Epistemic Game Theory

Lecture 1

ESSLLI'12, Opole

Eric Pacuit Olivier Roy

TiLPS, Tilburg University MCMP, LMU Munich
`ai.stanford.edu/~epacuit`
`http://olivier.amonbofis.net`

August 6, 2012

The Guessing Game

Plan for the week

Plan for the week

1. **Monday** Basic Concepts.
    - Basics of Game Theory.
    - The Epistemic View on Games.
    - Basics of Decision Theory

Plan for the week

1. **Monday** Basic Concepts.
2. **Tuesday** Epistemics.
    - Logical/qualitative models of beliefs, knowledge and higher-order attitudes.
    - Probabilistic/quantitative models of beliefs, knowledge and higher-order attitudes.

Plan for the week

1. **Monday** Basic Concepts.
2. **Tuesday** Epistemics.
3. **Wednesday** Fundamentals of Epistemic Game Theory.
    - Common knowledge of Rationality and iterated strict dominance in the matrix.
    - Common knowledge of Rationality and backward induction (strict dominance in the tree).

Plan for the week

1. **Monday** Basic Concepts.
2. **Tuesday** Epistemics.
3. **Wednesday** Fundamentals of Epistemic Game Theory.
4. **Thursday** Puzzles and Paradoxes.
    - Weak dominance and admissibility in the matrix.
    - Russell-style paradoxes in models of higher-order beliefs. (The Brandenburger-Kiesler paradox).

Plan for the week

1. **Monday** Basic Concepts.
2. **Tuesday** Epistemics.
3. **Wednesday** Fundamentals of Epistemic Game Theory.
4. **Thursday** Puzzles and Paradoxes.
5. **Friday** Extensions and New Directions.
    - Nash Equilibrium and mixted strategies.
    - Forward Induction.
    - Are the models normative or descriptive?
    - Theory of play.

Practicalities

- Course Website:
  - `ai.stanford.edu/~epacuit/esslli2012/epgmth.html`
- There you'll find handouts, reading material and additional references.
- In case of problem:
  - Olivier Roy: `Olivier.Roy@lmu.de`
  - Eric Pacuit: `E.J.Pacuit@uvt.nl`

Basics of Game Theory

Key Concepts

- Games in Strategic (matrix) and Extensive (tree) form.
- Strategies (pure and mixed).
- Solution Concepts: Iterated Strict Dominance, Iterated Weak Dominance, Nash Equilibrium,

The Matrix: games in strategic forms.

The Matrix: games in strategic forms.

		Alexei	
Strangelove			

Players,

The Matrix: games in strategic forms.

		Alexei	
		Disarm	Arm
Strangelove	Disarm		
	Arm		

Players, Actions or Strategies, Strategy profiles,

The Matrix: games in strategic forms.

		Alexei	
		Disarm	Arm
Strangelove	Disarm	3, 3	
	Arm		1, 1

Players, Actions or Strategies, Strategy profiles, Payoffs on profiles.

The Matrix: games in strategic forms.

		Alexei	
		Disarm	Arm
Strangelove	Disarm	3, 3	0, 4
	Arm	4, 0	1, 1

Players, Actions or Strategies, Strategy profiles, Payoffs on profiles.

A three players game

Fidel - D

		Alexei	
		D	A
Strglv	D	3, 3, 3	1, 4, 5
	A	4, 1, 1	2, 2, 2

Fidel - A

		Alexei	
		D	A
Strglv	D	3, 3, 2	1, 4, 4
	A	4, 1, 0	2, 2, 2

The Tree: games in extensive forms.

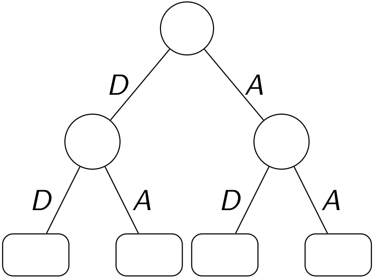

Actions,

The Tree: games in extensive forms.

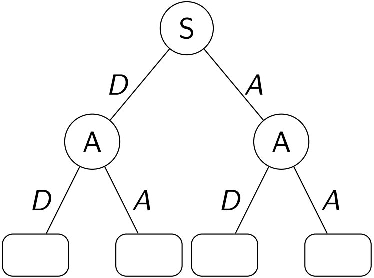

Actions, Players,

The Tree: games in extensive forms.

