Epistemic Game Theory

Lecture 2

Eric Pacuit

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

February 3, 2014

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, form his point of view, be described by statistical assumptions.

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, form 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)

It is a problem of *mutual expectations*:

"Not a single datum with which he [Crusoe] has to deal reflects another person's will or intention of an economic kind—based on motives of the same nature as his own.

It is a problem of *mutual expectations*:

"Not a single datum with which he [Crusoe] has to deal reflects another person's will or intention of an economic kind—based on motives of the same nature as his own. A participant in a social exchange economy, on the other hand, faces data of this last type as well: they are the product of other participants' actions and volitions (like prices).

It is a problem of *mutual expectations*:

"Not a single datum with which he [Crusoe] has to deal reflects another person's will or intention of an economic kind—based on motives of the same nature as his own. A participant in a social exchange economy, on the other hand, faces data of this last type as well: they are the product of other participants' actions and volitions (like prices). His actions will be influenced by his expectation of these, and they in turn reflect the other participants' expectation of his actions....it is this problem which the theory of "games of strategy" is mainly devised to meet."

(vNM, pg. 11, 12)

"...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

"Perhaps it is time to reunite the two streams of work descended from von Neumann and Morgenstern (1944), prescriptive theories of individual decision making and theories of strategically interactive decisions, and to look to other disciplines such as cognitive psychology for predictive theories of decisional behavior."

(Kadane and Larkey, pg. 118)

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?)

- What might the players' be thinking about?
- Do not confuse modeling with analyzing a game situation!
- Can the decision problem be *separated* from the game situation?
- Can a player assign subjective probabilities to strategies under the control of other players who have their own objectives?

Knowledge and beliefs in game situations

J. Harsanyi. *Games with incomplete information played by "Bayesian" players I-III*. *Management Science Theory* **14***: 159-182, 1967-68.*

Robert Aumann. *Agreeing to Disagree*. Annals of Statistics **4** (1976).

R. Aumann. *Interactive Epistemology I & II*. International Journal of Game Theory (1999).

P. Battigalli and G. Bonanno. *Recent results on belief, knowledge and the epistemic foundations of game theory*. Research in Economics (1999).

R. Myerson. *Harsanyi's Games with Incomplete Information*. Special 50th anniversary issue of *Management Science*, 2004.

John C. Harsanyi, nobel prize winner in economics, developed a theory of games with **incomplete information**.

J. Harsanyi. *Games with incomplete information played by "Bayesian" players I-III. Management Science Theory* **14***: 159-182, 1967-68.*

John C. Harsanyi, nobel prize winner in economics, developed a theory of games with **incomplete information**.

1. incomplete information: uncertainty about the *structure* of the game (outcomes, payoffs, strategy space)
2. imperfect information: uncertainty *within the game* about the previous moves of the players

J. Harsanyi. *Games with incomplete information played by "Bayesian" players I-III. Management Science Theory* **14***: 159-182, 1967-68.*

Harsanyi's Problem

A natural question following any game-theoretic analysis is

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*?

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*? How do we completely specify such a model?

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*? How do we completely specify such a model?

1. Suppose there is a parameter that some player i does not know

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*? How do we completely specify such a model?

1. Suppose there is a parameter that some player i does not know
2. i's uncertainty about the parameter must be included in the model (first-order beliefs)

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*? How do we completely specify such a model?

1. Suppose there is a parameter that some player i does not know
2. i's uncertainty about the parameter must be included in the model (first-order beliefs)
3. this is a new parameter that the other players may not know, so we must specify the players beliefs about this parameter (second-order beliefs)

Harsanyi's Problem

A natural question following any game-theoretic analysis is *how would the players react if some parameters of the model are not known to the players*? How do we completely specify such a model?

1. Suppose there is a parameter that some player i does not know
2. i's uncertainty about the parameter must be included in the model (first-order beliefs)
3. this is a new parameter that the other players may not know, so we must specify the players beliefs about this parameter (second-order beliefs)
4. but this is a new parameter, and so on....

Harsanyi's Problem

A (game-theoretic) **type** of a player summarizes everything the player knows privately at the beginning of the game which could affect his beliefs about payoffs in the game and about all other players' types.

(Harsanyi argued that all uncertainty in a game can be equivalently modeled as uncertainty about payoff functions.)

Information in games situations

- imperfect information about the play of the game
- incomplete information about the structure of the game

Information in games situations

- imperfect information about the play of the game
- incomplete information about the structure of the game
- strategic information (what will the other players do?)
- higher-order information (what are the other players thinking?)

Information in games situations

- Various states of information disclosure.

Information in games situations

- Various states of information disclosure.
  - *ex ante*, *ex interim*, *ex post*

Information in games situations

- Various states of information disclosure.
  - *ex ante*, *ex interim*, *ex post*

- Various "types" of information:

Information in games situations

- Various states of information disclosure.
  - *ex ante*, *ex interim*, *ex post*

- Various "types" of information:
  - imperfect information about the play of the game
  - incomplete information about the structure of the game
  - strategic information (what will the other players do?)
  - higher-order information (what are the other players thinking?)

