) ?

Perfect equilibrium

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


Completely mixed strategy: a mixed strategy in which every strategy gets some positive probability

Perfect equilibrium


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

Completely mixed strategy: a mixed strategy in which every strategy gets some positive probability

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*)

Stag-Hunt

		Bob	
		<i>S</i>	<i>H</i>
Ann	<i>S</i>	3, 3	0, 2
	<i>H</i>	2, 0	1, 1

Stag-Hunt

		Bob	
		<i>S</i>	<i>H</i>
Ann	<i>S</i>	3, 3	0, 2
	<i>H</i>	2, 0	1, 1

(S, S) and (H, H) are Nash equilibria

Stag-Hunt

		Bob	
		<i>S</i>	<i>H</i>
Ann	<i>S</i>	3, 3	0, 2
	<i>H</i>	2, 0	1, 1

(S, S) is Pareto-superior, but (H, H) is less risky

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	1, 1	2, 0	-2, 1
	<i>M</i>	0, 2	1, 1	2, 1
	<i>B</i>	1, -2	1, 2	1, 1

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	1, 1	2, 0	-2, 1
	<i>M</i>	0, 2	1, 1	2, 1
	<i>B</i>	1, -2	1, 2	1, 1

(T, L) is the unique pure-strategy Nash equilibrium

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	1, 1	2, 0	-2, 1
	<i>M</i>	0, 2	1, 1	2, 1
	<i>B</i>	1, -2	1, 2	1, 1

(T, L) is the unique pure-strategy Nash equilibrium

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	1, 1	2, 0	-2, 1
	<i>M</i>	0, 2	1, 1	2, 1
	<i>B</i>	1, -2	1, 2	1, 1

Why not play *B* and *R*?

Why play such an equilibrium?

“Let us now imagine that there exists a complete theory of the zero-sum two-person game which tells a player what to do, and which is absolutely convincing. If the players knew such a theory then each player would have to assume that his strategy has been “found out” by his opponent. The opponent knows the theory, and he knows that the player would be unwise not to follow it... a satisfactory theory can exist only if we are able to harmonize the two extremes...strategies of player 1 ‘found out’ or of player 2 ‘found out.’ ” (pg. 148)

J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Why play such an equilibrium?

“Von Neumann and Morgenstern are assuming that the *payoff matrix* is common knowledge to the players, but presumably the players' subjective probabilities might be private. Then each player might quite reasonably act to maximize subjective expected utility, believing that he will *not* be found out, with the result *not* being a Nash equilibrium.”

(Skyrms, pg. 14)

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

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

- ▶ Suppose that Ann believes Bob will play *L* with probability 1/4, for *whatever reason*. Then,

$$1 \times 0.25 + 4 \times 0.75 = 3.25 \geq 2 \times 0.25 + 3 \times 0.75 = 2.75$$

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

- ▶ Suppose that Ann believes Bob will play *L* with probability $1/4$, for whatever reason. Then,

$$1 \times 0.25 + 4 \times 0.75 = 3.25 \geq 2 \times 0.25 + 3 \times 0.75 = 2.75$$

- ▶ But, *L* is maximizes expected utility no matter what belief Bob may have:

$$p + 3 = 4 \times p + 3 \times (1 - p) \geq 1 \times p + 2 \times (1 - p) = 2 - p$$

Playing a Nash equilibrium is *required* by the players rationality and *common knowledge* thereof.

