Signalling Games and Pragmatics
Day II

Anton Benz
University of Southern Denmark,
IFKI, Kolding

The Course

- Day I: Introduction: From Grice to Lewis
- **Day II: Basics of Game and Decision Theory**
- Day III: Two Theories of Implicatures (Parikh, Jäger)
- Day IV: Best Answer Approach
- Day V: Utility and Relevance

Overview Day I

Introduction: From Grice to Lewis

- Gricean Pragmatics
  - General assumptions about conversation
  - Conversational implicatures
- Game and Decision Theory
- Lewis on Conventions
  - Examples of Conventions
  - Signalling conventions
  - Meaning in Signalling systems

Basics of Game and Decision Theory

Day 2 – August, 8th

Overview

- Elements of Decision Theory
  - Relevance as Informativity (Merin);
  - Relevance as Expected Utility (van Rooij).
- Game Theory
  - Strategic games in normal form
  - Equilibrium concepts
  - Games in extensive form
  - Signalling games
  - Application: Resolving Ambiguities (P. Parikh)

Game and Decision Theory

- **Decision theory:** Concerned with decisions of **individual** agents
- **Game theory:** Concerned with interdependent decisions of **several** agents.

Elements of Decision Theory

With application to measures of relevance

Decision Situations

- Take an umbrella with you when leaving the house.
- Choose between several candidates for a job.
- Decide where to look for a book which you want to buy.

A Classification of Decision Situations

One distinguishes between decision under:

- **Certainty:** The decision maker knows the outcome of each action with certainty.
- **Risk:** The decision maker knows of each outcome that it occurs with a certain probability.
- **Uncertainty:** No probabilities for outcomes of actions are known to the decision maker.

- We are only concerned with decisions under certainty or risk.
- Decisions may become risky because the decision maker does not know the true state of affairs.
- He may have expectations about the state of affairs.
- Expectations are standardly represented as probabilistic knowledge about a set of possible worlds.

Discrete Probability Space

A discrete probability space consists of:

- Ω: at most countable set.
- $P : \Omega \to [0, 1]$ a function such that

$$\sum_{v \in \Omega} P(v) = 1.$$

- Notation: $P(A) := \sum_{v \in A} P(v)$ for $A \subseteq \Omega$.

Representation of Decision Problem

A **decision problem** is a triple $((\Omega, P), A, u)$ such that:

- (Ω, P) is a discrete probability space,
- A a finite, non–empty set of **actions**.
- $u : \Omega \times A \to R$ a real valued function.

- A is called the action set, and its elements actions.
- u is called a payoff or utility function.

Taking an Umbrella with you

- Worlds:
  - w: rainy day.
  - v: cloudy but dry weather.
  - u: sunny day.
- Probabilities:
  - P(w)=1/3; P(v)=1/6; p(u)=1/2
- Actions:
  - a: taking umbrella with you; b: taking no umbrella.
- Utilities:
  - rainy day: u(w,a) = 1, u(w,b) = -1.
  - cloudy day: u(v,a) = -0.1, u(w,b) = 0.
  - sunny day: u(u,a) = -0.1, u(u,b) = 0..

Learning

How are expectations change by new information?

Example:

1. Before John looked out of window:

   $P(cloudy \wedge will\text{-}rain) = 1/3$; $P(cloudy) = 1/2$.

2. Looking out of window John learns that it is cloudy.
   - What is the new probability of will-rain?

Conditional Probabilities

- Let (Ω, P) be a discrete probability space representing **expectations prior** to new observation A.
- For any hypothesis H the **conditional probability** is defined as:

$$P(H|A) = P(H \cap A)/P(A) \text{ for } P(A)>0$$

Example:

1. Before John looked out of window:

$P(\text{cloudy} \wedge \text{will-rain}) = 1/3$; $P(\text{cloudy}) = 1/2$.