Information in games situations

- Various states of information disclosure.
  - *ex ante*, *ex interim*, *ex post*

- Various “types” of information:
  - imperfect information about the play of the game
  - incomplete information about the structure of the game
  - strategic information (what will the other players do?)
  - higher-order information (what are the other players thinking?)

- Varieties of informational attitudes

Information in games situations

- Various states of information disclosure.
  - *ex ante*, *ex interim*, *ex post*
- Various “types” of information:
  - imperfect information about the play of the game
  - incomplete information about the structure of the game
  - strategic information (what will the other players do?)
  - higher-order information (what are the other players thinking?)
- Varieties of informational attitudes
  - hard (“knowledge”)
  - soft (“beliefs”)

Models of Hard and Soft Information

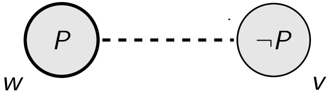

Epistemic Model: $\mathcal{M} = \langle W, \{\sim_i\}_{i\in\mathcal{A}}, V\rangle$

- $w \sim_i v$ means i cannot rule out v according to her information.

Language: $\varphi := p \mid \neg\varphi \mid \varphi \wedge \psi \mid K_i\varphi$

Truth:

- $\mathcal{M}, w \models p$ iff $w \in V(p)$ (p an atomic proposition)
- Boolean connectives as usual
- $\mathcal{M}, w \models K_i\varphi$ iff for all $v \in W$, if $w \sim_i v$ then $\mathcal{M}, v \models \varphi$

Models of Hard and Soft Information

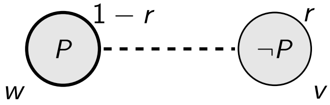

Epistemic-Probability Model: $\mathcal{M} = \langle W, \{\sim_i\}_{i\in\mathcal{A}}, \{\pi_i\}_{i\in\mathcal{A}}, V\rangle$

- $\pi_i : W \to [0,1]$ is a probability measure

Language: $\varphi := p \mid \neg\varphi \mid \varphi\wedge\psi \mid K_i\varphi \mid B^p\psi$

Truth:

- $[\![\varphi]\!]_{\mathcal{M}} = \{w \mid \mathcal{M}, w \models \varphi\}$
- $\mathcal{M}, w \models B^p\varphi$ iff $\pi_i([\![\varphi]\!]_{\mathcal{M}} \mid [w]_i) = \frac{\pi_i([\![\varphi]\!]_{\mathcal{M}} \cap [w]_i)}{\pi_i([w]_i)} \geq p$, $\mathcal{M}, v \models \psi$
- $\mathcal{M}, w \models K_i\varphi$ iff for all $v \in W$, if $w \sim_i v$ then $\mathcal{M}, v \models \varphi$

An Example

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

An Example

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

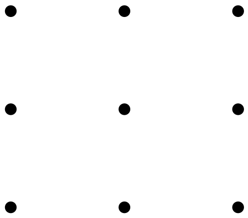

An Example

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

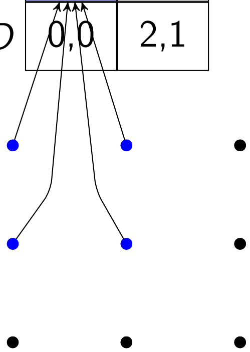

An Example

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

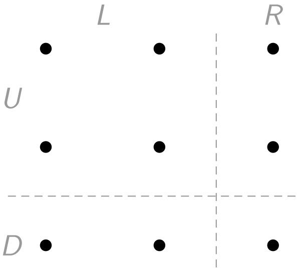

An Example

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

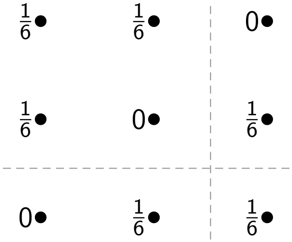

An Example

Bob

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

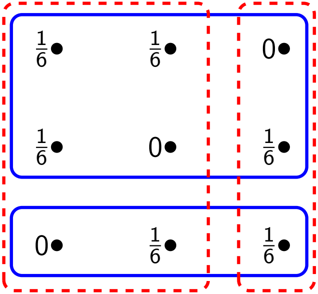

An Example

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

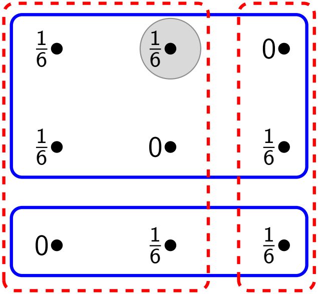

An Example

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

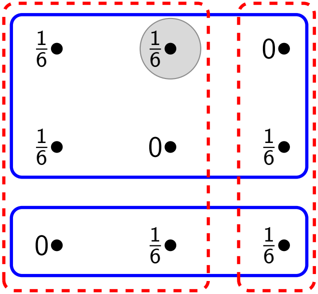

- Ann's choice is *optimal* (given her information)