- ▶ Players need not be *certain* of the other players' beliefs
- ▶ Strategies that are not an equilibrium may be *rationalizable*
- ▶ Sometimes considerations of riskiness trump the Nash equilibrium

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

(M, C) is the unique Nash equilibrium

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

T, *L*, *B* and *R* are **rationalizable**

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

T, *L*, *B* and *R* are **rationalizable**

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

Ann plays *B* because she thought Bob will play *R*

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

Bob plays *L* because she thought Ann will play *B*

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3
	<i>M</i>	0, 0	1, 1	0, 0
	<i>B</i>	2, 3	0, 0	3, 2

Bob was correct, but Ann was wrong

		Bob			
		<i>L</i>	<i>C</i>	<i>R</i>	<i>X</i>
Ann	<i>T</i>	3, 2	0, 0	2, 3	0, -5
	<i>M</i>	0, 0	1, 1	0, 0	200, -5
	<i>B</i>	2, 3	0, 0	3, 2	1, -3

Not every strategy is rationalizable: Ann can't play *M* because she thinks Bob will play *X*

“Analysis of strategic economic situations requires us, implicitly or explicitly, to maintain as plausible certain psychological hypotheses. The hypothesis that real economic agents universally recognize the salience of Nash equilibria may well be less accurate than, for example, the hypothesis that agents attempt to “out-smart” or “second-guess” each other, believing that their opponents do likewise.” (pg. 1010)

B. D. Bernheim. *Rationalizable Strategic Behavior*. *Econometrica*, 52:4, pgs. 1007 - 1028, 1984.

“The rules of a game and its numerical data are seldom sufficient for logical deduction alone to single out a unique choice of strategy for each player. *To do so one requires either richer information (such as institutional detail or perhaps historical precedent for a certain type of behavior) or bolder assumptions about how players choose strategies.* Putting further restrictions on strategic choice is a complex and treacherous task. But one’s intuition frequently points to patterns of behavior that cannot be isolated on the grounds of consistency alone.”
(pg. 1035)

D. G. Pearce. *Rationalizable Strategic Behavior*. *Econometrica*, 52, 4, pgs. 1029 - 1050, 1984.

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

		B	
		<i>L</i>	<i>R</i>
A	<i>U</i>	2,1	0,0
	<i>D</i>	0,0	1,2

Strategic Game

		B	
		<i>L</i>	<i>R</i>
A	<i>U</i>	2,1	0,0
	<i>D</i>	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

		B	
		L	R
A	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

		B	
		L	R
A	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

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**.

Models of Games

Suppose that G is a game.

- ▶ Outcomes of the game: $S = \prod_{i \in N} S_i$
- ▶ A profile is a vector $\vec{s} \in S$, specifying an action for each player
- ▶ Player i 's partial beliefs (or conjecture): $P_i \in \Delta(S_{-i})$

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

Models of Games, continued

$G = \langle N, (S_i, u_i)_{i \in N} \rangle$ is a strategic (form of a) game.

Models of Games, continued

$G = \langle N, (S_i, u_i)_{i \in N} \rangle$ is a strategic (form of a) game.

- ▶ W is a set of *possible worlds* (possible outcomes of the game)
- ▶ s is a function $s : W \rightarrow \prod_{i \in N} S_i$, write $s_i(w)$ for the i th component of $s(w)$

Models of Games, continued

$G = \langle N, (S_i, u_i)_{i \in N} \rangle$ is a strategic (form of a) game.

- ▶ W is a set of *possible worlds* (possible outcomes of the game)
- ▶ \mathbf{s} is a function $\mathbf{s} : W \rightarrow \prod_{i \in N} S_i$, write $\mathbf{s}_i(w)$ for the i th component of $\mathbf{s}(w)$
- ▶ If $\vec{s} \in \prod_{i \in N} S_i$, then $[\vec{s}] = \{w \mid \mathbf{s}(w) = \vec{s}\}$; if $s_i \in S_i$, then $[s_i] = \{w \mid \mathbf{s}_i(w) = s_i\}$; and if $X \subseteq S$, $[X] = \bigcup_{s \in X} [s]$.

Models of Games, continued

$G = \langle N, (S_i, u_i)_{i \in N} \rangle$ is a strategic (form of a) game.