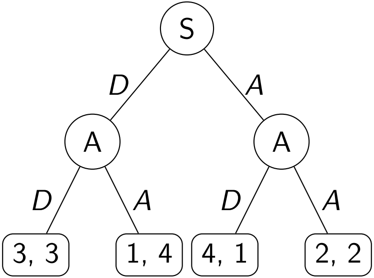

Actions, Players, Payoffs on leaves,

The Tree: games in extensive forms.

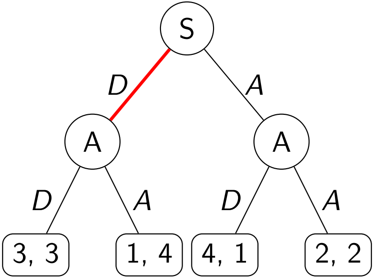

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

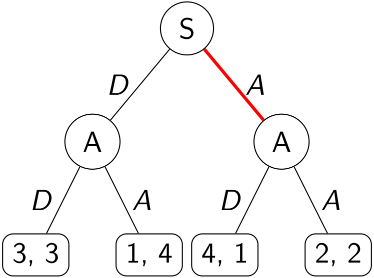

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

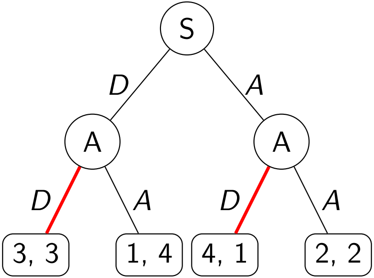

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

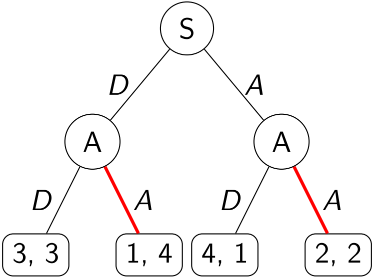

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

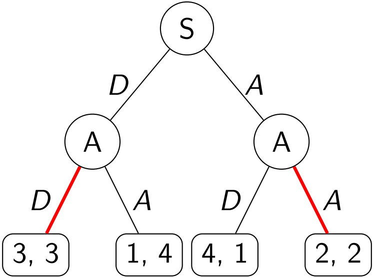

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

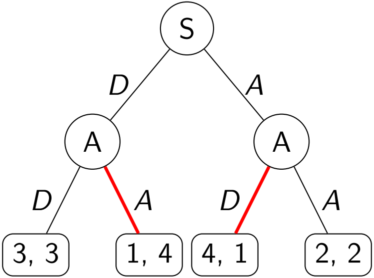

Actions, Players, Payoffs on leaves, Strategies

The Tree: games in extensive forms.

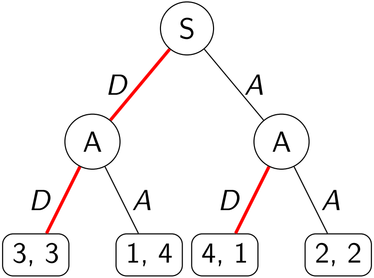

Actions, Players, Payoffs on leaves, Strategies

Extensive and strategic form games are related

		A	
		D	A
S	D	3, 3	1, 4
	A	4,1	2, 2

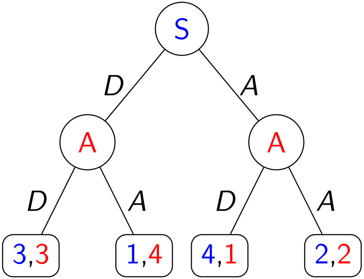

Extensive and strategic form games are related

		A	
		D	A
S	D	3, 3	1, 4
	A	4,1	2, 2

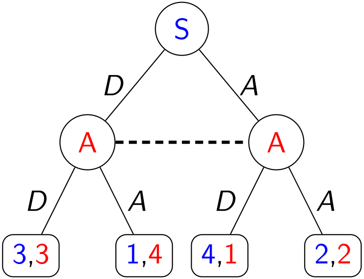

Some types of non-cooperative games of interest

- 2 players games.
- 2 players, zero-sum: if one player "wins" x then the other "looses" $-x$.
- 2 players, win-loose games.
- Perfect/imperfect information.

Pure and mixed strategies.

Pure and mixed strategies.

		Alexei	
		Head	Tail
Strangelove	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

Pure and mixed strategies.