An Example

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

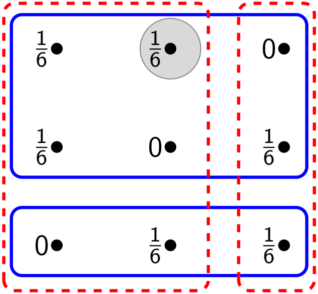

- Ann's choice is *optimal* (given her information)

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

An Example

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

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

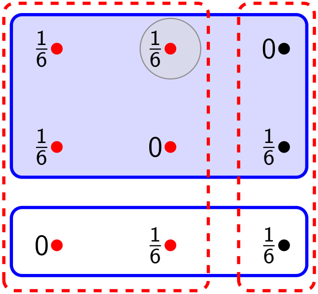

$$1\cdot P_A(L)+0\cdot P_A(R) \geq 0\cdot P_A(L)+2\cdot P_A(R)$$

An Example

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

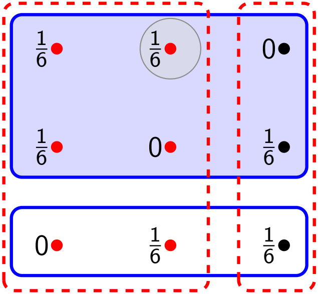

- Ann's choice is *optimal* (given her information)

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

An Example

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

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

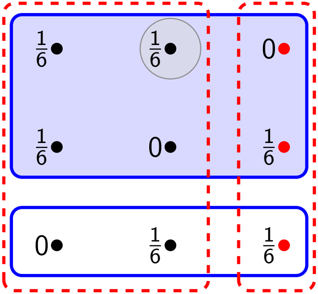

$$1 \cdot \frac{3}{4} + 0 \cdot \frac{1}{4} \geq 0 \cdot \frac{3}{4} + 2 \cdot \frac{1}{4}$$

An Example

Ann \ Bob	L	R
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)

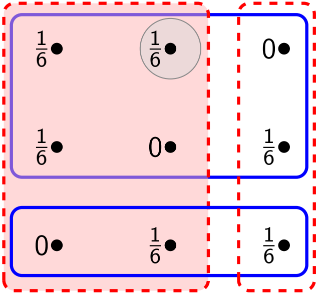

$$2 \cdot \frac{3}{4} + 0 \cdot \frac{1}{4} \geq 0 \cdot \frac{3}{4} + 1 \cdot \frac{1}{4}$$

An Example

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

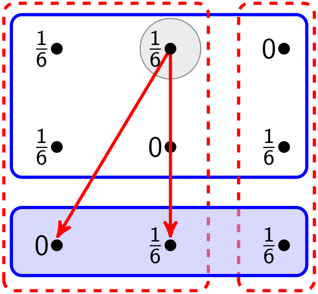

- Ann's choice is *optimal* (given her information)
- Bob's choice is *optimal* (given her information)
- Bob *considers it possible* Ann is *irrational*

$$1 \cdot \frac{1}{2} + 0 \cdot \frac{1}{2} \not\geq 0 \cdot \frac{1}{2} + 2 \cdot \frac{1}{2}$$

Modeling *Interactive* Beliefs

1. **Explicit description**: (infinite) sequence of σ-additive probability measures

$$\underbrace{S}_{\text{states}}$$

Modeling *Interactive* Beliefs

1. **Explicit description**: (infinite) sequence of σ-additive probability measures

$$\underbrace{S}_{\text{states}} \times \underbrace{\Delta(S)}_{\text{1st-order beliefs}}$$

Modeling *Interactive* Beliefs

1. **Explicit description**: (infinite) sequence of σ-additive probability measures

$$\underbrace{S}_{\text{states}} \quad \times \quad \underbrace{\Delta(S)}_{\text{1st-order beliefs}} \quad \times \quad \underbrace{\Delta(S \times \Delta(S))}_{\text{2nd-order beliefs}}$$

Modeling *Interactive* Beliefs

1. **Explicit description**: (infinite) sequence of σ-additive probability measures

$$\underbrace{S}_{\text{states}} \quad \times \quad \underbrace{\Delta(S)}_{\text{1st-order beliefs}} \quad \times \quad \underbrace{\Delta(S \times \Delta(S))}_{\text{2nd-order beliefs}} \quad \times \quad \cdots$$

Modeling *Interactive* Beliefs

1. **Explicit description**: (infinite) sequence of σ-additive probability measures

$$\underbrace{S}_{\text{states}} \quad \times \quad \underbrace{\Delta(S)}_{\text{1st-order beliefs}} \quad \times \quad \underbrace{\Delta(S \times \Delta(S))}_{\text{2nd-order beliefs}} \quad \times \quad \cdots$$

2. **Implicit description**: *Harsanyi type spaces* (sorted structure with maps between players' "states")

J. Harsanyi. *Games with incomplete information played by "bayesian" players I-III. Management Science Theory* **14**: *159-182, 1967-68.*

