

Puzzles and Paradoxes from Decision and Game Theory

Eric Pacuit

University of Maryland

pacuit.org

July 20, 2017

Plan

- ✓ Day 1: Rational Choice Theory, Decision Theory
- ✓ Day 2: Expected Utility Theory, Allais Paradox
- ✓ Day 3: Evidential and Causal Decision Theory,
 - ▶ Day 4: Introduction to (Epistemic) Game Theory, Common Knowledge, Backward Induction, Decisions over Time
 - ▶ Day 5: Paradoxes of Interactive Epistemology, Framing in Games and Decisions

Taking Stock

- ▶ Many choice rules: MEU, strict/weak dominance, maxmin, minmax regret
 - ▶ Which one is “best”?
 - ▶ What are the relationships between the different choice rules?
- ▶ Payoff is not the same as utility (von Neumann-Morgenstern utilities)
- ▶ Rational choice models should be applied with care (act-state dependence, deliberation, attitudes towards risk, attitudes toward ambiguity, ...)

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess?

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess? 100

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess? ~~100~~, 99

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess? ~~100~~, ~~99~~, ..., 67

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess? ~~100~~, ~~99~~, ..., ~~67~~, ..., 2, 1

The Guessing Game



Guess a number between 1 & 100.
The closest to $\frac{2}{3}$ of the average wins.

What number should you guess? ~~100~~, ~~99~~, ..., ~~67~~, ..., ~~2~~, **1**

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in Euros from the airline in compensation.

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in Euros from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in Euros from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.
4. The person that wrote the smaller number will receive that amount plus 2 EUR (as a reward), and the person that wrote the larger number will receive the smaller number minus 2 EUR (as a punishment).

Traveler's Dilemma

1. You and your friend write down an integer between 2 and 100 (without discussing).
2. If both of you write down the same number, then both will receive that amount in Euros from the airline in compensation.
3. If the numbers are different, then the airline assumes that the smaller number is the actual price of the luggage.
4. The person that wrote the smaller number will receive that amount plus 2 EUR (as a reward), and the person that wrote the larger number will receive the smaller number minus 2 EUR (as a punishment).

Suppose that you are randomly paired with another person here at ESSLLI. What number would you write down?

From Decisions to Games, I

Commenting on the difference between Robin Crusoe's maximization problem and the maximization problem faced by participants in a social economy, von Neumann and Morgenstern write:

“Every participant can determine the variables which describe his own actions but not those of the others. Nevertheless those “alien” variables cannot, from his point of view, be described by statistical assumptions.

From Decisions to Games, I

Commenting on the difference between Robin Crusoe's maximization problem and the maximization problem faced by participants in a social economy, von Neumann and Morgenstern write:

“Every participant can determine the variables which describe his own actions but not those of the others. Nevertheless those “alien” variables cannot, from his point of view, be described by statistical assumptions. This is because the others are guided, just as he himself, by rational principles—whatever that may mean—and no *modus procedendi* can be correct which does not attempt to understand those principles and the interactions of the conflicting interests of all participants.”

(vNM, pg. 11)

Game Situations

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

	Bob	
	<i>L</i>	<i>R</i>
	1	0

1. a group of *self-interested* agents (players) involved in some interdependent *decision problem*

Game Situations

	Bob	
<i>L</i>		<i>R</i>
1	1	0
0	0	1

1. a group of *self-interested* agents (players) involved in some interdependent *decision problem*

Game Situations

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

1. a **group** of *self-interested* agents (players) involved in some interdependent decision problem

Game Situations

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

1. a **group** of *self-interested* agents (players) involved in some **interdependent decision problem**

Game Situations

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

1. a group of *self-interested* agents (players) involved in some interdependent decision problem

Just Enough Game Theory

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

- ▶ the actions the players *can* take
- ▶ the players' interests (i.e., preferences),
- ▶ the “structure” of the decision problem

Just Enough Game Theory

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

- ▶ the actions the players *can* take
- ▶ the players' interests (i.e., preferences),
- ▶ the “structure” of the decision problem


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

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



Questions

- ▶ Do players maximize (expected) utilities when playing games?

Questions

- ▶ Do players maximize (expected) utilities when playing games?
 - ▶ How, exactly, do you apply [revealed preference theory](#) to game theory?
 - ▶ How, exactly, do you apply [von Neumann-Morgenstern utility theory](#) to game theory?
 - ▶ How, exactly, do you apply [Savage's subjective expected utility theory](#) to game theory?
 - ▶ How, exactly, do you apply [Kahneman and Tversky's prospect theory](#) to game theory?