2. John learns that it is cloudy. The posterior probability P^+ is defined as:

$$P^+(\text{will-rain}) := P(\text{will-rain}|\text{cloudy})$$
$$= P(\text{will-rain} \cap \text{cloudy}) / P(\text{cloudy})$$
$$= 1/3 : 1/2 = 2/3$$

Relevance as Informativity

(Arthur Merin)

The Argumentative view

- Speaker tries to persuade the hearer of a hypothesis H.
- Hearers expectations given by (Ω, P).
- Hearer's decision problem:

- $\mathcal{A} = \{H, \bar{H}\}$ with $\bar{H} = \Omega \setminus H$;
- $u(v, H) = \begin{cases} 1 \textit{ iff } v \in H \\ 0 \textit{ iff } v \notin H \end{cases}$ and $u(v, \bar{H}) = \begin{cases} 1 \textit{ iff } v \in \bar{H} \\ 0 \textit{ iff } v \notin \bar{H} \end{cases}$

Example

If Eve has an interview for a job she wants to get, then

- her goal is to convince the interviewer that she is qualified for the job (H).
- Whatever she says is the more relevant the more it favours H and disfavours the opposite proposition.

Measuring the Update Potential of an Assertion A.

- Hearer's inclination to believe H prior to learning A:

$$P(H)/P(H^{-})$$

- Inclination to believe H after learning A:

$$P^{+}(H)/P^{+}(H^{-}) = P(H|A)/P(H^{-}|A) =$$

$$= P(H)/P(H^{-}) \times P(A|H)/P(A|H^{-})$$

Using log (just a trick!) we get:

$\log P^{+}(H)/P^{+}(H^{-}) = \log P(H)/P(H^{-}) + \log P(A|H)/P(A|H^{-})$

New = Old + update

$\log P(A|H)/P(A|H^{-})$ can be seen as the update potential of proposition A with respect to H.

Relevance (Merin)

Intuitively: A proposition A is the more relevant to a hypothesis H the more it increases the inclination to believe H.

$$r_H(A) := \log P(A|H)/P(A|\bar{H})$$

- It is $r_{H^-}(A) = - r_H(A)$;
- If $r_H(A) = 0$, then A does not change the prior expectations about H.

Relevance as Expected Utility

(Robert van Rooij)

An Example (Job interview)

v_1: Eve has ample of job experience and can take up a responsible position immediately.

v_2: Eve has done an internship and acquired there job relevant qualifications but needs some time to take over responsibility.

v_3: Eve has done an internship but acquired no relevant qualifications and needs heavy training before she can start on the job.

v_4: Eve has just finished university and needs extensive training.

- Interviewer's decision problem:
  - a_1: Employ Eve.
  - a_2: Don't employ Eve.

u	v_1	v_2	v_3	v_4
a_1	10	3	−1	−5
a_2	0	0	0	0

All worlds equally probable

How to decide the decision problem?

Expected Utility

Given a decision problem ((Ω, P),A,u), the expected utility of an action *a* is:

$$EU(a) = \sum_{v \in \Omega} P(v) \cdot u(v, a)$$

In our Example

$$EU(a_1) = \frac{1}{4} \cdot 10 + \frac{1}{4} \cdot 3 + \frac{1}{4} \cdot (-1) + \frac{1}{4} \cdot (-5) = 1\frac{3}{4}$$

$$EU(a_2) = \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot 0 + \frac{1}{4} \cdot 0 = 0.$$

Decision Criterion

- It is assumed that rational agents are **Bayesian utility maximisers**.
- If an agent chooses an action, then the action's expected utility must be maximal.

In our example: As $EU(a_1) > EU(a_2)$ it follows that the interviewer will employ Eve.

The Effect of Learning

- If an agent learns that A, how does this change expected utilities?

$$EU(a|A) = \sum_{v \in \Omega} P(v|A) \cdot u(v, a).$$

Our Example

- What happens if the interviewer learns that Eve did an internship ($A=\{v_1,v_2\}$)?

$$EU(a_1|A) = \sum_{v \in \Omega} P(v|A) \cdot u(v, a_1) = \frac{1}{2} \cdot 3 + \frac{1}{2} \cdot (-1) = 1.$$

- Similarly, we find $EU(a_2|A) = 0$.
- ➢ The interviewer will decide not to employ Eve.

Measures of Relevance I
(van Rooij)

(Sample Value of Information)

New information A is relevant if

- it leads to a different choice of action, and
- it is the more relevant the more it increases thereby expected utility.