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

- Each **type** is assigned a joint probability over the space of types and actions

$$\lambda_i : T_i \to \Delta(T_{-i} \times S_{-i})$$

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

- Each **type** is assigned a joint probability over the space of types and actions

$$\lambda_i : \underbrace{T_i}_{\text{Player } i\text{'s types}} \to \Delta(T_{-i} \times S_{-i})$$

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

- Each **type** is assigned a joint probability over the space of types and actions

$$\lambda_i : T_i \to \Delta(T_{-i} \times S_{-i})$$

The set of all probability distributions

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

- Each **type** is assigned a joint probability over the space of types and actions

$$\lambda_i : T_i \to \Delta(\underbrace{T_{-i}}_{\text{The other players' types}} \times S_{-i})$$

Harsanyi Type Space

Based on the work of John Harsanyi on games with *incomplete information*, game theorists have developed an elegant formalism that makes precise talk about beliefs, knowledge and rationality:

- A **type** is everything a player knows privately at the beginning of the game which could affect his beliefs about payoffs and about all other players' possible types.

- Each **type** is assigned a joint probability over the space of types and actions

$$\lambda_i : T_i \to \Delta(T_{-i} \times S_{-i})$$

The other players' choices

Returning to the Example: A Game Model

		Bob H	Bob M
Ann	H	3,3	0,0
Ann	M	0,0	1,1

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

Ann \ Bob	H	M
H	3,3	0,0
M	0,0	1,1

- One type for Ann (t_A) and two types for Bob (t_B, u_B)

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

Bob

Ann	H	M
H	3,3	0,0
M	0,0	1,1

- One type for Ann (t_A) and two types for Bob (t_B, u_B)
- A **state** is a tuple of choices and types: (M, M, t_A, t_B)

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

Bob

Ann	H	M
H	3,3	0,0
M	0,0	1,1

- One type for Ann (t_A) and two types for Bob (t_B, u_B)
- A **state** is a tuple of choices and types: (M, t_A, M, u_B)
- Calculate **expected utility** in the usual way...

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

Bob

		H	M
Ann	H	3,3	0,0
	M	0,0	1,1

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

		Bob	
		H	M
Ann	H	3,3	0,0
	M	0,0	1,1

- M is **rational** for Ann (t_A)
  $0 \cdot 0.2 + 1 \cdot 0.8 \geq 3 \cdot 0.2 + 0 \cdot 0.8$

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

Ann \ Bob	H	M
H	3,3	0,0
M	0,0	1,1

- M is **rational** for Ann (t_A)
  $0 \cdot 0.2 + 1 \cdot 0.8 \geq 3 \cdot 0.2 + 0 \cdot 0.8$
- M is **rational** for Bob (t_B)
  $0 \cdot 0 + 1 \cdot 1 \geq 3 \cdot 0 + 0 \cdot 1$

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

Returning to the Example: A Game Model

		Bob	
		H	M
Ann	H	3,3	0,0
	M	0,0	1,1

- M is **rational** for Ann (t_A)
  $0 \cdot 0.2 + 1 \cdot 0.8 \geq 3 \cdot 0.2 + 0 \cdot 0.8$
- M is **rational** for Bob (t_B)
  $0 \cdot 0 + 1 \cdot 1 \geq 3 \cdot 0 + 0 \cdot 1$
- Ann thinks Bob may be irrational

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_A

	H	M
t_A	0	1

t_B

	H	M
t_A	0.4	0.6

u_B

Returning to the Example: A Game Model

Bob

Ann	H	M
H	3,3	0,0
M	0,0	1,1

- M is **rational** for Ann (t_A)
  $0 \cdot 0.2 + 1 \cdot 0.8 \geq 3 \cdot 0.2 + 0 \cdot 0.8$
- M is **rational** for Bob (t_B)
  $0 \cdot 0 + 1 \cdot 1 \geq 3 \cdot 0 + 0 \cdot 1$
- Ann thinks Bob may be irrational
  $P_A(Irrat[B]) = 0.3$, $P_A(Rat[B]) = 0.7$

t_A

	H	M
t_B	0	0.5
u_B	0.2	0.3

t_B

	H	M
t_A	0	1

u_B

	H	M
t_A	0.4	0.6

More on Types

For simplicity, we assume $S = \times_{i \in \mathcal{A}} S_i$, where each S_i is a strategy space for agent i in some fixed game G. In this case, $\lambda_i : T_i \rightarrow \Delta(S_{-i} \times T_{-i})$.

A fixed state $(s_1, t_1, s_2, t_2, \ldots, s_n, t_n)$ specifies the strategies and each player's *entire hierarchy of beliefs*:

More on Types

For simplicity, we assume $S = \times_{i \in \mathcal{A}} S_i$, where each S_i is a strategy space for agent i in some fixed game G. In this case, $\lambda_i : T_i \to \Delta(S_{-i} \times T_{-i})$.