Questions

- ▶ Do players maximize (expected) utilities when playing games?
 - ▶ How, exactly, do you apply [revealed preference theory](#) to game theory?
 - ▶ How, exactly, do you apply [von Neumann-Morgenstern utility theory](#) to game theory?
 - ▶ How, exactly, do you apply [Savage's subjective expected utility theory](#) to game theory?
 - ▶ How, exactly, do you apply [Kahneman and Tversky's prospect theory](#) to game theory?
- ▶ What is game theory trying to accomplish?
(predictions? recommendations? explanations? analytical results?)

I. Gilboa and D. Schmeidler. *A Derivation of Expected Utility Maximization in the Context of a Game*. Games and Economic Behavior, 44, pgs. 184 - 194, 2003.

M. Mariotti. *Decisions in games: why there should be a special exemption from Bayesian rationality*. Journal of Economic Methodology, 4: 1, pgs. 43 - 60, 1997.

P. Hammond. *Expected Utility in Non-Cooperative Game Theory*. in *Handbook of Utility Theory*, 2004.

J. Kadane and P. Larkey. *Subjective Probability and the Theory of Games*. Management Science, Volume 28, 1982.

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.

Solution Concepts

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

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

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

Let $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ be a finite strategic game (each S_i is finite and the set of players N is finite).

A **strategy profile** is an element $\sigma \in S = S_1 \times \cdots \times S_n$

σ is a **Nash equilibrium** provided for all i , for all $s_i \in S_i$,

$$u_i(\sigma) \geq u_i(s_i, \sigma_{-i})$$

Zero-Sum Games

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

What should Ann *do*?

Zero-Sum Games

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

What should Ann *do*? *Bob best choice in Ann's worst choice*

Zero-Sum Games

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

What should Ann do? *Security strategy: minimize over each row and choose the maximum value*

Zero-Sum Games

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

What should Bob *do*? *Security strategy: minimize over each column and choose the maximum value*

Zero-Sum Games

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

The profile of security strategies (D, L) is a Nash equilibrium

Matching Pennies

		Bob	
		<i>H</i>	<i>T</i>
Ann	<i>H</i>	1, -1	-1, 1
	<i>T</i>	-1, 1	1, -1

There are no pure strategy Nash equilibria.

Mixed Strategies

		Bob	
		<i>H</i>	<i>T</i>
Ann	<i>H</i>	1, -1	-1, 1
	<i>T</i>	-1, 1	1, -1

A **mixed strategy** is a probability distribution over the set of pure strategies. For instance:

- ▶ $(1/2H, 1/2T)$
- ▶ $(1/3H, 2/3T)$
- ▶ ...

Matching Pennies

		Bob	
		<i>H</i>	<i>T</i>
Ann	<i>H</i>	1, -1	-1, 1
	<i>T</i>	-1, 1	1, -1

The mixed strategy $([1/2 : H, 1/2 : T], [1/2 : H, 1/2 : T])$ is the only Nash equilibrium.

Theorem (Von Neumann). For every two-player zero-sum game with finite strategy sets S_1 and S_2 , there is a number v , called the **value** of the game such that:

1. $v = \max_{p \in \Delta(S_1)} \min_{q \in \Delta(S_2)} U_1(p, q) = \min_{q \in \Delta(S_2)} \max_{p \in \Delta(S_1)} U_1(p, q)$
2. The set of mixed Nash equilibria is nonempty. A mixed strategy profile (p, q) is a Nash equilibrium if and only if

$$p \in \operatorname{argmax}_{p \in \Delta(S_1)} \min_{q \in \Delta(S_2)} U_1(p, q)$$

$$q \in \operatorname{argmax}_{q \in \Delta(S_2)} \min_{p \in \Delta(S_1)} U_1(p, q)$$

3. For all mixed Nash equilibria (p, q) , $U_1(p, q) = v$

In zero-sum games

- ▶ There exists a mixed strategy Nash equilibrium
- ▶ There may be more than one Nash equilibria
- ▶ Security strategies are always a Nash equilibrium
- ▶ Components of Nash equilibria are interchangeable: If σ and σ' are Nash equilibria in a 2-player game, then (σ_1, σ'_2) is also a Nash equilibria.

Finding the rational choice...

		Bob	
		<i>H</i>	<i>T</i>
Ann	<i>H</i>	1,-1	-1,1
	<i>T</i>	-1,1	1,-1

What is a rational choice for Ann (Bob)?

Finding the rational choice...

		Bob	
		<i>H</i>	<i>T</i>
Ann	<i>H</i>	1,-1	-1,1
	<i>T</i>	-1,1	1,-1

What is a rational choice for Ann (Bob)? *Flip a coin!*

Finding the rational choice...