Measures of Relevance I
(van Rooij)

(Sample Value of Information)

- Let ((Ω, P),A,u) be a given decision problem.
- Let a* be the action with maximal expected utility before learning *A*.

Utility Value or **Relevance** of *A*:

$$UV(A) = \max_{a \in \mathcal{A}} EU(a|A) - EU(a^*|A).$$

Measures of Relevance II (van Rooij)

New information A is relevant if

- it increases expected utility.
- it is the more relevant the more it increases it.

$$UV'(A) = \max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a).$$

Measures of Relevance III (van Rooij)

New information A is relevant if

- it changes expected utility.
- it is the more relevant the more it changes it.

$$UV''(A) = \left| \max_{a \in \mathcal{A}} EU(a|A) - \max_{a \in \mathcal{A}} EU(a) \right|.$$

Application
(van Rooij)

Somewhere in the streets of Amsterdam...

1. J: Where can I buy an Italian newspaper?
2. E: At the station and at the Palace but nowhere else. (S)
3. E: At the station. (A) / At the Palace. (B)

The answers S, A and B are equally useful with respect to conveyed information and the inquirer's goals.

Game Theory

Overview

- Strategic games in normal form
- Equilibrium concepts
- Games in extensive form
- Signalling games
- Application: Resolving Ambiguities (Parikh)

Strategic games in normal form

- **Strategic games in normal form**
- Equilibrium concepts
- Games in extensive form
- Signalling games
- Application: Resolving Ambiguities

Basic distinctions in game theory

Static vs. dynamic games:

- **Static game:** In a static game every player performs only one action, and all actions are performed simultaneously.
- **Dynamic game:** In dynamic games there is at least one possibility to perform several actions in sequence.

Basic distinctions in game theory

Cooperative v.s. non–cooperative games

- **Cooperative:** In a cooperative game, players are free to make binding agreements in pre-play communications. Especially, this means that players can form coalitions.
- **Non–cooperative:** In non–cooperative games no binding agreements are possible and each player plays for himself.

Basic distinctions in game theory

Normal form vs. extensive form

- **Normal form:** Representation in matrix form.
- **Extensive form:** Representation in tree form. It is more suitable for dynamic games.

A strategic game in normal form

Components:

1. **Players:** games are played by players. If there are n players, then we represent them by the numbers 1, . . . , n.
2. **Action sets:** each player can choose from a set of actions. It may be different for different players. Hence, if there are n players, then there are n action sets $A_1, \ldots, A_n$.
3. **Payoffs:** each player has preferences over choices of actions. We represent the preferences by **payoff functions** u_i.

Representation of Strategic Games

A static game can be represented by a **payoff matrix**.

Column Player

Row Player

	a	b	*actions*
c	payoffs	payoffs	
d	payoffs	payoffs	
actions			

Representation of Strategic Games

In case of two–player games with two possible actions for each player:

	b_1	b_2
a_1	$u_1(a_1, b_1), u_2(a_1, b_1)$	$u_1(a_1, b_2), u_2(a_1, b_2)$
a_2	$u_1(a_2, b_1), u_2(a_2, b_1)$	$u_1(a_2, b_2), u_2(a_2, b_2)$

Row player's payoff

Column player's payoff

Prisoner’s dilemma

- Player: Two imprisoned criminals
- Actions: c *cooperate*; d *defect*

	c	d
c	$1, 1$	$-1, 2$
d	$2, -1$	$0, 0$

Prisoner’s dilemma

Battle of the sexes

- Player: A man (row) and a woman (column).
- Actions: b *go to boxing*; c *go to concert.*

	b	c
b	$2, 1$	$0, 0$
c	$0, 0$	$1, 2$

Battle of the sexes

Stag hunt

- Player: Two hunter.
- Actions: s *hunting stag*; r *hunting rabbit.*

	s	r
s	$2, 2$	$0, 1$
r	$1, 0$	$1, 1$

Stag hunt

Chicken

- Player: Two young guys.
- Actions: r *racing*; s *swerve.*

	r	s
r	$-1, -1$	$2, 0$
s	$0, 2$	$1, 1$

Chicken

Equilibrium concepts

- Strategic games in normal form
- **Equilibrium concepts**
- Games in extensive form
- Signalling games
- Application: Resolving Ambiguities

- Weak and strong dominance
- Nash equilibrium
- Pareto Optimality

Weak and Strong Dominance

- An action *a* of player i **strictly dominates** an action *b* iff the utility of playing *a* is strictly higher than the utility of playing *b* <u>whatever</u> actions the other players choose.
- An action *a* of player i **weakly dominates** an action *b* iff the utility of playing *a* is is at least as high as the utility of playing *b* <u>whatever</u> actions the other players choose.