A fixed state $(s_1, t_1, s_2, t_2, \ldots, s_n, t_n)$ specifies the strategies and each player's *entire hierarchy of beliefs*:

1. i's first-order beliefs: $T_i \mapsto \Delta(S_{-i} \times T_{-i}) \mapsto \Delta(S_{-i})$ (marginalizing)

More on Types

For simplicity, we assume $S = \times_{i \in \mathcal{A}} S_i$, where each S_i is a strategy space for agent i in some fixed game G. In this case, $\lambda_i : T_i \rightarrow \Delta(S_{-i} \times T_{-i})$.

A fixed state $(s_1, t_1, s_2, t_2, \ldots, s_n, t_n)$ specifies the strategies and each player's *entire hierarchy of beliefs*:

1. i's first-order beliefs: $T_i \mapsto \Delta(S_{-i} \times T_{-i}) \mapsto \Delta(S_{-i})$ (marginalizing)

2. i's second-order beliefs: $T_i \mapsto \Delta(S_{-i} \times T_{-i}) \mapsto$
   $\Delta(S^{-i} \times \times_{i \neq j} \Delta(S_{-j} \times T_{-j})) \mapsto \Delta(S_{-i} \times \times_{j \neq i} \Delta(S_{-j}))$
   (marginalizing)

More on Types

- For any given set S of external states we can use a type space on S to provide consistent representations of the players' beliefs.

More on Types

- For any given set S of external states we can use a type space on S to provide consistent representations of the players' beliefs.

- Every state in a belief model or type space induces an infinite hierarchy of beliefs, but *not all consistent and coherent infinite hierarchies are in any finite model*. It is not obvious that even in an infinite model that all such hierarchies of beliefs can be represented.

More on this later...

- What might the players' be thinking about?
- Do not confuse modeling with analyzing a game situation!
- Can the decision problem be *separated* from the game situation?
- Can a player assign subjective probabilities to strategies under the control of other players who have their own objectives?

Revealed Preference Theory: Preference vs. Preference*

Let X be a set and $P(X)$ the set of non-empty finite subsets (*menus*) of X.

Choice function: $C : P(X) \rightarrow P(X)$ where for all $Y \in P(X)$, $C(Y) \subseteq Y$

For $x, y \in X$, xVy iff there is a $S \in P(X)$ such that $x, y \in S$ and $x \in C(S)$.

Revealed Preference Theory: The Revelation Theorem

Theorem. V is a preference relation (complete, reflexive and transitive) iff C satisfies the weak axiom of revealed preference (WARP).

- α For all $A, B \in P(X)$, if $x \in A \subseteq B$ and $x \in C(B)$, then $x \in C(A)$.
- β For all $A, B \in P(X)$, if $A \subseteq B$, $x, y \in C(A)$ and $y \in C(B)$, then $x \in C(B)$

Applying revealed preference theory to game theory

D. Hausman. *Revealed Preference, Belief, and Game Theory*. Economics and Philosophy, 16:1, pgs. 99-115, 2000.

A. Lehtinen. *The Revealed-Preference Interpretation of Payoffs in Game Theory*. Homo Oeconomicus, 28:3, pgs. 265 - 296, 2011.

This can't be right...

"Modern utility theory makes tautology of the fact that action B will be chosen rather than A when the former yields a higher payoff by *defining* the payoff of B to be larger than the payoff of A if B is chosen when A is available." (Binmore, pg. 169)

K. Binmore. *Game Theory and the Social Contract: Playing Fair*. The MIT Press, 1994.

Reading the Normal Form

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

The numbers must represent the subjective preferences, not the revealed preferences. I.e., Ann *believes* that Bob will play L if he *believes* that she will play U not Ann knows that Bob will play L if she plays U.

Reading the Normal Form

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

The numbers must represent the subjective preferences, not the revealed preferences. I.e., Ann *believes* that Bob will play L if he *believes* that she will play U not Ann knows that Bob will play L if she plays U.

Questions about how to play games should be sharply separated from questions about what games people are playing.

What's in a game?

"We adhere to the classical point of view that the game under consideration fully describes the real situation — that any (pre) commitment possibilities, any repetitive aspect, any probabilities of error, or any possibility of jointly observing some random event, have already been modeled in the game tree." (Kohlberg and Mertens, pg. 1005)

E. Kohlberg and J.-F. Mertens. *On the strategic stability of equilibria*. Econometrica, 54, pgs. 1003 - 1038, 1986.

Modelling is hard

"Modelling requires intuition, common sense, and empirical data in order to determine the relevant factors entering into the players' strategic considerations and should thus be included in the model. This requirement makes the application of game theory more of an art than a mechanical algorithm." (Rubinstein, pg. 919)

A. Rubinstein. *Interpretations of Game Theory*. Econometrica, 59:4, 1991, pgs. 901 - 924.

Three games

- Prisoner's dilemma
- Ultimatum game
- Dictator game

Prisoner's Dilemma

Two people commit a crime.