- ▶ W is a set of *possible worlds* (possible outcomes of the game)
- ▶ s is a function $s : W \rightarrow \prod_{i \in N} S_i$, write $s_i(w)$ for the i th component of $s(w)$
- ▶ If $\vec{s} \in \prod_{i \in N} S_i$, then $[\vec{s}] = \{w \mid \mathbf{s}(w) = \vec{s}\}$; if $s_i \in S_i$, then $[s_i] = \{w \mid s_i(w) = s_i\}$; and if $X \subseteq S$, $[X] = \bigcup_{s \in X} [s]$.
- ▶ **ex ante beliefs:** For each $i \in N$, let $P_i \in \Delta(W)$ (the set of probability measures on W). Two assumptions:
 - ▶ $[s]$ is measurable for all strategy profiles $s \in S$
 - ▶ $P_i([s_i]) > 0$ for all $s_i \in S_i$

ex interim beliefs: $P_{i,w} \in \Delta(S_{-i})$

- ▶ ...given player i 's choice: $P_{i,w}(\cdot) = P_i(\cdot \mid [\mathbf{s}_i(w)])$
- ▶ ...given all player i knows: $P_{i,w}(\cdot) = P_i(\cdot \mid K_i)$, $K_i \subseteq [\mathbf{s}_i(w)]$
- ▶ ...given all player i fully believes: $P_{i,w}(\cdot) = P_i(\cdot \mid B_i)$, $B_i \subseteq [\mathbf{s}_i(w)]$

ex interim beliefs: $P_{i,w} \in \Delta(S_{-i})$

- ▶ ...given player i 's choice: $P_{i,w}(\cdot) = P_i(\cdot \mid [\mathbf{s}_i(w)])$
- ▶ ...given all player i knows: $P_{i,w}(\cdot) = P_i(\cdot \mid K_i)$, $K_i \subseteq [\mathbf{s}_i(w)]$
- ▶ ...given all player i fully believes: $P_{i,w}(\cdot) = P_i(\cdot \mid B_i)$, $B_i \subseteq [\mathbf{s}_i(w)]$

Expected utility of strategy $s_i \in S_i$: Given $P \in \Delta(S_{-i})$,

$$EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i}) u_i(s_i, s_{-i})$$

ex interim beliefs: $P_{i,w} \in \Delta(S_{-i})$

- ▶ ...given player i 's choice: $P_{i,w}(\cdot) = P_i(\cdot \mid [\mathbf{s}_i(w)])$
- ▶ ...given all player i knows: $P_{i,w}(\cdot) = P_i(\cdot \mid K_i)$, $K_i \subseteq [\mathbf{s}_i(w)]$
- ▶ ...given all player i fully believes: $P_{i,w}(\cdot) = P_i(\cdot \mid B_i)$, $B_i \subseteq [\mathbf{s}_i(w)]$

Expected utility of strategy $s_i \in S_i$: Given $w \in W$,


$$EU_{i,w}(s_i) = \sum_{s_{-i} \in S_{-i}} P_{i,w}([s_{-i}]) u_i(s_i, s_{-i})$$

An Example

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


An Example

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




An Example

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



An Example

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




An Example

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

$\frac{1}{6}$ •	$\frac{1}{6}$ •	0•
$\frac{1}{6}$ •	0•	$\frac{1}{6}$ •
0•	$\frac{1}{6}$ •	$\frac{1}{6}$ •


An Example

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




An Example

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



An Example

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




- ▶ Ann's choice is *optimal* (given her information)

An Example

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

- Ann's choice is *optimal* (given her information)




$$1 \cdot P_A(L) + 0 \cdot P_A(R) \geq 0 \cdot P_A(L) + 2 \cdot P_A(R)$$

An Example

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

- Ann's choice is *optimal* (given her information)




$$1 \cdot P_A(L) + 0 \cdot P_A(R) \geq 0 \cdot P_A(L) + 2 \cdot P_A(R)$$

An Example

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

- Ann's choice is *optimal* (given her information)




$$1 \cdot \frac{3}{4} + 0 \cdot P_A(R) \geq 0 \cdot \frac{3}{4} + 2 \cdot P_A(R)$$