		Alexei	
		Head	Tail
Strangelove	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

- Strangelove has two pure strategies: Head and Tail.

Pure and mixed strategies.

		Alexei	
		Head	Tail
Strangelove	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

- Strangelove has two pure strategies: Head and Tail.
- A mixed strategy is a probability distribution over the set of pure strategies. For instance:
  - (1/2 Head, 1/2 Tail)
  - (1/3 Head, 2/3 Tail)
  - ...

Pure and mixed strategies.

		Alexei	
		Head	Tail
Strangelove	Head	1, -1	-1, 1
	Tail	-1, 1	1, -1

- Strangelove has two pure strategies: Head and Tail.
- A mixed strategy is a probability distribution over the set of pure strategies. For instance:
  - (1/2 Head, 1/2 Tail)
  - (1/3 Head, 2/3 Tail)
  - ...
- Additional subtleties in extensive games. (mixing at a node vs mixing whole strategies).

Interpretation of mixed strategies

Interpretation of mixed strategies

1. Real randomizations:
    - Side of goal in penalty kicks.
    - Serving side in tennis.
    - Luggage check at the airport.

Interpretation of mixed strategies

1. Real randomizations:
    - Side of goal in penalty kicks.
    - Serving side in tennis.
    - Luggage check at the airport.
2. Epistemic interpretation:
    - Mixed strategies as beliefs of the other player(s) about what you do.

Solution Concepts

Solution Concepts

- Set of profiles or outcome of the game that are intuitively viewed as “rational”.

Solution Concepts

- Set of profiles or outcome of the game that are intuitively viewed as “rational”.
- Three well-known solution concepts in the matrix:
  - Nash Equilibrium.
  - Iterated elimitation of:
    - Strictly dominated strategies.
    - Weakly dominated strategies.

Solution Concepts

- Set of profiles or outcome of the game that are intuitively viewed as "rational".
- Three well-known solution concepts in the matrix:
    - Nash Equilibrium.
    - Iterated elimitation of:
        - Strictly dominated strategies.
        - Weakly dominated strategies.
- In the tree we will focus on one:
    - Backward induction.

Nash Equilibrium

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- The profile **aA** is a *Nash equilibrium* of that game.

Nash Equilibrium

	A	B
a	1, 1	0, 0
b	0, 0	1, 1

- The profile **aA** is a *Nash equilibrium* of that game.

Definition

A strategy profile σ is a *Nash equilibrium* iff for all i and all $s'_i \neq \sigma_i$:

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

Some Facts about Nash Equilibrium

- Nash equilibria in Pure Strategies do not always exist.
- Every game in strategic form has a Nash equilibrium in mixed strategies.
    - The proof of this make use of Kakutani's Fixed point thm.
- Some games have multiple Nash equilibria.

von Neumann's minimax theorem

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:

$$\begin{aligned} v &= \max_{p\in\Delta(S_1)} \min_{q\in\Delta(S_2)} u_1(s_1, s_2) \\ &= \min_{q\in\Delta(S_2)} \max_{p\in\Delta(S_1)} u_1(s_1, s_2) \end{aligned}$$

Furthermore, a mixed strategy profile (s_1, s_2) is a Nash equilibrium if and only if