Prisoner's Dilemma

Two people commit a crime. The 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. The 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. The 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. The 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. The 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: Confess (C), Don't Confess (D)

Prisoner's Dilemma

Two options: Confess (C), Don't Confess (D)

Possible outcomes:

Prisoner's Dilemma

Two options: Confess (C), Don't Confess (D)

Possible outcomes: We both confess (C, C),

Prisoner's Dilemma

Two options: Confess (C), Don't Confess (D)

Possible outcomes: We both confess (C, C), I confess but my partner doesn't (C, D),

Prisoner's Dilemma

Two options: Confess (C), Don't Confess (D)

Possible outcomes: We both confess (C, C), I confess but my partner doesn't (C, D), My partner confesses but I don't (D, C),

Prisoner's Dilemma

Two options: Confess (C), Don't Confess (D)

Possible outcomes: We both confess (C, C), I confess but my partner doesn't (C, D), My partner confesses but I don't (D, C), neither of us confess (D, D).

Prisoner's Dilemma

		Bob	
		D	C
Ann	D		
	C		

Prisoner's Dilemma

		Bob D	Bob C
Ann	D	3	1
Ann	C	4	2

Ann's preferences

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	3	4
	C	1	2

Bob's preferences

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

What should Ann (Bob) do?

Dominance Reasoning

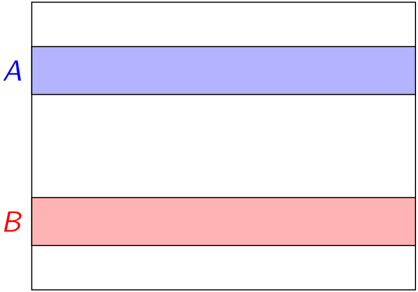

Dominance Reasoning

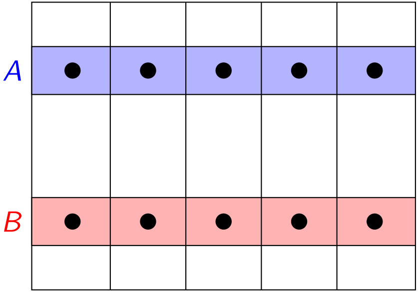

Dominance Reasoning

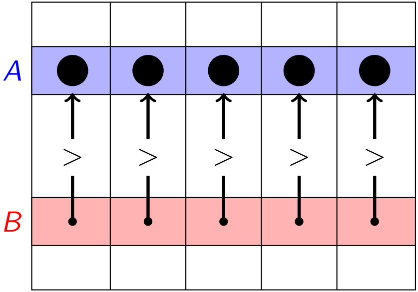

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

What should Ann (Bob) do?

Prisoner's Dilemma

		Bob D	Bob C
Ann	D	3,3	1,4
Ann	C	4,1	2,2

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

Prisoner's Dilemma

		Bob	
		D	*C*
Ann	*D*	3,3	1,4
	C	4,1	2,2

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

Prisoner's Dilemma

Bob

Ann		D	C
	D	3,3	1,4
	C	4,1	2,2

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

Prisoner's Dilemma

Bob

Ann

	D	C
D	3	2.5
C	2.5	2

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

Prisoner's Dilemma

Ann \ Bob	D	C
D	3,3	1,4
C	4,1	2,2

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

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

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

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	ϵ,ϵ	1,4
	C	4,1	2,2

What should Ann (Bob) do? *Change the game* (eg., Symbolic Utilities)

		Bob	
		D	C
Ann	D	4,4	1,3
	C	3,1	2,2

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	4,4	1,3
	C	3,1	2,2

Assurance Game

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

		Bob	
		D	C
Ann	D	4,4	1,3
	C	3,1	2,2

Assurance Game

What should/will Ann (Bob) do?

Nozick: Symbolic Utility

"Yet the symbolic value of an act is not determined solely by *that* act.

Nozick: Symbolic Utility

"Yet the symbolic value of an act is not determined solely by *that* act. The act's meaning can depend upon what other acts are available with what payoffs and what acts also are available to the other party or parties.

Nozick: Symbolic Utility

"Yet the symbolic value of an act is not determined solely by *that* act. The act's meaning can depend upon what other acts are available with what payoffs and what acts also are available to the other party or parties. What the act symbolizes is something it symbolizes when done in *that* particular situation, in preference to *those* particular alternatives.

Nozick: Symbolic Utility

"Yet the symbolic value of an act is not determined solely by *that* act. The act's meaning can depend upon what other acts are available with what payoffs and what acts also are available to the other party or parties. What the act symbolizes is something it symbolizes when done in *that* particular situation, in preference to *those* particular alternatives. If an act symbolizes "being a cooperative person," it will have that meaning not simply because it has the two possible payoffs it does

Nozick: Symbolic Utility

"Yet the symbolic value of an act is not determined solely by *that* act. The act's meaning can depend upon what other acts are available with what payoffs and what acts also are available to the other party or parties. What the act symbolizes is something it symbolizes when done in *that* particular situation, in preference to *those* particular alternatives. If an act symbolizes "being a cooperative person," it will have that meaning not simply because it has the two possible payoffs it does but also because it occupies a particular position within the two-person matrix — that is, being a dominated action that (when joined with the other person's dominated action) yield a higher payoff to each than does the combination of dominated actions. " (pg. 55)