An Example

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


- ▶ Ann's choice is *optimal* (given her information)
- ▶ Bob's choice is *optimal* (given her information)



$$2 \cdot \frac{3}{4} + 0 \cdot \frac{1}{4} \geq 0 \cdot \frac{3}{4} + 1 \cdot \frac{1}{4}$$

An Example

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



- ▶ Ann's choice is *optimal* (given her information)
- ▶ Bob's choice is *optimal* (given her information)
- ▶ Bob *considers it possible* Ann is *irrational*

$$1 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2} \neq 0 \cdot \frac{1}{2} + 2 \cdot \frac{1}{2}$$

For any $P \in \Delta(S_{-i})$ and $s_i \in S_i$, $EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i})u_i(s_i, s_{-i})$

For any $P \in \Delta(S_{-i})$ and $s_i \in S_i$, $EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i})u_i(s_i, s_{-i})$

For any $w \in W$ and $s_i \in S_i$, $EU_{i,w}(s_i) = \sum_{s_{-i} \in S_{-i}} P_{i,w}([s_{-i}])u_i(s_i, s_{-i})$

For any $P \in \Delta(S_{-i})$ and $s_i \in S_i$, $EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i})u_i(s_i, s_{-i})$

For any $w \in W$ and $s_i \in S_i$, $EU_{i,w}(s_i) = \sum_{s_{-i} \in S_{-i}} P_{i,w}([s_{-i}])u_i(s_i, s_{-i})$

$\text{Rat}_i = \{w \mid EU_{i,w}(s_i(w)) \geq EU_{i,w}(s_i) \text{ for all } s_i \in S_i\}$

For any $P \in \Delta(S_{-i})$ and $s_i \in S_i$, $EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i})u_i(s_i, s_{-i})$

For any $w \in W$ and $s_i \in S_i$, $EU_{i,w}(s_i) = \sum_{s_{-i} \in S_{-i}} P_{i,w}([s_{-i}])u_i(s_i, s_{-i})$

$\text{Rat}_i = \{w \mid EU_{i,w}(s_i(w)) \geq EU_{i,w}(s_i) \text{ for all } s_i \in S_i\}$

Each $P \in \Delta(W)$ is associated with $P^S \in \Delta(S)$ as follows: for all $s \in S$, $P^S(s) = P([s])$

For any $P \in \Delta(S_{-i})$ and $s_i \in S_i$, $EU_{i,P}(s_i) = \sum_{s_{-i} \in S_{-i}} P(s_{-i})u_i(s_i, s_{-i})$

For any $w \in W$ and $s_i \in S_i$, $EU_{i,w}(s_i) = \sum_{s_{-i} \in S_{-i}} P_{i,w}([s_{-i}])u_i(s_i, s_{-i})$

$\text{Rat}_i = \{w \mid EU_{i,w}(s_i(w)) \geq EU_{i,w}(s_i) \text{ for all } s_i \in S_i\}$

Each $P \in \Delta(W)$ is associated with $P^S \in \Delta(S)$ as follows: for all $s \in S$, $P^S(s) = P([s])$

A mixed strategy $\sigma \in \prod_{i \in N} \Delta(S_i)$, $P_\sigma \in \Delta(S)$, $P_\sigma(s) = \sigma_1(s_1) \cdots \sigma_n(s_n)$

Characterizing Nash Equilibria

Theorem (Aumann). σ is a Nash equilibrium of G iff there exists a model $\mathcal{M}^G = \langle W, (P_i)_{i \in N}, \mathbf{s} \rangle$ such that:

- ▶ for all $i \in N$, $\text{Rat}_i = W$;
- ▶ for all $i, j \in N$, $P_i = P_j$; and
- ▶ for all $i \in N$, $P_i^S = P_\sigma$.