$$s_1 \in \mathrm{argmax}_{p\in\Delta(S_1)} \min_{q\in\Delta(S_2)} u_1(p, q)$$

$$s_2 \in \mathrm{argmax}_{q\in\Delta(S_2)} \min_{p\in\Delta(S_1)} u_1(p, q)$$

Finally, for all mixed Nash equilibria (p, q), $u_1(p, q) = v$

Strictly Dominated Strategies

Strictly Dominated Strategies

		A	
		D	A
S	D	3, 3	1, 4
	A	4,1	2, 2

Strictly Dominated Strategies

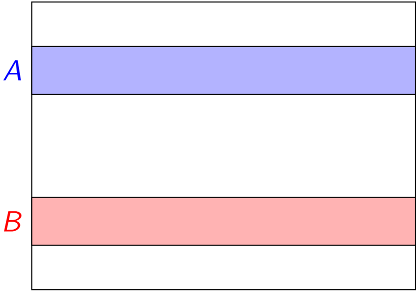

Strictly Dominated Strategies

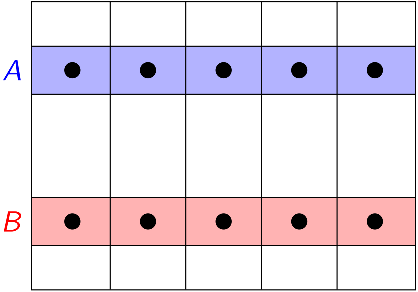

Strictly Dominated Strategies

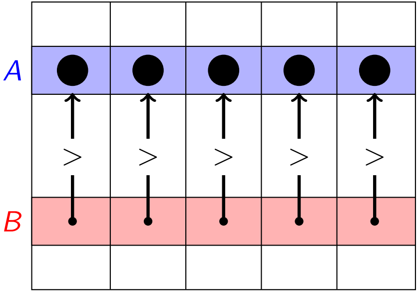

Strictly Dominated Strategies

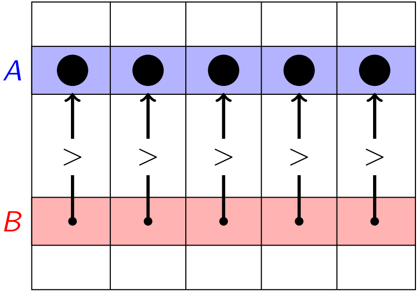

In general, the idea applies to both mixed and pure strategies.

Iterated Elimination of Strictly Dominated Strategies

Bob

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

Iterated Elimination of Strictly Dominated Strategies

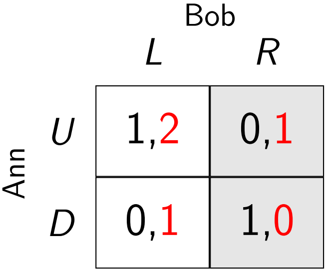

		Bob	
		L	R
Ann	U	1,2	0,1
	D	0,1	1,0

Iterated Elimination of Strictly Dominated Strategies

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

Iterated Elimination of Strictly Dominated Strategies

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

Facts about IESDS

- *The algorithm always terminates on finite games*. Intuition: this is a decreasing (in fact, monotonic) function on sub-games. It thus has a fixed-point by the Knaster-Tarski thm.
- *The algorithm is order independent*: One can eliminate SDS one player at the time, in difference order, or all simultaneously. The fixed-point of the elimination procedure will always be the same.
- *All Nash equilibria survive IESDS*. But not all profile that survive IESDS are Nash equilibria.

Weak Dominance

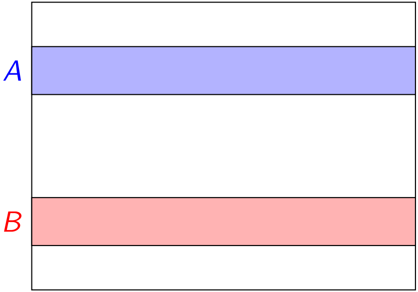

Weak Dominance

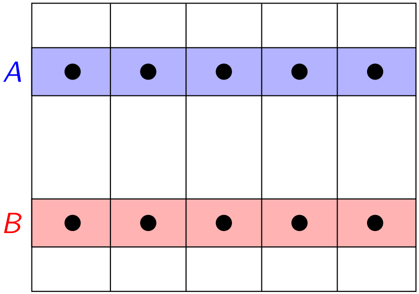

Weak Dominance

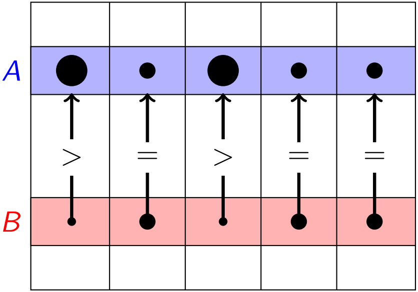

Weak Dominance

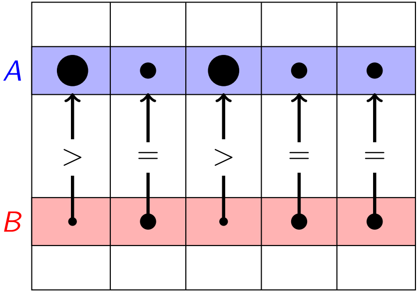

- All strictly dominated strategies are weakly dominated.

Iterated Elimination of Weakly Dominated Strategies

Bob

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

Iterated Elimination of Weakly Dominated Strategies

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

Iterated Elimination of Weakly Dominated Strategies

Bob

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

Iterated Elimination of Weakly Dominated Strategies

Bob

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

Facts about IEWDS

- *The algorithm always terminates on finite games*.
- *The algorithm is order dependent!*: Eliminating simultaneously all WDS at each round need not to lead to the same result as eliminating only some of them.
- *Not all Nash equilibria survive IESDS*.