R. Nozick. *The Nature of Rationality.* Princeton University Press, 1993.

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

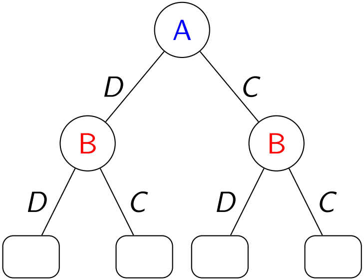

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

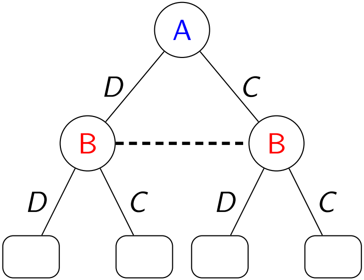

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

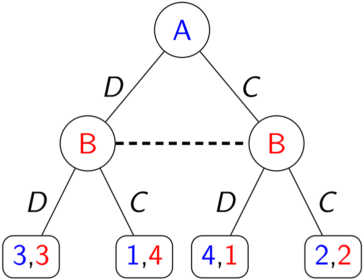

What should/will Ann (Bob) do?

		Bob	
		D	*C*
Ann	*D*	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

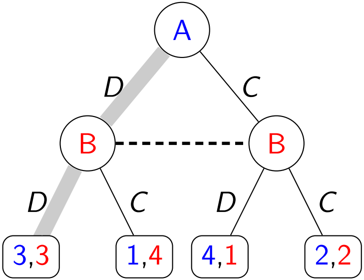

What should/will Ann (Bob) do?

		Bob	
		D	*C*
Ann	*D*	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

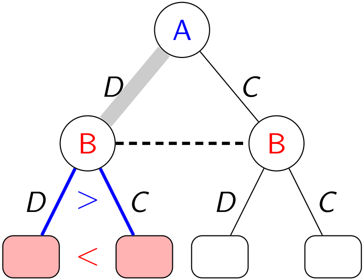

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

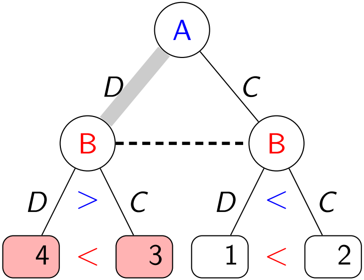

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

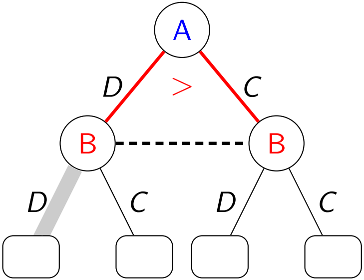

What should/will Ann (Bob) do?

		Bob	
		D	C
Ann	D	3,3	1,4
	C	4,1	2,2

Prisoner's Dilemma

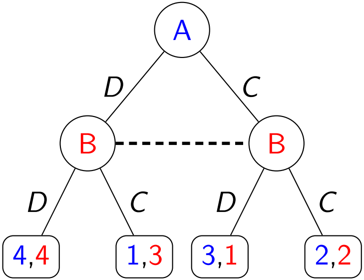

What should/will Ann (Bob) do?

“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.

Ultimatum Game

There is a good (say an amount of money) to be divided between two players.

Ultimatum Game

There is a good (say an amount of money) to be divided between two players. In order for either player to get the money, both players must agree to the division.

Ultimatum Game

There is a good (say an amount of money) to be divided between two players. In order for either player to get the money, both players must agree to the division. One player is selected by the experimenter to go first and is given all the money (call her the "Proposer"): the Proposer gives and ultimatum of the form "I get x percent and you get y percent — take it or leave it!".

Ultimatum Game

There is a good (say an amount of money) to be divided between two players. In order for either player to get the money, both players must agree to the division. One player is selected by the experimenter to go first and is given all the money (call her the "Proposer"): the Proposer gives and ultimatum of the form "I get x percent and you get y percent — take it or leave it!". No negotiation is allowed ($x + y$ must not exceed 100%).

Ultimatum Game

There is a good (say an amount of money) to be divided between two players. In order for either player to get the money, both players must agree to the division. One player is selected by the experimenter to go first and is given all the money (call her the "Proposer"): the Proposer gives and ultimatum of the form "I get x percent and you get y percent — take it or leave it!". No negotiation is allowed ($x + y$ must not exceed 100%). The second player is the Disposer: she either accepts or rejects the offer. If the Disposer rejects, then both players get 0 otherwise they get the proposed division.

Ultimatum Game

There is a good (say an amount of money) to be divided between two players. In order for either player to get the money, both players must agree to the division. One player is selected by the experimenter to go first and is given all the money (call her the "Proposer"): the Proposer gives and ultimatum of the form "I get x percent and you get y percent — take it or leave it!". No negotiation is allowed ($x + y$ must not exceed 100%). The second player is the Disposer: she either accepts or rejects the offer. If the Disposer rejects, then both players get 0 otherwise they get the proposed division.