Rationalizability

A **best reply set** (BRS) is a sequence $(B_1, B_2, \dots, B_n) \subseteq S = \prod_{i \in N} S_i$ such that for all $i \in N$, and all $s_i \in B_i$, there exists $\mu_{-i} \in \Delta(B_{-i})$ such that s_i is a best response to μ_{-i} : i.e.,

$$b_i = \arg \max_{s_i \in S_i} EU_{i, \mu_{-i}}(s_i)$$

		2			
		b_1	b_2	b_3	b_4
1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

		2			
		b_1	b_2	b_3	b_4
1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

- ▶ (a_2, b_2) is the unique Nash equilibria, hence $(\{a_2\}, \{b_2\})$ is a BRS

		2			
		b_1	b_2	b_3	b_4
1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

- ▶ (a_2, b_2) is the unique Nash equilibria, hence $(\{a_2\}, \{b_2\})$ is a BRS
- ▶ $(\{a_1, a_3\}, \{b_1, b_3\})$ is a BRS

		2			
		b_1	b_2	b_3	b_4
1	a_1	0, 7	2, 5	7, 0	0, 1
	a_2	5, 2	3, 3	5, 2	0, 1
	a_3	7, 0	2, 5	0, 7	0, 1
	a_4	0, 0	0, -2	0, 0	10, -1

- ▶ (a_2, b_2) is the unique Nash equilibria, hence $(\{a_2\}, \{b_2\})$ is a BRS
- ▶ $(\{a_1, a_3\}, \{b_1, b_3\})$ is a BRS
- ▶ $(\{a_1, a_2, a_3\}, \{b_1, b_2, b_3\})$ is a full BRS

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	
		<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*.

		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

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

		<i>L</i>	Bob	<i>R</i>
Ann	<i>U</i>	5,*	1,*	
	<i>M</i>	1,*	5,*	
	<i>D</i>	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 $\text{supp}(P) = X$.

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

Is *R* rationalizable?

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

Is *R* rationalizable?

There is no *full support* probability such that *R* is a best response

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

Is *R* rationalizable?

There is no *full support* probability such that *R* is a best response

Should Ann assign probability 0 to *R* or probability > 0 to *R*?

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).

Common Knowledge of Admissibility

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>T</i>	1,1	1,0
	<i>B</i>	1,0	0,1

There is no model of this game with *common knowledge* of admissibility.

Common Knowledge of Admissibility

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>T</i>	1,1	1,0
	<i>B</i>	1,0	0,1

The "full" model of the game

Common Knowledge of Admissibility

		Bob	
		L	R
Ann	T	1,1	1,0
	B	1,0	0,1

The "full" model of the game: B is not admissible given Ann's information

Common Knowledge of Admissibility

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>T</i>	1,1	1,0
	<i>B</i>	1,0	0,1

What is wrong with this model?

Common Knowledge of Admissibility

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>T</i>	1,1	1,0
	<i>B</i>	1,0	0,1


The diagram illustrates the concept of common knowledge of admissibility. It shows a 2x2 normal form game matrix for Ann and Bob. The nodes of the game tree are represented by colored boxes with a central dot. The nodes (T,L) and (T,R) are grouped in a red dashed box. The nodes (B,L) and (B,R) are grouped in a blue solid box. The nodes (T,L) and (B,R) are also grouped in a red dashed box. The nodes (B,L) and (T,R) are also grouped in a blue solid box.

Privacy of Tie-Breaking/No Extraneous Beliefs: If a strategy is *rational* for an opponent, then it cannot be “ruled out”.


Summary

- ▶ Game models describe the *informational context* of a game.
- ▶ Game models can be used to *characterize* different solution concepts (e.g., iterated strict dominance, iterated weak dominance, Nash equilibrium, correlated equilibrium,...)


Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




Backward Induction




BI Puzzle




BI Puzzle




BI Puzzle




BI Puzzle




BI Puzzle




BI Puzzle




BI Puzzle



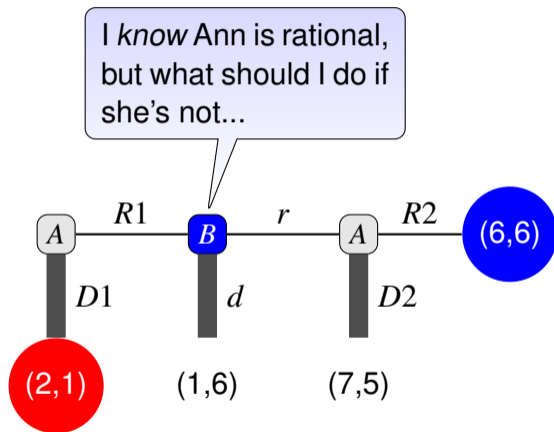
BI Puzzle





BI Puzzle?





BI Puzzle?













		Bob	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	1,0	1,0
	<i>L</i>	1,1	0,1



- ▶ The strategies of both players are rationalizable.
- ▶ Only *T* is *perfectly rational* for Ann and *t* is *perfectly rational* for Bob.

		Bob	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	3,3
	<i>LL</i>	1,1	0,0



Materially Rational: every choice actually made is optimal (i.e., maximizes subjective expected utility).


Substantively Rational: 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: every choice actually made is optimal (i.e., maximizes subjective expected utility).

Substantively Rational: 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	
		<i>t</i>	<i>l</i>
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	3,3
	<i>LL</i>	1,1	0,0




- ▶ Suppose that Bob believes that Ann will choose *T* with probability 1; what should he do? This depends on what he thinks Ann would do on the hypothesis that his belief about her is mistaken.
- ▶ Suppose that if Bob were surprised by her, then he concludes she is irrational, selecting *L* on her second move. Bob's choice of *t* is perfectly rational.

		<i>t</i> Bob <i>l</i>	
Ann	<i>T</i>	2,2	2,2
	<i>LT</i>	1,1	3,3
	<i>LL</i>	1,1	0,0



- ▶ Suppose Ann is sure that Bob will choose t , which is the only perfectly rational choice for Bob. Then, Ann's only rational choice is T .
- ▶ So, it might be that Ann and Bob both know each other's beliefs about each other, and are both perfectly rational, but they still fail to coordinate on the optimal outcome for both.



- ▶ Perhaps if Bob believed that Ann would choose L are her second move then he wouldn't believe she was fully rational, *but it is not suggested that he believes this.*
- ▶ Divide Ann's strategy T into two TT : T , and "I would choose T again on the second move if I were faced with that choice" and TL : " T , but I would choose L on the second move..."
- ▶ Of these two only TT is rational
- ▶ But if Bob **learned he was wrong**, he would conclude she is playing LL .

“To think there is something incoherent about this combination of beliefs and belief revision policy is to confuse epistemic with causal counterfactuals—it would be like thinking that because I believe that if Shakespeare hadn’t written Hamlet, it would have never been written by anyone, I must therefore be disposed to conclude that Hamlet was never written, were I to learn that Shakespeare was in fact not its author”

(pg. 152, Stalnaker)

1. If Shakespeare had not written Hamlet, it would never have been written.
2. If Shakespeare didn't write Hamlet, someone else did.

1. is a causal counterfactual, and 2. is an expression of a belief revision policy.

1. General Smith is a shrewd judge of character—he knows (better than I) who is brave and who is not.
2. The general sends only brave people into battle.
3. Private Jones is cowardly.

I believe that (1) Jones would run away if they were sent into battle and (2) if Jones *is* sent into battle, then they won't run away.

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.

		<i>t</i> Bob <i>l</i>	
<i>T</i>		2,2	2,2
<i>LT</i>	<i>Ann</i>	1,1	3,3
<i>LL</i>		1,1	0,0



- ▶ Bob believes that Ann is playing *TT*: Initially chooses *T*, and choose *T*, if she would have a second chance to move.
- ▶ If Bob *learns* that Ann does not play *T*, then he believes that she will play *LL*
- ▶ Ann believes all of the above about Bob.

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.