Prisoner's dilemma

defect (*d*) **strictly dominates** all other actions:

	c	d
c	$1, 1$	$-1, 2$
d	$2, -1$	$0, 0$

Prisoner's dilemma

Nash equilibrium
(2 player)

An action pair (*a*,*b*) is a **weak Nash equilibrium** iff

1. there is no action *a'* such that

$$u_1(a',b) > u_1(a,b)$$

2. there is no action *b'* such that

$$u_2(a,b') > u_2(a,b)$$

Nash equilibrium
(2 player)

An action pair (a,b) is a **strong Nash equilibrium** iff

1. for all actions $a' \neq a$:

$$u_1(a',b) < u_1(a,b)$$

2. for all actions $b' \neq b$:

$$u_2(a,b') < u_2(a,b)$$

Battle of the sexes

- None of the actions is strictly dominating.
- Two strict Nash equilibria: (*b*,*b*), (*c*,*c*)

	b	c
b	$2, 1$	$0, 0$
c	$0, 0$	$1, 2$

Battle of the sexes

Pareto Nash equilibrium
(2 player)

An action pair (a,b) is a **Pareto Nash equilibrium** iff there is no other Nash equilibrium (a',b') such that 1. or 2. holds:

1. $u_1(a',b') > u_1(a,b)$ and $u_2(a',b') \geq u_2(a,b)$
2. $u_1(a',b') \geq u_1(a,b)$ and $u_2(a',b') > u_2(a,b)$

Stag hunt

- Two Nash equilibria: (*s*,*s*), (*r*,*r*).
- One Pareto Nash equilibrium: (*s*,*s*).

	s	r
s	$2, 2$	$0, 1$
r	$1, 0$	$1, 1$

Stag hunt

Games in extensive form

- Strategic games in normal form
- Equilibrium concepts
- **Games in extensive form**
- Signalling games
- Application: Resolving Ambiguities