		Bob	
		C1	C2
Ann	P1	1,-1	-1,1
	P2	-1,1	1,-1

What is a rational choice for Ann (Bob)?

Finding the rational choice...

		Bob	
		C1	C2
Ann	P1	1,-1	-1,1
	P2	-1,1	1,-1

		Bob	
		C1	C2
Ann	P1	1,-1	1,-1
	P2	1,-1	1,-1

What is a rational choice for Ann (Bob)? *Play a different game!*

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 $\sigma \in \Sigma = \Sigma_1 \times \cdots \times \Sigma_n$:

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


		Column	
		L	R
Row	U	$(2, 1)$	$(0, 0)$
	D	$(0, 0)$	$(1, 2)$

- ▶ $N = \{Row, Column\}$
- ▶ $A_{Row} = \{U, D\}, A_{Column} = \{L, R\}$
- ▶ $u_{Row} : A_{Row} \times A_{Column} \rightarrow \{0, 1, 2\}, u_{Column} : A_{Row} \times A_{Column} \rightarrow \{0, 1, 2\}$ with
 $u_{Row}(U, L) = 2; u_{Column}(U, L) = 1; u_{Row}(D, R) = 1; u_{Column}(D, R) = 2,$
and $u_x(U, L) = u_x(D, R) = 0$ for $x \in N$.


Mixed Extension



Mixed Extension



Mixed Extension



Theorem. Suppose that σ is a Nash equilibrium in mixed strategies for a game $G = \langle \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$. Suppose that $s_i, s_i^* \in S_i$ are two pure strategies such that $\sigma_i(s_i) > 0$ and $\sigma_i(s_i^*) > 0$, then

$$U_i(s_i, \sigma_{-i}) = U_i(s_i^*, \sigma_{-i})$$

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).

Mixed Strategies

“We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes.”

Mixed Strategies

“We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes.”

- ▶ One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.

Mixed Strategies

“We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes.”

- ▶ One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.
- ▶ *Harsanyi's purification theorem*: A player's mixed strategy is thought of as a plan of action which is dependent on private information which is not specified in the model. Although the player's behavior appears to be random, it is actually deterministic.

Mixed Strategies

“We are reluctant to believe that our decisions are made at random. We prefer to be able to point to a reason for each action we take. Outside of Las Vegas we do not spin roulettes.”

- ▶ One can think about a game as an interaction between large populations...a mixed strategy is viewed as the distribution of the pure choices in the population.
- ▶ *Harsanyi's purification theorem*: A player's mixed strategy is thought of as a plan of action which is dependent on private information which is not specified in the model. Although the player's behavior appears to be random, it is actually deterministic.
- ▶ Mixed strategies are beliefs held by all *other* players concerning a player's actions.

Not all equilibrium are created equal...

Why play Nash equilibrium?

Self-Enforcing Agreements: Nash equilibria are recommended by being the only strategy combinations on which the players could make self-enforcing agreements, i.e., agreements that each has reason to respect, even without external enforcement mechanisms.

M. Risse. *What is rational about Nash equilibria?*. Synthese, 124:3, pgs. 361 - 384, 2000.

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	4, 6	5, 4	0, 0
	<i>M</i>	5, 7	4, 8	0, 0
	<i>B</i>	0, 0	0, 0	1, 1

		Bob		
		<i>L</i>	<i>C</i>	<i>R</i>
Ann	<i>T</i>	4, 6	5, 4	0, 0
	<i>M</i>	5, 7	4, 8	0, 0
	<i>B</i>	0, 0	0, 0	1, 1

(B, R) is a Nash equilibrium, but it is **not self-enforcing**

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	0, 0	4, 2
	<i>D</i>	2, 4	3, 3

		Bob	
		<i>L</i>	<i>R</i>
Ann	<i>U</i>	0, 0	4, 2
	<i>D</i>	2, 4	3, 3

(D,R) is self-enforcing, but **not a Nash equilibrium**

Self-Enforcing Agreements: Nash equilibria are recommended by being the only strategy combinations on which the players could make self-enforcing agreements, i.e., agreements that each has reason to respect, even without external enforcement mechanisms.

- ▶ Not all Nash equilibria are “equally” self-enforcing
- ▶ There are Nash equilibria that are not self-enforcing
- ▶ There are self-enforcing outcomes that are not Nash equilibria

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.

Pure Coordination Game

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

The profiles **(U, L)** and **(D, R)** are Nash equilibria.

Prisoner's Dilemma

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

What should Ann (Bob) do?