> *Hey, 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 its an important concept, but its not the most basic concept. The most basic concept should be: to maximise your utility given your information. Its in a game just like in any other situation. Maximise your utility given your information!*
>
> *Robert Aumann, 5 Questions on Epistemic Logic, 2010*

> *Hey, 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 its an important concept, but its not the most basic concept. The most basic concept should be: to maximise your utility given your information. Its in a game just like in any other situation. Maximise your utility given your information!*
>
> *Robert Aumann, 5 Questions on Epistemic Logic, 2010*

Two views on games:

> *Hey, 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 its an important concept, but its not the most basic concept. The most basic concept should be: to maximise your utility given your information. Its in a game just like in any other situation. Maximise your utility given your information!*
>
> *Robert Aumann, 5 Questions on Epistemic Logic, 2010*

Two views on games:

- Based on solution Concepts.

> *Hey, 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 its an important concept, but its not the most basic concept. The most basic concept should be: to maximise your utility given your information. Its in a game just like in any other situation. Maximise your utility given your information!*
>
> *Robert Aumann, 5 Questions on Epistemic Logic, 2010*

Two views on games:

- Based on solution Concepts.
- Classical, decision-theoretic.

Component of a Game

A game in strategic form:

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

A coordination game:

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

$$G = \langle \mathrm{Ag}, \{(S_i, \pi_i)_{i\in \mathrm{Ag}}\}\rangle$$

- Ag is a finite set of agents.
- S_i is a finite set of strategies, one for each agent $i \in \mathrm{Ag}$.
- $u_i : \Pi_{i\in \mathrm{Ag}} S_i \longrightarrow \mathbb{R}$ is a payoff function defined on the set of outcomes of the game.

Solutions/recommendations: Nash Equilibrium, Elimination of strictly dominated strategies, of weakly dominated strategies...

A Decision Problem: Leonard's Omelette

	Egg Good	Egg Rotten
Break with other eggs	4	0
Separate bowl	2	1

A Decision Problem: Leonard's Omelette

	Egg Good	Egg Rotten
Break with other eggs	4	0
Separate bowl	2	1

- Agent, actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

	Egg Good	Egg Rotten
Break with other eggs	4	0
Separate bowl	2	1

- Agent, actions, states, payoffs, beliefs.
- Ex.: Leonard's beliefs: $p_L(EG) = 1/2$, $p_L(ER) = 1/2$.

A Decision Problem: Leonard's Omelette

	Egg Good	Egg Rotten
Break with other eggs	4	0
Separate bowl	2	1

- Agent, actions, states, payoffs, beliefs.
- Ex.: Leonard's beliefs: $p_L(EG) = 1/2$, $p_L(ER) = 1/2$.
- Solution/recommendations: choice rules. Maximization of Expected Utility, Dominance, Minmax...

The Epistemic or Bayesian View on Games

- Traditional game theory:
  Actions, outcomes, preferences, **solution concepts**.
- Decision theory:
  Actions, outcomes, preferences **beliefs, choice rules**.

The Epistemic or Bayesian View on Games

- Traditional game theory:
  Actions, outcomes, preferences, solution concepts.
- Decision theory:
  Actions, outcomes, preferences **beliefs, choice rules**.
- Epistemic game theory:
  Actions, outcomes, preferences, **beliefs**, choice rules.

The Epistemic or Bayesian View on Games

- Traditional game theory:
  Actions, outcomes, preferences, solution concepts.
- Decision theory:
  Actions, outcomes, preferences beliefs, choice rules.
- Epistemic game theory:
  $:=$ (interactive) decision problem and choice rule $+$ **higher-order information**.

Basics of Decision Theory

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.

A Decision Problem: Leonard's Omelette

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

- Actions, states, payoffs, beliefs.
- Solution/recommendations: choice rules.
    - Which choice rule is normatively or descriptively appropriate depends on what kind of information are at the agent's disposal, and what kind of attitude she has.