A Tree

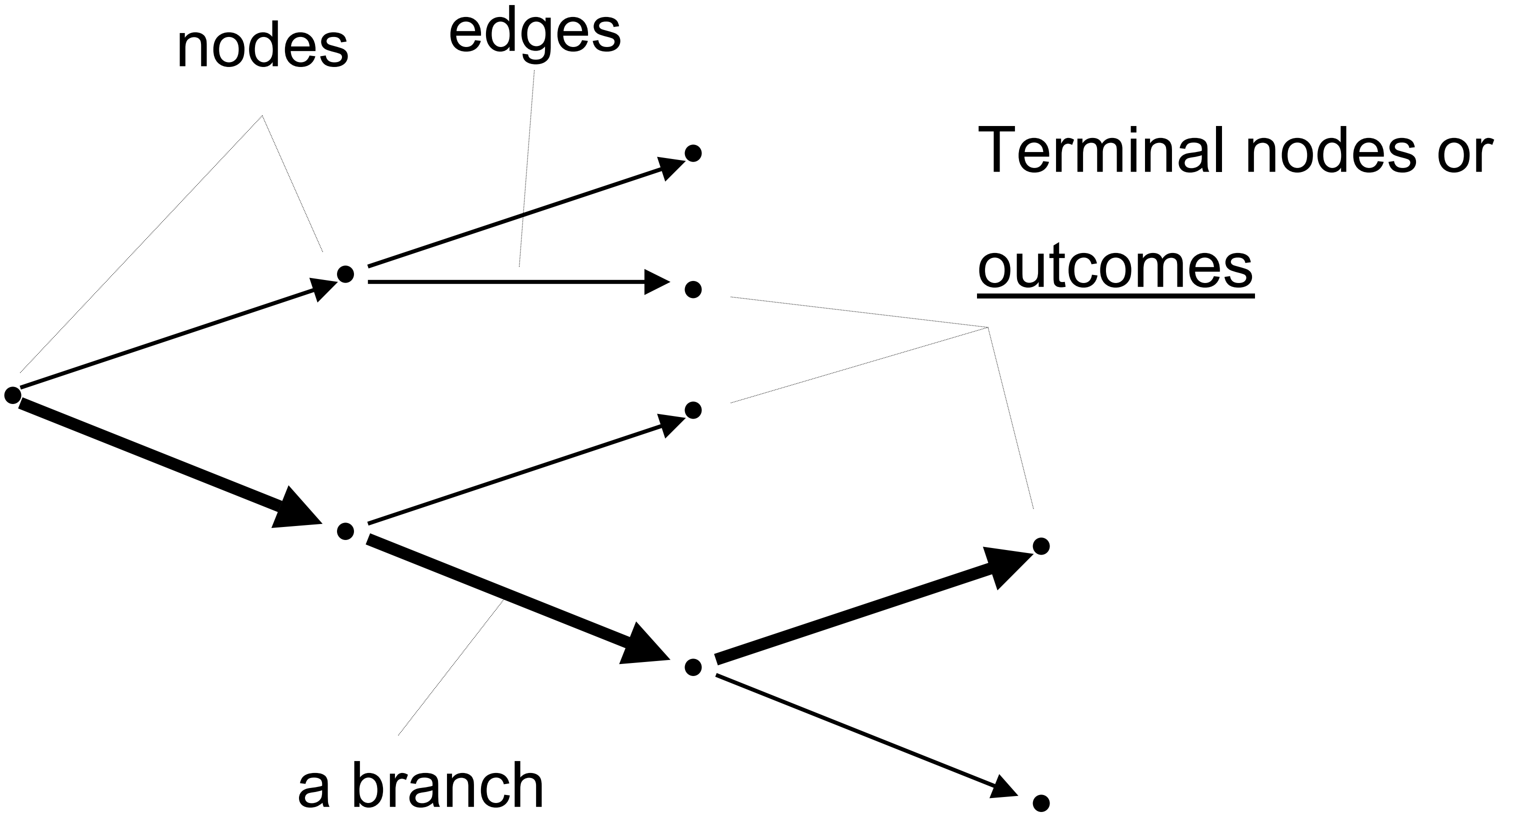

Components of a Game in Extensive Form

1. **Players:** $N = \{1,\ldots,n\}$ a set of n players.
   - Nature is a special player with number 0.
   - Each node in a game tree is assigned to a player.
2. **Moves:** Each edge in a game tree is labelled by an action.
3. **Information sets:** To each node n which is assigned to a player $i \in N$, a set of nodes is given which represents i's knowledge at n.

4. **Outcomes:** There is a set of *outcomes.* Each terminal node represents one outcome.
5. **Payoffs:** For each player $i \in N$ there exists a payoff (or utility) function u_i which assigns a real value to each of the outcomes.
6. Nodes assigned to 0 (Nature) are nodes where random moves can occur.