Prisoner's Dilemma

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

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

Prisoner's Dilemma

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

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

Prisoner's Dilemma

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

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

Traveler's Dilemma

	2	3	4	...	99	100
2	(2, 2)	(4, 0)	(4, 0)	...	(4, 0)	(4, 0)
3	(0, 4)	(3, 3)	(5, 1)	...	(5, 1)	(5, 1)
4	(0, 4)	(1, 5)	(4, 4)	...	(6, 2)	(6, 2)
⋮	⋮	⋮	⋮	⋮	⋮	⋮
99	(0, 4)	(1, 5)	(2, 6)	...	(99, 99)	(101, 97)
100	(0, 4)	(1, 5)	(2, 6)	...	(97, 101)	(100, 100)

Traveler's Dilemma

	2	3	4	...	99	100
2	(2, 2)	(4, 0)	(4, 0)	...	(4, 0)	(4, 0)
3	(0, 4)	(3, 3)	(5, 1)	...	(5, 1)	(5, 1)
4	(0, 4)	(1, 5)	(4, 4)	...	(6, 2)	(6, 2)
⋮	⋮	⋮	⋮	⋮	⋮	⋮
99	(0, 4)	(1, 5)	(2, 6)	...	(99, 99)	(101, 97)
100	(0, 4)	(1, 5)	(2, 6)	...	(97, 101)	(100, 100)

(2, 2) is the only Nash equilibrium.

Traveler's Dilemma

	2	3	4	...	99	100
2	(2, 2)	(4, 0)	(4, 0)	...	(4, 0)	(4, 0)
3	(0, 4)	(3, 3)	(5, 1)	...	(5, 1)	(5, 1)
4	(0, 4)	(1, 5)	(4, 4)	...	(6, 2)	(6, 2)
⋮	⋮	⋮	⋮	⋮	⋮	⋮
99	(0, 4)	(1, 5)	(2, 6)	...	(99, 99)	(101, 97)
100	(0, 4)	(1, 5)	(2, 6)	...	(97, 101)	(100, 100)

The analysis is insensitive to the amount of reward/punishment.

In an arbitrary (finite) strategic games

- ▶ There exists a mixed strategy Nash equilibrium
- ▶ Security strategies are not necessarily a Nash equilibrium
- ▶ There may be more than one Nash equilibrium
- ▶ Components of Nash equilibrium are not interchangeable.

“...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

The Epistemic Program in Game Theory

“...the analysis constitutes a fleshing-out of the textbook interpretation of equilibrium as ‘rationality plus correct beliefs.’...this suggests that equilibrium behavior cannot arise out of strategic reasoning alone. ”


E. Dekel and M. Siniscalchi. *Epistemic Game Theory*. manuscript, 2013.

A. Brandenburger. *The Power of Paradox*. International Journal of Game Theory, 35, pgs. 465 - 492, 2007.


EP and O. Roy. *Epistemic Game Theory*. Stanford Encyclopedia of Philosophy, 2015.

A. Perea. *Epistemic Game Theory*. Cambridge University Press, 2012.


Extensive Form




Extensive Form




Extensive Form




Extensive Form




Extensive Form



Extensive Form




Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Normal form vs. Extensive form




	l if L	r if L
L	$-1, -1$	$1, 1$
R	$0, 0$	$0, 0$

Normal form vs. Extensive form



	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$


Normal form vs. Extensive form




	l if L	r if L
L	$-1,-1$	$1,1$
R	$0,0$	$0,0$

Incredible threat


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



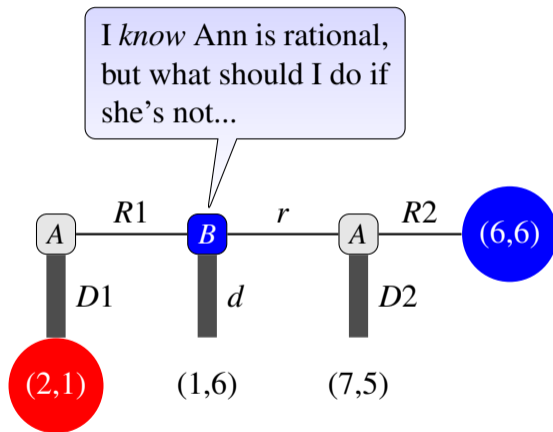
BI Puzzle



BI Puzzle?





BI Puzzle?





Experimentally, 92% of participants choose to continue at the first node. This is perhaps attributed to a social norm of reciprocity - If player 1 continues at the first node, it is more likely that player 2 will also play continue at the second node. Given this behavior, the optimal choice (the one that yields the highest payoff) is actually for player 1 to play continue: Given the distribution of actual play in the laboratory, the ones who play stop are actually making a mistake!


McKelvey and Palfrey. *An experimental study of the centipede game*. Games and Economic Behavior, 1992.














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




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

		Bob	
		<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