Decision Under Risk

When the agent has probabilistic beliefs, or that her beliefs can be represented probabilistically.

u_i	P	$\neg P$
A	4	0
B	2	1

p_i	P	$\neg P$
A	1/8	3/8
B	1/8	3/8

Expected Utility: Given an agent's beliefs and desires, the **expected utility** of an **action** leading to a set of outcomes *Out* is:

$$\sum_{o \in Out} [\text{ subjective prob. of } o] \times [\text{utility of } o]$$

Why don't we just give our best guess of wet or dry?
Often people want to make a decision, such as whether to put out their washing to dry, and would like us to give a simple yes or no. However, this is often a simplification of the complexities of the forecast and may not be accurate.

Why don't we just give our best guess of wet or dry?

Often people want to make a decision, such as whether to put out their washing to dry, and would like us to give a simple yes or no. However, this is often a simplification of the complexities of the forecast and may not be accurate. By giving PoP we give a more honest opinion of the risk and allow you to make a decision depending on how much it matters to you.

Why don't we just give our best guess of wet or dry?

Often people want to make a decision, such as whether to put out their washing to dry, and would like us to give a simple yes or no. However, this is often a simplification of the complexities of the forecast and may not be accurate. By giving PoP we give a more honest opinion of the risk and allow you to make a decision depending on how much it matters to you. For example, if you are just hanging out your sheets that you need next week you might take the risk at 40% probability of precipitation, whereas if you are drying your best shirt that you need for an important dinner this evening then you might not hang it out at more than 10% probability.

Why don't we just give our best guess of wet or dry?
Often people want to make a decision, such as whether to put out their washing to dry, and would like us to give a simple yes or no. However, this is often a simplification of the complexities of the forecast and may not be accurate. By giving PoP we give a more honest opinion of the risk and allow you to make a decision depending on how much it matters to you. For example, if you are just hanging out your sheets that you need next week you might take the risk at 40% probability of precipitation, whereas if you are drying your best shirt that you need for an important dinner this evening then you might not hang it out at more than 10% probability. PoP allows you to make the decisions that matter to you.

http://www.metoffice.gov.uk/news/in-depth/science-behind-probability-of-precipitation

Maximization of Expected Utility

Let $DP = \langle S, O, u, p\rangle$ be a decision problem. S is a finite set of states and O a set of outcomes. An action $a : S \longrightarrow O$ is a function from states to outcomes, u_i a real-valued utility function on O, and p_i a probability measure over S. The **expected utility** of $a \in A$ with respect to p_i is defined as follows:

$$EU_p(a) \ := \ \Sigma_{s\in S}\, p(s)u(a(s))$$

An action $a \in A$ **maximizes expected utility** with respect to p_i provided for all $a' \in A$, $EU_p(a) \geq EU_p(a')$. In such a case, we also say a is a **best response** to p in game DP.

Decision under Ignorance

What to do when the agent cannot assign probabilities states? Or when we can't represent his beliefs probabilistically? Many alternatives proposed:

- Dominance Reasoning
- Admissibility
- Minimax
- ...

Dominance Reasoning

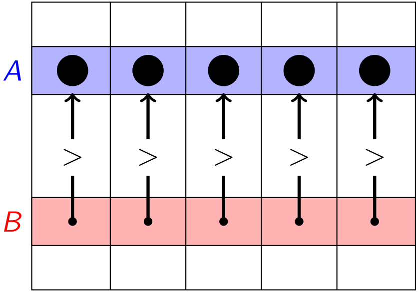

Some facts about strict dominance

- **Strict dominance is downward monotonic**: If a_i is strictly dominated with respect to $X \subseteq S$ and $X' \subseteq X$, then a_i is strictly dominated with respect to X'.

Some facts about strict dominance

- **Strict dominance is downward monotonic**: If a_i is strictly dominated with respect to $X \subseteq S$ and $X' \subseteq X$, then a_i is strictly dominated with respect to X'.
    - Intuition: the condition of being strictly dominated can be written down in a first-order formula of the form $\forall x \varphi(x)$, where $\varphi(x)$ is quantifier-free. Such formulas are downward monotonic: If $\mathcal{M}, s \models \forall x \varphi(x)$ and $\mathcal{M}' \subseteq \mathcal{M}$ then $\mathcal{M}', s \models \forall x \varphi(x)$

Some facts about strict dominance

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

Some facts about admissibility

- **Admissibility is NOT downward monotonic**: If a_i is not admissible with respect to $X \subseteq S$ and $X' \subseteq X$, it can be that a_i is admissible with respect to X'.