A Game Tree

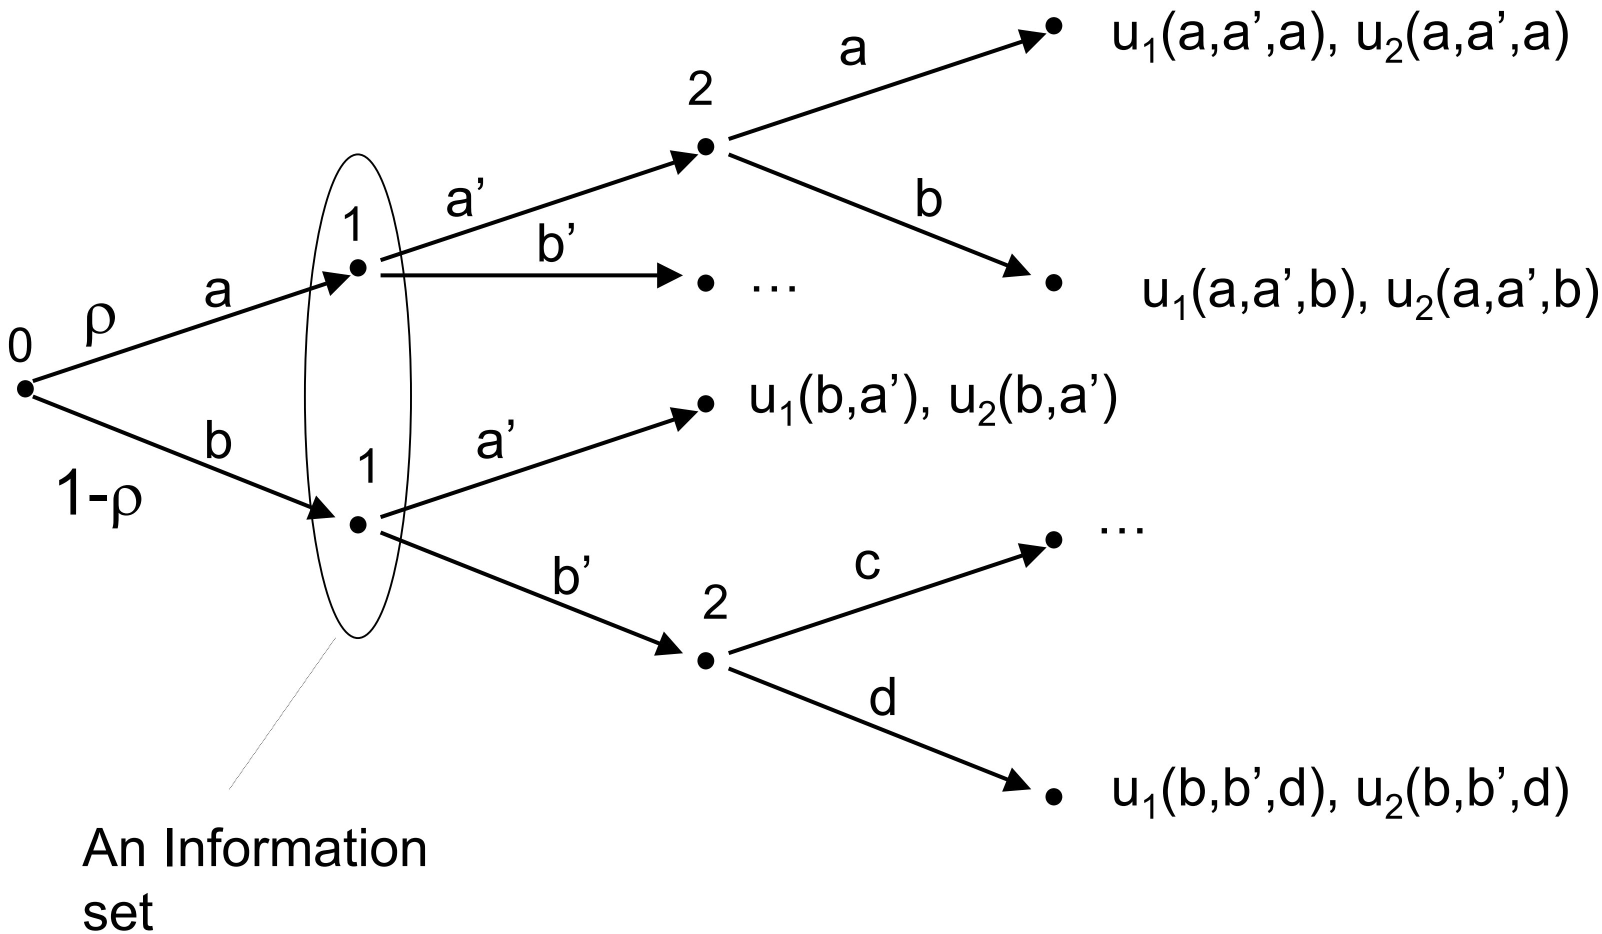

0
ρ
a
1-ρ
b
1
1
a'
b'
a'
b'
2
2
a
b
c
d
…
…
u1(a,a',a), u2(a,a',a)
u1(a,a',b), u2(a,a',b)
u1(b,a'), u2(b,a')
u1(b,b',d), u2(b,b',d)
An Information set

Signalling games

- Strategic games in normal form
- Equilibrium concepts
- Games in extensive form
- **Signalling games**
- Application: Resolving Ambiguities

- We consider only signalling games with two players:
  - a speaker S,
  - a hearer H.
- Signalling games are Bayesian games in extensive form; i.e. players may have private knowledge.