Suppose the players meet only once. It would seem that the Proposer should propose 99% for herself and 1% for the Disposer. And if the Disposer is instrumentally rational, then she should accept the offer.

Ultimatum Game

But this is not what happens in experiments: if the Disposer is offered 1%, 10% or even 20%, the Disposer very often rejects. Furthermore, the proposer tends demand only around 60%.

Ultimatum Game

But this is not what happens in experiments: if the Disposer is offered 1%, 10% or even 20%, the Disposer very often rejects. Furthermore, the proposer tends demand only around 60%.

A typical explanation is that the players' utility functions are not simply about getting funds to best advance their goals, but about acting according to some norms of fair play.

Ultimatum Game

But this is not what happens in experiments: if the Disposer is offered 1%, 10% or even 20%, the Disposer very often rejects. Furthermore, the proposer tends demand only around 60%.

A typical explanation is that the players' utility functions are not simply about getting funds to best advance their goals, but about acting according to some norms of fair play. But acting according to norms of fair play does not seem to be a goal: it is a principle to which a person wishes to conform.

Dictator Game

Similar to the ultimatum game, there is a proposer and a second player. The proposer determines an allocation of some pot of money (say \$100). The second player simply receives the portion of the money from the proposer (i.e., the second player is completely passive).

Dictator Game

Similar to the ultimatum game, there is a proposer and a second player. The proposer determines an allocation of some pot of money (say \$100). The second player simply receives the portion of the money from the proposer (i.e., the second player is completely passive).

Proposers often allocate some money to the second player...

D. Kahneman, J. Knetsch, and R. Thaler. *Fairness And The Assumptions Of Economics.*. The Journal of Business, 59, pgs. 285- 300, 1986.

Can the decision problem be *separated* from the game situation?

Can the decision problem be *separated* from the game situation?

Are strategies merely neutral access routes to consequences?

Separability: Let G be any game, and let D be the problem that a given player in G would face, were the outcomes of the available strategies in G conditioned not bye the choices of another player but rather by some "natural" turn of events in the world, so that the player faces (in effect) a classic problem of individual decision making under conditions of risk or uncertainty. Suppose further that the player's expectation with regard to the conditioning events corresponds to the expectations held with regard to the choice that the other player will make in G. Then the first player's preference ordering over the options in G must correspond to that player's preference ordering over the options in D.

E. McClennen. *Rational choice in the context of ideal games*. in *Knowledge, Belief and Strategic Interaction*, pgs. 47-60, 1992.

utility must be measured *in the context of the game itself*.

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.

The following two outcomes are not equivalent:

- "I get \$90"
- "I get \$90 and choose to leave \$10 to my opponent"

The following two outcomes are not equivalent:

- "I get \$10 and player one gets \$90, and this was decided by Nature"
- "I get \$10, player one gets \$90 and this was decided by Player one".

Can a player assign subjective probabilities to strategies under the control of other players who have their own objectives?

M. Mariotti. *Is Bayesian Rationality Compatible with Strategic Rationality?*. The Economic Journal, 105: 432, pgs. 1099 - 1109, 1995.

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.

Games as consequences: "A decision maker prefers to be player i in game G_1 to being player j in game G_2"

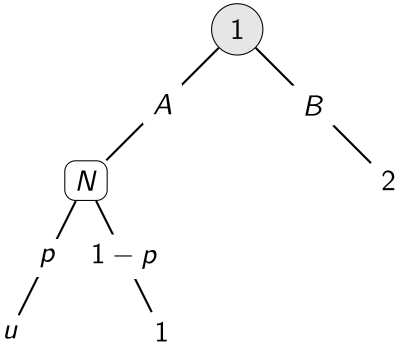
1
A
B
N
2
p
1 − p
u
1

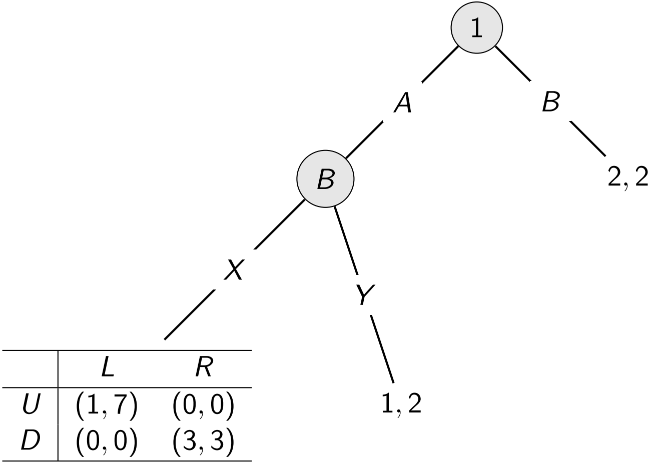

	L	R
U	$(1, 7)$	$(0, 0)$
D	$(0, 0)$	$(3, 3)$