Some facts about admissibility

- **Admissibility is NOT downward monotonic**: If a_i is not admissible with respect to $X \subseteq S$ and $X' \subseteq X$, it can be that a_i is admissible with respect to X'.
    - Intuition: the condition of being inadmissible can be written down in a first-order formula of the form $\forall x \varphi(x) \wedge \exists x \psi(x)$, where $\varphi(x)$ and $\psi(x)$ are quantifier-free. The existential quantifier breaks the downward monotonicity.

Some facts about admissibility

- **Relation with MEU**:
  Suppose that $G = \langle N, \{S_i\}_{i \in N}, \{u_i\}_{i \in N} \rangle$ is a strategic game. A strategy $s_i \in S_i$ is weakly dominated (possibly by a mixed strategy) with respect to $X \subseteq S_{-i}$ iff there is **no full support probability measure** $p \in \Delta^{>0}(X)$ such that s_i is a best response with respect to p.

Road Map again

1. **Today** Basic Concepts.
    - Basics of Game Theory.
    - The Epistemic View on Games.
    - Basics of Decision Theory

Road Map again

1. **Today** Basic Concepts.
    - Basics of Game Theory.
    - The Epistemic View on Games.
    - Basics of Decision Theory
2. **Tomorrow** Epistemics.
    - Logical/qualitative models of beliefs, knowledge and higher-order attitudes.
    - Probabilistic/quantitative models of beliefs, knowledge and higher-order attitudes.

Strategic Games

Definition

A game in strategic form $\mathbb{G}$ is a tuple $\langle \mathcal{A}, S_i, u_i \rangle$ such that :

- $\mathcal{A}$ is a finite set of agents.
- S_i is a finite set of *actions* or *strategies* for i. A *strategy profile* $\sigma \in \Pi_{i \in \mathcal{A}} S_i$ is a vector of strategies, one for each agent in I. The strategy s_i which i plays in the profile σ is noted σ_i.
- $u_i : \Pi_{i \in \mathcal{A}} S_i \longrightarrow \mathbb{R}$ is an *utility function* that assigns to every strategy profile $\sigma \in \Pi_{i \in \mathcal{A}} S_i$ the utility valuation of that profile for agent i.

Extensive form games

Definition

A *game in extensive form* $\mathcal{T}$ is a tuple $\langle I, T, \tau, \{u_i\}_{i \in I} \rangle$ such that:

- T is finite set of finite sequences of *actions*, called *histories*, such that:
  - The empty sequence $\emptyset$, the *root* of the tree, is in T.
  - T is prefix-closed: if $(a_1, \ldots, a_n, a_{n+1}) \in T$ then $(a_1, \ldots, a_n) \in T$.
- A history h is *terminal in T* whenever it is the sub-sequence of no other history $h' \in T$. Z denotes the set of terminal histories in T.
- $\tau : (T - Z) \longrightarrow I$ is a *turn function* which assigns to every non-terminal history h the player whose turn it is to play at h.
- $u_i : Z \longrightarrow \mathbb{R}$ is a *payoff function* for player i which assigns i's payoff at each terminal history.

Strategies

Definition

- A *strategy* s_i for agent i is a function that gives, for every history h such that $i = \tau(h)$, an action $a \in A(h)$. S_i is the set of strategies for agent i.
- A *strategy profile* $\sigma \in \Pi_{i \in I} S_i$ is a combination of strategies, one for each agent, and $\sigma(h)$ is a shorthand for the action a such that $a = \sigma_i(h)$ for the agent i whose turn it is at h.
- A history h' is *reachable* or *not excluded* by the profile σ from h if $h' = (h, \sigma(h), \sigma(h, \sigma(h)), ...)$ for some finite number of application of σ.
- We denote $u_i^h(\sigma)$ the value of $util_i$ at the unique terminal history reachable from h by the profile σ.

Nash Equilibrium - General Definition

Definition

A profile of mixed strategy σ is a *Nash equilibrium* iff for all i and all mixed strategy $\sigma_i' \neq \sigma_i$:

$$EU_i(\sigma_i, \sigma_{-i}) \geq EU_i(\sigma_i', \sigma_{-i})$$

Where EU_i, the *expected utility of the strategy* σ_i against σ_{-i} is calculated as follows ($\sigma = (\sigma_i, \sigma_{-i})$):

$$EU_i(\sigma) = \Sigma_{s \in \Pi_j S_j}\left((\Pi_{j \in \mathrm{Ag}} \sigma_j(s_j)) u_i(s)\right)$$