Private knowledge

- We consider only cases where the **speaker** has additional private knowledge.
- Whatever the **hearer** knows is common knowledge.
- The private knowledge of a player is called the player’s **type**.
- It is assumed that the hearer has certain expectations about the speaker’s type.

Signalling Game

A signalling game is a tuple:

$$\langle N, \Theta, p, (A_1, A_2), (u_1, u_2) \rangle$$

- N: Set of two players S,H.
- Θ: Set of types representing the speakers private information.
- p: A probability measure over Θ representing the hearer's expectations about the speaker's type.

- (A_1,A_2): the speaker's and hearer's action sets.
- (u_1,u_2): the speaker's and hearer's payoff functions with

$$u_i: A_1 \times A_2 \times \Theta \to \mathbf{R}$$

Playing a signalling game

1. At the root node a type is assigned to the speaker.
2. The game starts with a move by the speaker.
3. The speaker’s move is followed by a move by the hearer.
4. This ends the game.

Strategies in a Signalling Game

- Strategies are functions from the agents information sets into their action sets.
- The speaker's information set is identified with his type $\theta \in \Theta$.
- The hearer's information set is identified with the speaker's previous move $a \in A_1$.

$$S : \Theta \to A_1 \text{ and } H : A_1 \to A_2$$

Resolving Ambiguities

Prashant Parikh

- Strategic games in normal form
- Equilibrium concepts
- Games in extensive form
- Signalling games
- **Application: Resolving Ambiguities**

The Standard Example

a) Every ten minutes a man gets mugged in New York. (A)
b) Every ten minutes some man or other gets mugged in New York. (F)
c) Every ten minutes a particular man gets mugged in New York. (F’)

- How to read the quantifiers in a)?

Abbreviations

- φ: Meaning of `every ten minutes some man or other gets mugged in New York.'
- φ': Meaning of `Every ten minutes a particular man gets mugged in New York.'
- θ_1: State where the speaker knows that φ.
- θ_2: State where the speaker knows that φ'.

A Representation

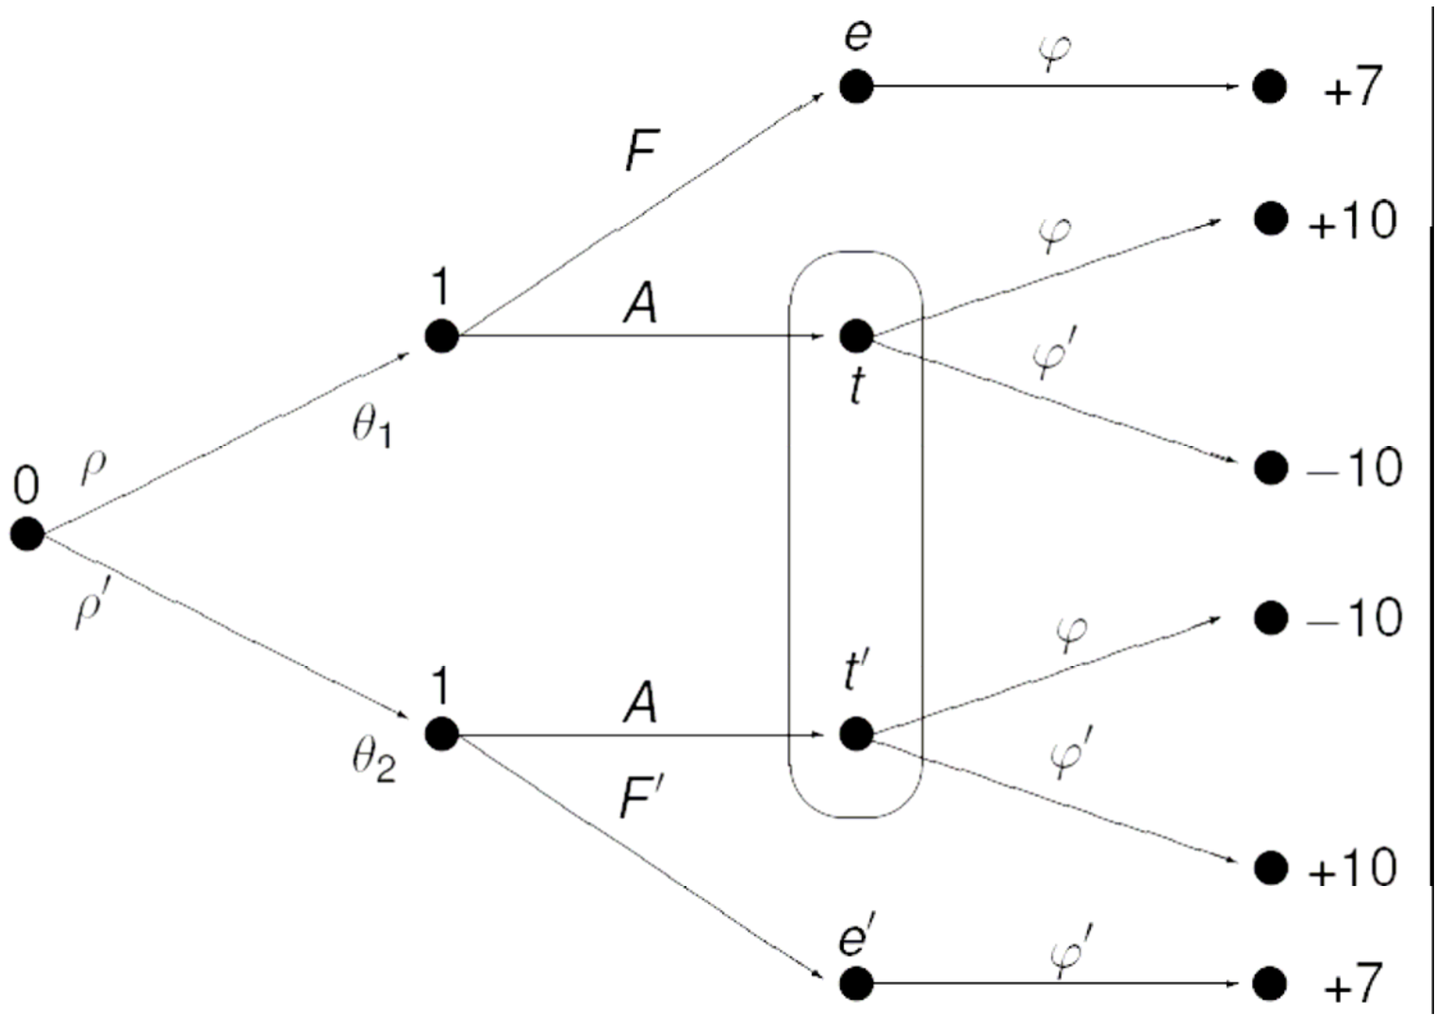

The Strategies

	θ_1	θ_2
S	A	A
S'	A	F'
S''	F	A
S'''	F	F'

	A	F	F'
H	φ	φ	φ'
H'	φ'	φ	φ'

The Strategies

Speaker: $\mathcal{S}_1 = \{S, S', S'', S'''\}$

Hearer: $\mathcal{S}_2 = \{H, H'\}$

The Payoffs

θ_1	H	H'
S	10	-10
S'	10	-10
S''	7	7
S'''	7	7

θ_2	H	H'
S	-10	10
S'	7	7
S''	-10	10
S'''	7	7

The Payoffs

Left: In situation θ_1

Right: In situation θ_2

Expected Payoffs

	H	H'
S	8	-8
S'	9.7	-8.3
S''	5.3	7.3
S'''	7	7

The Expected Payoffs

Probability of θ_1: $\rho = 0.9$

Probability of θ_2: $\rho' = 0.1$

Analysis

- There are two Nash equilibria

  (S',H) and (S'',H')
- The first one is also a Pareto Nash equilibrium.
- With (S',H) the utterance (A) should be interpreted as meaning (F):

(A) Every ten minutes a man gets mugged in New York.

(F) Every ten minutes some man or other gets mugged in New York.